Digital Discovery

ROYAL SOCIETY OF CHEMISTRY

PERSPECTIVE

Published on 16 October 2026
Licensed under CC-BY 4.0

Check for updates

Cite this: DOI: 10.1039/d6dd00494f

Quantum multiscale modeling: a hierarchy of algorithms for complex chemical systems

Seenivasan Hariharan, *[ab] Kareljan Schoutens, [ab] Sachin Kinge[c] and Lucas Visscher [d]

Multiscale modeling of complex chemical systems requires algorithms that operate coherently across electronic, atomistic, mesoscopic, and continuum scales. While quantum algorithms have been proposed for each regime, no systematic framework exists to compose them across scale boundaries. Here, we identify the conditions under which fault-tolerant quantum algorithms might preserve scale-specific quantum advantages. We map quantum phase estimation, Hamiltonian simulation with Gibbs state preparation, quantum random walks, and quantum partial differential equation solvers onto electronic structure, molecular dynamics, mesoscopic kinetics, and continuum reactor physics, respectively. Crucially, these correspondences do not imply unconditional end-to-end quantum advantage; speedups depend heavily on state preparation, memory architectures, matrix conditioning, and classical readout costs. Six unresolved questions define this composition problem, illustrated *via* a quantum hierarchy for CO oxidation over Pt(111). We propose viewing inter-scale transfer as a quantum channel composition problem at the interface of algorithm design and non-equilibrium statistical mechanics, and ask whether information loss at scale boundaries is intrinsic to multiscale modeling or merely a consequence of lossy classical transduction between algorithmic layers. The resulting roadmap suggests that multiscale quantum advantage is governed primarily by the structure of information transfer between algorithmic layers, rather than by performance at individual scales alone.

Received 13th July 2026
Accepted 14th October 2026

DOI: 10.1039/d6dd00494f

rsc.li/digitaldiscovery

Introduction

A defining challenge in computational chemistry is that the physical phenomena governing the performance of a catalyst,[1] a drug molecule,[2] or a functional material[3] unfold simultaneously across multiple scales that differ by many orders of magnitude in both length and time.[4] Consider heterogeneous catalysis—the enabling chemistry behind ammonia synthesis for fertilizer production, automotive emission control, hydrogen production and utilization, carbon dioxide conversion, fuel cells, petrochemical refining, and the synthesis of sustainable fuels and commodity chemicals—as the canonical case.[5] At the electronic scale (Scale I; ∼1 Å, ∼1 fs), quantum mechanical interactions between adsorbate electrons and surface d-bands determine whether a molecule binds, dissociates, or passes unreacted. At the atomistic scale (Scale II; ∼10 nm, ∼1 ns), thermal fluctuations drive surface diffusion, reconstruct the active site, and set the pre-exponential factors (k_0) that appear in Arrhenius rate expressions. At the mesoscopic scale (Scale III; ∼1 μm, ∼1 ms), the interplay of adsorption, desorption, surface reaction, and lateral interactions governs steady-state coverages and the turnover frequency that experiments measure. At the continuum scale (Scale IV; ∼1 cm, ∼1 s), species and energy transport couple with surface chemistry to determine conversion, selectivity, and *operando* stability.[1] No single computational method resolves all four scales in heterogeneous catalysis modeling; multiscale modeling is, fundamentally, the science of crossing the boundaries between them. While hierarchical multiscale modeling links distinct physical scales through sequential information transfer, hybrid and embedding schemes (such as quantum mechanics/molecular mechanics (QM/MM) and quantum–classical dynamics) partition a single scale, treating important regions with higher-fidelity representations. This perspective has been articulated as a general computational paradigm in multiscale modeling and simulation, where the focus lies on the systematic representation and transformation of information across scales.[6] In this work, we extend this viewpoint by evaluating whether this paradigm can be formulated as a formal framework for multiscale modeling based on fault-tolerant quantum algorithms. Unlike prior multiscale quantum chemistry frameworks that focus primarily on embedding quantum resources within localized active spaces or isolated algorithmic layers,[7–9] our Perspective

[a]Institute for Theoretical Physics, University of Amsterdam, Amsterdam, The Netherlands. E-mail: hseeni@gmail.com
[b]QuSoft, CWI, Amsterdam, The Netherlands
[c]Toyota Motor Europe, Materials Engineering Division, Zaventem, Belgium
[d]Department of Chemistry and Pharmaceutical Sciences, Vrije Universiteit, Amsterdam, The Netherlands

© 2026 The Author(s). Published by the Royal Society of Chemistry

treats cross-scale information transfer itself as the central computational object.

The classical hierarchy for this problem is well-established. Density functional theory (DFT) supplies detailed structural information, static adsorption energies and activation barriers, while coupled-cluster methods can provide a high-accuracy refinement of the energy for small systems.[10] To capture finite-temperature effects and structural fluctuations, *ab initio* molecular dynamics (AIMD) can be used to sample DFT potential energy surfaces directly by generating nuclear trajectories.[11] However, these explicit vibrational effects and temperature-dependent anharmonicities are rarely transferred intact to upper scales. Instead, kinetic Monte Carlo (kMC) typically evolves the chemical reaction network (CRN) stochastically using rigid, harmonic rate constants derived from lower-scale calculations to produce coverages and turnover frequencies.[12,13] Reactor-scale partial differential equations (PDEs) close the hierarchy by coupling surface source terms with bulk transport.[14] Each level is computationally demanding in its own right: electronic structure methods range from DFT, scaling formally as $O(N^3)$ and often insufficient for strongly correlated adsorbates,[15] to wavefunction-based approaches with steep polynomial or even exponential scaling that limits their application to small molecular clusters.[16] At the atomistic scale, molecular dynamics requires long-time sampling of nuclear trajectories to capture rare events, which routinely bounds timescales to the nanosecond or microsecond regime.[17] Moving to the mesoscale, kMC simulations become highly intractable when rate-constant hierarchies span many orders of magnitude. In such cases, the simulation spends most of its time sampling rapid, reversible transitions that contribute little to the overall reaction dynamics.[18] Finally, macroscopic reactor-scale PDE discretizations introduce severe numerical challenges, frequently involving 10^6–10^8 degrees of freedom to capture complex, multi-dimensional fluid transport and chemical species fields.[19] While the computational bottlenecks differ across scales, every transition necessarily involves approximating the information transferred between models.

It must be emphasized that the classical baseline is itself a moving target. Machine-learned interatomic potentials now reproduce *ab initio* forces at a small fraction of the cost of explicit DFT evaluations,[20,21] effectively collapsing the cost barrier between Scales I and II, and machine-learned surrogates increasingly accelerate the construction and solution of kinetic models at Scale III.[22,23] These approaches, however, inherit the accuracy ceiling of the electronic structure data they are trained on and offer no systematic route to the strongly correlated regimes where DFT itself fails.[24] The case for quantum algorithms in the multiscale hierarchy therefore rests not only on systematically improvable accuracy at Scale I, but equally on preserving that additional information as it propagates across successive scales.

What is rarely addressed is a structural question about the entire hierarchy: when an algorithm at one scale passes information to an algorithm at the next, what form does this information take, and should this mapping change fundamentally when both algorithms are quantum? Quantum algorithms have been proposed to accelerate individual scales of the multiscale hierarchy,[25] from using quantum phase estimation (QPE) to obtain high-accuracy electronic structure data,[26–28] to near-optimal Hamiltonian simulation for atomistic dynamics.[29,30] At the mesoscale, multidimensional quantum walks on mass-action system graphs provide a quantum analogue of kMC that leverages reaction network structure for Markov-chain speedups,[31,32] while quantum singular value transformation (QSVT_Backspace)-based linear solvers offer exponential improvements in system size for reactor-scale PDEs.[33–35] However, while each proposal is individually well-grounded, every classical algorithm in the hierarchy is merely replaced by a quantum counterpart. The critical scientific question is not whether these individual analogues exist, but how to compose them coherently across scale boundaries, shifting the objective from isolated speedups to the integration of quantum-accelerated routines into unified simulation pipelines.[36]

Even if every individual quantum subroutine achieves its intended quantum advantage, that advantage can only be realized end-to-end if the information exchanged across successive scales preserves the quantities that matter. Rather than proposing a new quantum algorithm for a single scale, this Perspective establishes inter-scale composition as the central question in quantum multiscale modeling. Specifically, we define the scale boundary $\mathcal{I}_{k\to k+1}$ as an explicit information channel between quantum subroutines, identify six open questions governing the mathematical structure and uncertainty propagation of these channels, and map this hierarchy to a concrete case study of CO oxidation over Pt(111). Ultimately, we argue that preserving end-to-end quantum advantage depends critically on the composition of these inter-scale boundaries, rather than on isolated accelerations within individual algorithmic layers. Identifying these hierarchical bottlenecks precisely, rather than focusing on modular algorithm substitutions, constitutes the central contribution of this work.

The classical–quantum algorithm correspondence

Table 1 presents the central idea of this Perspective. The four rows map each physical scale onto its canonical classical method and its primary quantum analogue, together with the computational complexity of the quantum analogue and, most importantly, the boundary information passed to the next coarser scale. What this mapping makes clear in the rightmost column is that the information passed upward at each boundary is, in the classical hierarchy, a lossy, classical observable—such as an energy, a rate constant, or a coverage. Our core thesis is that this classical bottleneck fundamentally can limit end-to-end quantum advantage, motivating the investigation of coherent quantum information channels where the downstream task permits.

In conventional multiscale workflows, the output of one computational scale is typically converted into classical observables before being passed to the next scale. This measurement-mediated interface is entirely appropriate when

© 2026 The Author(s). Published by the Royal Society of Chemistry

Table 1 Fault-tolerant quantum multiscale modeling hierarchy and classical analogues. Complexities reflect the dominant bottleneck sub-task; N = basis functions, $\|H\|$ = Hamiltonian norm, ε = target precision, and n = system dimension. The information rows define the proposed scale-boundary channels that are the focus of this Perspective[a]

Approach	Algorithmic framework, complexity, & channel information
SCALE I: electronic (∼1 Å/1 fs)	
Classical	Baseline: Self-consistent field (SCF) and correlated electronic-structure calculations
Quantum	QPE on qubitized Hamiltonian: $\mathcal{O}(\|H\|/\varepsilon)$ queries to the block-encoding. Avoids representing of the exponentially large FCI vector; advantage conditional on preparing an initial state with non-vanishing ground-state overlap *Challenge*: high state-preparation overhead
Information	Extracts ground-state adsorption energies ΔE_{ads}, activation barriers E_{a}, and pre-exponential factor k_0; parameterized into Arrhenius expressions for Scale II/III
SCALE II: atomistic (∼10 nm/1 ns)	
Classical	Baseline: *ab initio* MD (AIMD) and canonical ensemble (*NVT*) thermal sampling
Quantum	Fault-tolerant Gibbs state preparation followed by coherent Hamiltonian simulation: $\mathcal{O}(t\cdot\|H\|\cdot\mathrm{poly}\log(1/\varepsilon))$ queries for real-time propagation of thermal observables; polynomial state preparation speedup *Challenge*: Gibbs state-preparation overhead
Information	Extracts temperature-dependent free-energy diffusion barriers $\Delta G^{\ddagger}_{\mathrm{diff}}(T)$ and lateral interaction parameters ω_{ij}; compiled as a coverage-dependent rate matrix
SCALE III: mesoscopic (∼1μm/1 ms)	
Classical	Baseline: stochastic kMC and non-linear mass-action kinetics
Quantum	Multidimensional QRW on mass-action system graphs *via* circuit theory: provable query speedups *via* QRAM to decide species reachability, sample reachable states, and approximate steady-state fluxes under network perturbations *Challenge*: dependence on QRAM and CRN topology sensitivity
Information	Extracts steady-state surface coverages θ_i^*, turnover frequencies (TOF), and selectivities; injected as local source terms $\dot{\omega}_i$ into Scale IV.
SCALE IV: continuum (∼1 cm/1 s)	
Classical	Baseline: Computational Fluid Dynamics (CFD) and macroscopic reactor PDEs
Quantum	QSVT-based sparse linear PDE solvers and quantum ODE/PDE integrators: $\mathcal{O}(\mathrm{poly}\log(n)\cdot\kappa\cdot\log(1/\varepsilon))$ gate complexity for accessing global observables of suitably structured sparse systems, compared with $\mathcal{O}(n)$ classical scaling *Challenge*: speedup conditional on condition number κ and efficient observable readout
Information	Resolves macroscale reactor concentration profiles $c_i(\mathbf{x},t)$, temperature fields $T(\mathbf{x},t)$, and integrated terminal yields/selectivities

[a] Abbreviations: DFT = density functional theory; CCSD(T) = coupled-cluster singles and doubles with perturbative triples; FCI = full configuration interaction; AIMD = *ab initio* molecular dynamics; kMC = kinetic Monte Carlo; CRN = chemical reaction network; QPE = quantum phase estimation; QRW = quantum random walk; QRAM = quantum random access memory; QSVT = quantum singular value transformation; PDE = partial differential equation; NVT = the canonical ensemble, keeps the number of particles (N), volume (V), and temperature (T) constant.

the downstream model requires only those observables. However, measurement can also discard information contained in the quantum representation of the intermediate calculation, including amplitudes, phases, correlations, or entanglement. When the downstream computation can operate directly on this representation, a quantum-native composition of adjacent scales may avoid the intermediate conversion

$$\text{Quantum state} \rightarrow \text{classical data} \rightarrow \text{new quantum state.}$$

The potential benefit of such coherent composition is therefore not simply the elimination of state-preparation cost. Rather, it is the possibility of preserving and processing information across computational scales without first reducing the intermediate result to a prescribed set of classical observables. Whether this provides a computational advantage is problem dependent and requires comparison with the costs of measurement, classical processing or reconstruction, and subsequent state preparation. Thus, intermediate measurement should not be regarded as intrinsically disadvantageous; instead, the appropriate interface depends on which information is required by the subsequent scale and on the relative cost of classical and coherent composition.

Classical computational scaling across the multiscale hierarchy

The classical computational cost of the four scales is not determined solely by the formal scaling of an individual computational operation. At each scale, the total cost depends on both the cost of an elementary computational step and the number of steps, events, statistically independent samples, or solver iterations required to achieve the target physical accuracy. We therefore characterize the classical baseline in terms of these two contributions, as summarized in Table 2.

The resulting expressions should be interpreted as representative cost models rather than universal asymptotic complexities. At Scale I, formal polynomial scaling provides a useful characterization of the dominant electronic-structure cost. For example, conventional pure and hybrid density-functional calculations exhibit characteristic cubic and

© 2026 The Author(s). Published by the Royal Society of Chemistry

Table 2 Classical computational scaling across the four scales. The classical cost at each scale is determined by both the cost of an elementary computational step and the number of steps, events, samples, or solver iterations required to achieve the desired physical accuracy. The expressions shown represent characteristic cost models rather than universal asymptotic complexities. Here, MP2 denotes second-order Møller–Plesset perturbation theory, N denotes the system size, N_{traj} the number of independent trajectories, N_{steps} the number of integration steps, R the number of reaction channels, S the number of kinetic states or species, N_{events} the number of reaction events, N_{t} the number of time steps, N_{iter} the number of solver iterations, N_{DOF} the number of discretized degrees of freedom and $C_{(\cdot)}$ the cost of the corresponding elementary operation

Scale and classical calculation	Representative classical cost
I: Electronic SCF and correlated electronic-structure calculations	$C_{\mathrm{ES}}(N)$; *e.g.*, $O(N^3)$ for conventional pure DFT $O(N^4)$ for hybrid DFT $O(N^5)$ for MP2 $O(N^7)$ for CCSD(T)
II: Atomistic *Ab initio* MD (AIMD) and canonical ensemble (NVT) thermal sampling	$N_{\mathrm{traj}}N_{\mathrm{steps}}C_{\mathrm{force}}(N)$, with $C_{\mathrm{force}}(N) = O(N^3)$ for DFT-based AIMD; for classical force fields, $O(N^2)$ for naive pairwise interactions, approaching $O(N)$ for short-range interactions and approximately $O(N \log N)$ for particle-mesh Ewald methods
III: Mesoscopic Stochastic kMC and non-linear mass-action kinetics	$N_{\mathrm{events}}C_{\mathrm{event}}(R)$ for kMC, or $N_t C_{\mathrm{RHS}}(S, R)N_{\mathrm{iter}}$ for time-dependent kinetic equations
IV: Continuum Computational Fluid Dynamics (CFD) and macroscopic reactor PDEs	$N_{\mathrm{DOF}}N_{\mathrm{iter}}C_{\mathrm{operator}}$, where the scaling depends on the discretization and numerical solver

quartic scaling, respectively, while correlated methods such as (multireference) perturbation theory and coupled cluster approaches exhibit substantially steeper scaling with system size.[37] In practice, these formal scalings are however often substantially reduced using density fitting or Cholesky decomposition,[38] local correlation methods[39] and tensor compression techniques,[40] which make the actual scaling less transparent and to some extent system-specific. At the subsequent scales, the picture is still less clear. The computational bottleneck then depends increasingly on the number of trajectories, events, samples, or numerical iterations required to obtain a converged physical observable. Consequently, the nominal cost of an individual operation does not by itself characterize the total computational burden.

For Scale II, for example, the cost of molecular dynamics can be written schematically as

$$C_{\mathrm{MD}} \sim N_{\mathrm{traj}}N_{\mathrm{steps}}C_{\mathrm{force}}(N). \tag{1}$$

The force-evaluation cost may scale as $O(N^2)$ for naive pairwise interactions, approach $O(N)$ for short-range interactions with appropriate neighbor-list methods, or approximately $O(N \log N)$ for particle-mesh treatments of long-range electrostatics[41] (for AIMD, C_{force} is instead dominated by the $O(N^3)$ cost of each DFT evaluation). However, this per-step scaling does not capture the full cost of obtaining statistically converged thermodynamic or free-energy quantities. The number of statistically independent configurations is reduced by temporal correlations, with the statistical inefficiency determined by the integrated autocorrelation time. Consequently, for a target statistical error ε, the total sampling cost can be expressed schematically as

$$C_{\mathrm{MD}} = O\left[N_{\mathrm{traj}}\frac{\tau_{\mathrm{int}}}{\Delta t}\varepsilon^{-2}C_{\mathrm{force}}(N)\right], \tag{2}$$

where τ_{int} is the integrated autocorrelation time and Δt is the integration timestep. Thus, although force evaluation may have relatively low polynomial scaling in N, the computational cost of converging a free-energy or rare-event observable can instead be dominated by the long-time sampling required to decorrelate configurations and adequately explore the relevant configuration space.

At Scale III, the classical cost is associated with the propagation of a kinetic network over the relevant physical time and with the resolution of the reaction channels that contribute to the dynamics. For kinetic Monte Carlo (kMC), a representative cost is

$$C_{\mathrm{kMC}} \sim N_{\mathrm{events}}C_{\mathrm{event}}(R), \tag{3}$$

where N_{events} is the number of reaction events required to reach the target physical time[42] and $C_{\mathrm{event}}(R)$ denotes the cost of constructing or sampling from the set of R reaction channels. Importantly, N_{events} is not determined solely by the number of reactions in the network: it also depends on the reaction propensities and on the separation between fast and slow processes. Large disparities in reaction timescales can therefore lead to a large number of microscopic events being required to resolve a much slower chemical process.[12] For deterministic kinetic descriptions, the corresponding cost can be expressed schematically as

$$C_{\mathrm{ODE}} \sim N_{\mathrm{t}}C_{\mathrm{RHS}}(S,R)N_{\mathrm{iter}}, \tag{4}$$

where N_{t} is the number of integration steps, S is the number of kinetic states or species, and N_{iter} accounts for the iterations required by the numerical integration scheme. Thus, for Scale III, network size, kinetic stiffness, and the physical timescale that must be resolved can all contribute to the classical computational cost.

© 2026 The Author(s). Published by the Royal Society of Chemistry

At Scale IV, the classical cost is determined primarily by the discretization of the reactor equations and the numerical solution of the resulting linear or nonlinear system. A representative cost model is

$$C_{\text{reactor}} \sim N_{\text{DOF}} N_{\text{iter}} C_{\text{operator}}, \tag{5}$$

where N_{DOF} is the number of spatial and, where applicable, temporal degrees of freedom, N_{iter} is the number of iterations required for convergence, and C_{operator} is the cost of applying the discretized differential operator or solving the associated linear system. The resulting scaling depends strongly on the spatial discretization, boundary conditions, stiffness of the coupled equations, and choice of numerical solver.[43,44] Direct sparse solvers, Krylov methods, and multigrid methods can therefore exhibit substantially different scaling with N_{DOF}. Consequently, a single universal polynomial complexity is not representative of the continuum scale; the dominant classical cost is instead determined by the size and conditioning of the discretized problem together with the convergence requirements of the solver.

Taken together, these scaling considerations highlight that the classical cost of multiscale modelling is governed by different computational bottlenecks at different levels of the hierarchy. Electronic-structure calculations are often dominated by the scaling of the underlying many-electron method, whereas atomistic simulations can be limited by long-time sampling, mesoscopic kinetics by the number of reaction events and separation of timescales, and continuum models by the size and conditioning of the discretized equations. This distinction is important when assessing potential quantum advantage: an improvement in the asymptotic scaling of an individual subroutine does not necessarily translate into an end-to-end multiscale advantage unless the dominant sampling, state-preparation, conditioning, communication, and measurement costs are also addressed.

A quantum multiscale hierarchy for CO oxidation over Pt(111)

To illustrate how the proposed hierarchy appears in practice, we consider catalytic CO oxidation over a close-packed platinum surface (Pt(111)),

$$2\text{CO(g)} + \text{O}_2\text{(g)} \xrightarrow{\text{Pt(111)}} 2\text{CO}_2\text{(g)} \tag{6}$$

which proceeds predominantly through a Langmuir–Hinshelwood mechanism involving adsorption, oxygen dissociation, surface diffusion, and the subsequent reaction of adsorbed intermediates. In its minimal representation, the elementary reaction network consists of

$$\text{CO(g)} + * \rightleftharpoons \text{CO}*, \tag{7}$$

$$\text{O}_2\text{(g)} + 2* \rightleftharpoons 2\text{O}*, \tag{8}$$

$$\text{CO}* + \text{O}* \rightarrow \text{CO}_2\text{(g)} + 2*. \tag{9}$$

where * denotes an empty surface site. This reaction represents a prototypical benchmark in heterogeneous catalysis whose mechanism, surface intermediates, and non-equilibrium kinetic behavior have been studied extensively through classical electronic-structure calculations, kinetic modeling, and *operando* experiments.[1,13]

This minimal network provides a concrete example of how observables extracted at one scale become inputs to subsequent scales. Furthermore, this workflow could be scaled up to handle the exact structural complexities, such as single-atom or magnetic active sites, identified by recent industry-academic roadmaps for industrial catalyst screening.[25] In particular, we can formalize the information transfer across the Scale I–IV hierarchy as the systematic compilation of a multiscale parameter vector:

$$\lambda = \left\{ \Delta E_{\text{ads}},\ E_{\text{a}},\ \Delta G^{\ddagger}_{\text{diff}}(T),\ \omega_{ij},\ k_j,\ \theta_i^{*},\ \text{TOF} \right\} \tag{10}$$

where these components correspond respectively to adsorption energies, activation barriers, diffusion barriers, lateral interaction parameters, rate constants, steady-state surface coverages, and turnover frequencies. The explicit purpose of this case study is to show precisely how physical information crosses each scale boundary, what specific chemical observables are transferred, and where the structural composition questions of quantum multiscale modeling manifest (Section Open questions in quantum multiscale modeling).

Scale I – Electronic structure of CO/Pt(111)

The initial task is determining the ground-state adsorption energies, activation barriers, and transition-state geometries for elementary events including CO adsorption, O_2 dissociation, and CO_2 formation. Classically, these quantities are obtained from DFT or correlated wavefunction methods. The drawback is well-documented: approximate exchange-correlation functionals introduce severe localized errors on transition-metal surfaces (*e.g.*, the infamous DFT puzzle of incorrect CO site-preference on Pt(111)),[15,45] while systematically improvable classical methods scale exponentially.

In the fault-tolerant hierarchy, QPE provides a systematically improvable alternative for high-accuracy electronic structure, evaluating eigenvalues with a query complexity scaling as $\mathcal{O}(\|H\|/\varepsilon)$ and potentially avoiding explicit classical representation of exponentially large wavefunction vectors.[26,27] To formalize this interface mathematically for the CO/Pt(111) system, we consider a position-dependent formulation where the electronic Hamiltonian is block-conditional on a nuclear configuration register:[46,47]

$$\hat{H}_{\text{elec}} = \sum_i |i\rangle\langle i| \otimes \hat{H}_i \tag{11}$$

where $|i\rangle$ indexes a discrete grid of nuclear configurations corresponding to local adsorption geometries on the platinum slab, and $\hat{H}_i$ represents the corresponding electronic Hamiltonian at that geometry. This expression provides an idealized coherent description of a discretized Born–Oppenheimer energy landscape, in which the nuclear degrees of freedom are

© 2026 The Author(s). Published by the Royal Society of Chemistry

represented quantum mechanically in a discretized configuration basis $\{|i\rangle\}$, while the electronic Hamiltonian is parameterized by each nuclear configuration. By preparing a spatial superposition over these surface configurations and executing conditional QPE to produce a coherent electronic energy register,[48] the algorithm outputs an entangled state vector across the joint nuclear, electronic, and energy registers:

$$|\psi_{\mathrm{I}}\rangle = \sum_i c_i |i\rangle |\phi_i\rangle \left|\Delta E_{\mathrm{ads}}^{(i)}\right\rangle \tag{12}$$

where $|\phi_i\rangle$ represents the electronic state and the energy register stores the precise adsorption energy eigenvalues $\Delta E_{\mathrm{ads}}^{(i)}$ corresponding to each specific configuration $|i\rangle$. Thus, Scale I does not assume a classical sampling of fixed molecular geometries; the nuclear configuration register remains in coherent superposition until the Scale I → II transformation.

In principle, the coherent energy register produced by QPE could serve directly as an input to quantum sampling procedures rather than being measured immediately. Rather than collapsing this state projectively *via* intermediate classical measurement–which can reintroduce severe input/output bottlenecks–this entangled register is instead used coherently as the control register for a Scale II Gibbs reweighting procedure U_{Gibbs}.[49] By appending an ancilla qubit and executing a coherent conditional rotation controlled on the energy register,

$$|i\rangle|\phi_i\rangle\left|\Delta E_{\mathrm{ads}}^{(i)}\right\rangle|0\rangle \mapsto |i\rangle|\phi_i\rangle\left|\Delta E_{\mathrm{ads}}^{(i)}\right\rangle\left(\mathrm{e}^{-\beta\Delta\tilde{E}^{(i)}/2}|0\rangle + \sqrt{1-\mathrm{e}^{-\beta\Delta\tilde{E}^{(i)}}}\,|1\rangle\right) \tag{13}$$

with $\Delta\tilde{E}(i) = \Delta E_{\mathrm{ads}}^{(i)} - E_{\min}$ shifted by the minimal register energy to keep the rotation well-defined. Post-selecting the ancilla on $|0\rangle$ and tracing out the electronic and energy registers yields the mixed density operator over the nuclear configuration space:

$$\rho_{\mathrm{II}} = \frac{1}{\tilde{Z}}\sum_i |c_i|^2 \mathrm{e}^{-\beta\Delta\tilde{E}^{(i)}}|i\rangle\langle i| = \frac{1}{\tilde{Z}'}\sum_i |c_i|^2 \mathrm{e}^{-\beta\Delta E_{\mathrm{ads}}^{(i)}}|i\rangle\langle i|, \tag{14}$$

where $\tilde{Z} = \sum_i |c_i|^2 \mathrm{e}^{-\beta\Delta\tilde{E}^{(i)}}$, $\tilde{Z}' = \sum_i |c_i|^2 \mathrm{e}^{-\beta\Delta E_{\mathrm{ads}}^{(i)}}$, and $\beta = (k_{\mathrm{B}}T)^{-1}$. The constant energy shift $E_{\min}$ factors out and cancels exactly under normalization, leaving the relative thermal weights invariant. The Scale I amplitudes c_i do not disappear: ρ_{II} reduces to the exact Gibbs state only when the input superposition is uniform ($|c_i|^2 = 1/N$); otherwise the thermal weights remain modulated by the upstream preparation. Since ρ_{II} is diagonal (for non-degenerate, resolved energies), the coherent route transfers a clasical Boltzmann distribution; its benefit is avoiding intermediate sampling an re-enc. Moreover, the post-selection succeeds with probability $P_{\mathrm{succ}} = \sum_i |c_i|^2 \mathrm{e}^{-\beta\Delta E^{(i)}}$, which, in unfavorable cases, can become exponentially small in $\beta\cdot\mathrm{spread}(\Delta E_{\mathrm{ads}})$, requiring $\mathcal{O}(1/\sqrt{P_{\mathrm{succ}}})$ rounds of amplitude amplification. This construction therefore bypasses intermediate classical measurement and re-encoding of the energy information, but at a quantifiable boundary cost–precisely the query-complexity trade-off that Question Q1 identifies as the central object of the coherent channel $\mathcal{I}_{\mathrm{I}\to\mathrm{II}}$.

Scale II – Adsorbate dynamics on Pt(111)

At the atomistic scale, the objective shifts from isolated ground-state energies to computing the temperature-dependent parameters that govern the mobility and spatial distribution of species across the Pt(111) surface. Specifically, the model requires the free-energy diffusion barriers $\Delta G_{\mathrm{diff}}^{\ddagger}(T)$ for CO* hopping between adjacent top and bridge sites, and the lateral interaction parameters ω_{ij} that quantify how neighboring CO* or O* spectator species alter local adsorption and reaction energetics.

Rather than running trillions of classical force evaluations to log classical trajectories, the fault-tolerant quantum framework uses the density matrix passed from Scale I, ρ_{II}, as the initial quantum representation from which a Gibbs state associated with the interacting multi-adsorbate Hamiltonian is prepared *via* block-encoding methods. By configuring this state to represent co-adsorbed CO* and O* at varying coordinate distances, the quantum computer deploys quantum algorithms for free-energy estimation to evaluate coverage-dependent free-energy differences and interaction energies. Coherent, real-time Hamiltonian simulation is subsequently used to sample the finite-temperature dynamics of the interacting adsorbate system.

From this simulation pipeline, the target observables are extracted *via* quantum expectation values, yielding the temperature-dependent free-energy diffusion barriers $\Delta G_{\mathrm{diff}}^{\ddagger}(T)$ and a matrix of lateral interaction coefficients ω_{ij}. This directly raises the first Question Q1: if the structural relaxation of the underlying Pt(111) lattice happens on a timescale comparable to the adsorbate diffusion hops, the coarse-graining step injects memory effects, requiring $\mathcal{I}_{\mathrm{II}\to\mathrm{III}}$ to be formalized as a non-Markovian quantum channel rather than a memoryless rate assignment.

Scale III – CO oxidation kinetics

The classical mesoscopic description accounts for spatial heterogeneity and correlations across the catalyst surface by evolving the master equation over the combinatorial state space of the Pt(111) lattice network. Here, the rate matrix $\mathbf{W}$ is compiled using the parameters passed from the previous scale, where local event frequencies are updated dynamically based on coverage:

$$k_\alpha = \frac{k_{\mathrm{B}}T}{h}\exp\left[-\beta\Delta G_\alpha^{\ddagger}(T,\boldsymbol{\theta})\right] \tag{15}$$

with the total coverage-dependent activation free energy formalized as

$$\Delta G_\alpha^{\ddagger}(T,\boldsymbol{\theta}) = \Delta G_{\alpha,0}^{\ddagger}(T) + \sum_l \omega_{\alpha l}\theta_l. \tag{16}$$

Here, $\Delta G_{\alpha,0}^{\ddagger}(T)$ is the intrinsic reference free-energy barrier of elementary event α, and $\omega_{\alpha l}$ denotes the energetic barrier correction arising from lateral interactions with neighboring species l. For catalytic CO oxidation, extreme mathematical stiffness arises because the rate constant for rapid CO* diffusion hops is several orders of magnitude larger than the rate

© 2026 The Author(s). Published by the Royal Society of Chemistry

constant for the rare, slow recombination step to form CO_2. This causes classical kinetic Monte Carlo (kMC) to expend a dominant fraction of computational effort simulating rapid, non-reactive diffusion back-and-forth steps, contributing substantially to the computational cost through high-frequency reversible events without significantly advancing the overall reaction coordinate.

In the quantum multiscale hierarchy, this stiffness bottleneck is addressed by representing the CRN as a multidimensional quantum random walk (QRW) mapped directly onto the mass-action system graph–rather than the full lattice configuration graph–using circuit theory equivalents and QRAM.[32] The walk operates on species-count states of the mass-action network, and the corresponding electrical network mappings can accelerate the computation of steady-state properties under favorable graph and spectral conditions. An important qualification applies here: the mass-action representation is a mean-field description in species numbers, whereas lattice kMC resolves explicit spatial correlations – including the coverage-dependent lateral interaction parameters from Scale II, which enter the kinetic model through the barrier corrections $\omega_{\alpha l}$. Recovering spatial resolution within the quantum walk framework requires either enlarging the network with site-resolved pseudo-species or developing new walk constructions on lattice configuration graphs. Quantifying what is lost in the mean-field contraction, and at what cost it can be restored, is part of the composition problem formalized below in Question Q5.

Rather than executing a full classical readout of the state space – or converging to a stationary distribution as a classical Markov chain would–the quantum walk encodes steady-state properties in the spectral structure of its walk operator for suitably constructed reversible or graph-embedded Markov processes, from which quantum phase and amplitude estimation extract two specific macro-kinetic observables: the steady-state surface coverages θ^*_{CO} and θ^*_O, and the net turnover frequency (TOF) of emitted CO_2 per active site. As highlighted in Question Q5, the core structural challenge is verifying that the sparse, anisotropic graph topology of the actual CO/O/Pt(111) surface network preserves the required QRW spectral gap when bound to these specific input and output parameters.

Scale IV – The CO oxidation reactor

At the macroscopic continuum scale, the catalytic surface is embedded into a physical chemical reactor model (*e.g.*, a stagnation-flow boundary layer reactor or a packed-bed channel over the Pt catalyst). Here, gas-phase transport and mass transfer limitations are coupled to the local surface kinetics *via* conservation partial differential equations (PDEs):

$$\frac{\partial c_i}{\partial t} + \nabla \cdot \mathbf{J}_i = \dot{\omega}_i\left(c_{CO}, c_{O_2}, T, \boldsymbol{\theta}^*\right) \tag{17}$$

where c_i represents the macroscopic gas-phase concentrations of CO, O_2, and CO_2. The coupling between the microscopic catalytic scales and the reactor model is introduced through the volumetric chemical source terms $\dot{\omega}_i = \nu_i \Gamma_s a_s \mathrm{TOF}$, where ν_i is

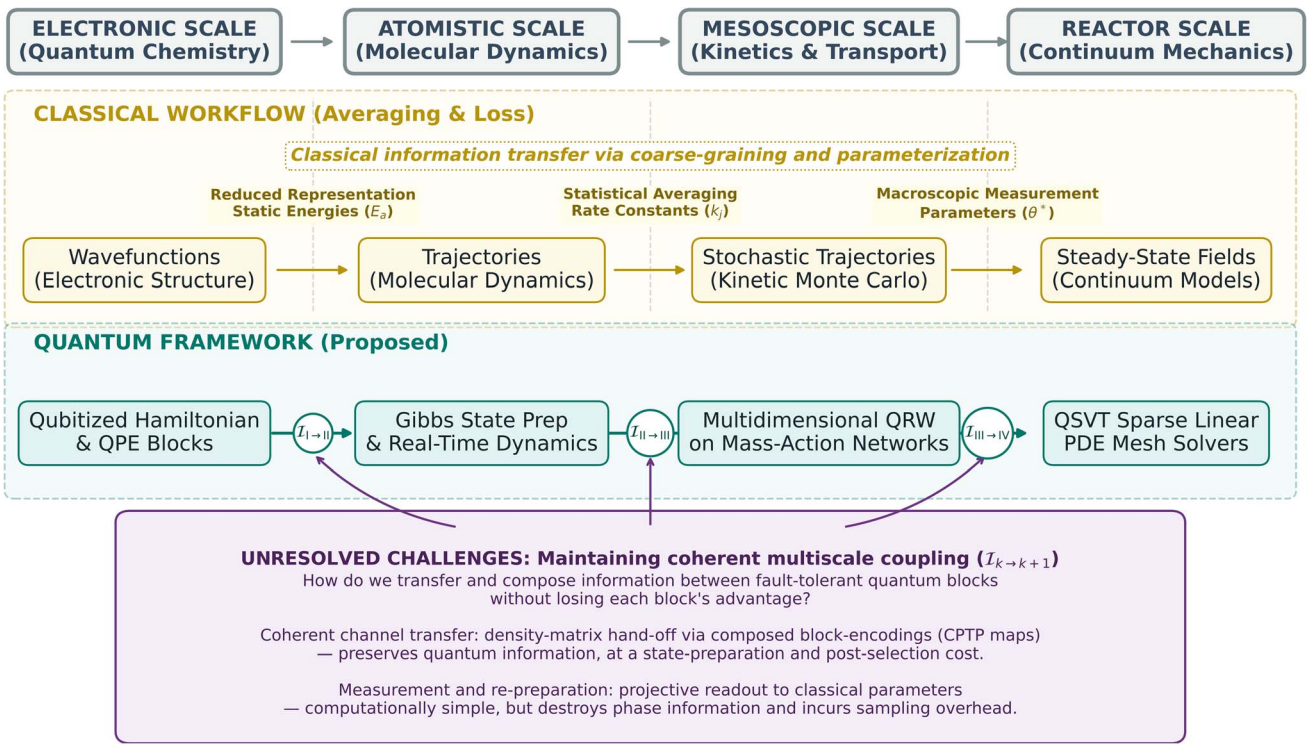


Fig. 1 Comparison of classical and fault-tolerant quantum workflows across electronic, atomistic, mesoscopic, and continuum scales. Classical approaches progressively compress information through averaging, reduced representations, and measurement. The proposed quantum hierarchy aims to compose coherent transformations across scales using fault-tolerant quantum algorithms. Interface nodes $\mathcal{I}_{k\to k+1}$ denote unresolved challenges in maintaining coherent multiscale coupling.

© 2026 The Author(s). Published by the Royal Society of Chemistry

the stoichiometric coefficient, Γ_s is the molar density of active sites per unit catalyst area, a_s represents the catalyst surface-area density, and TOF is the quantum-estimated surface turnover frequency obtained within the reduced kinetic representation of Scale III.

Under a hybrid classical–quantum Newton-QSVT approach, spatial discretization of the reactor geometry on a mesh of n nodes reduces each linearized implicit step of the nonlinear fluid dynamics solver to a sparse linear system $A\mathbf{x} = \mathbf{b}$ solved *via* the QSVT.[34,35] The true quantum advantage at this scale is shaped by Question Q6: if the engineering objective requires the full spatial visualization of the concentration field $c_i(\mathbf{x},t)$ throughout the entire reactor volume, the classical readout bottleneck can eliminate the practical advantage associated with the favorable polylogarithmic system-size dependence of the idealized QSVT linear solver.

Instead, the pipeline preserves its potential quantum advantage by avoiding full tomography, using quantum expectation measurements to extract a single global macroscopic engineering metric: the total integrated CO_2 production rate at the reactor outlet, defined by the boundary operator:

$$\dot{N}_{CO_2,out} = \int_{outlet} \mathbf{J}_{CO_2} d\mathbf{A} \tag{18}$$

By compressing the multiscale quantum simulation down to a single target scalar observable, the global pipeline completes its handoff while avoiding classical reconstruction of the high-dimensional intermediate representations generated at the electronic, atomistic, and mesoscopic scales.

The CO oxidation on Pt(111) case study reinforces that the central difficulty lies not in proposing individual quantum counterparts, but in composing them: resolving how cross-scale channels ($\mathcal{I}_{k\to k+1}$) handle information transfer without destroying the individual advantage of each computational layer. This architectural divergence – and the structural questions it raises at each interface – is mapped conceptually in Fig. 1.

Generalizing the hierarchy across chemical systems

The CO oxidation pipeline is one instance of a construction that applies to any chemical system whose physics spans the four scales of Table 1. In this section we abstract away from Pt(111) and state, for each scale, the general computational task, the fault-tolerant algorithm that addresses it, and the complexity and hardware conditions under which its localized advantage survives. The emphasis shifts from *what* information crosses each boundary, which is established above, to *at what cost* it does so for arbitrary reaction networks and reactor geometries.

Scale I – Electronic structure

At the electronic scale (Scale I), the central computational task is solving the many-electron Schrödinger equation for an active site embedded in a larger model such as a lattice with periodic boundary conditions. DFT handles this at polynomial cost (formally $\mathcal{O}(N^3)$), while wavefunction-based coupled-cluster theory restricts calculations to clusters of tens of atoms due to its steep $\mathcal{O}(N^7)$ scaling, and exact Full Configuration Interaction (FCI) scales exponentially with system size. Among these approaches, DFT remains the undisputed workhorse of heterogeneous catalysis modeling.

As illustrated for CO/Pt(111) above, approximate exchange-correlation functionals fail precisely for the strongly correlated, spin-polarized surfaces that catalyze industrially important reactions.[15,25] To break these classical limitations, fault-tolerant quantum algorithms have been proposed at this scale, with QPE positioned to compute precise electronic ground states and activation barriers. The query complexity of QPE on a qubitized Hamiltonian scales as $\mathcal{O}(\|H\|/\varepsilon)$, where $\|H\|$ is the spectral norm and ε the target precision.[26,27] But this does not constitute an unconditional exponential speedup. The total cost is governed by the overlap of the prepared initial state with the true ground state, and whether this overlap decays with system size for chemically relevant problems remains actively debated.[50] A landmark resource estimate by Reiher *et al.* showed that QPE could determine the catalytic mechanism of N_2 fixation by the FeMo-cofactor of the nitrogenase enzyme, a prototypically strongly correlated active site, using on the order of 10^2 logical qubits and $\sim 10^{15}$ T gates. Although daunting, these estimates define the regime at which quantum advantage becomes chemically relevant.[28] Subsequent algorithmic developments have reduced these requirements by several orders of magnitude; for example, spectrum-amplification estimates for FeMoco require ~ 1.1–1.5×10^3 logical qubits and $\sim 3 \times 10^8$ to 10^9 Toffoli gates (Table 3), making fault-tolerant quantum simulations of chemically relevant systems increasingly promising.[51,52]

The information transmitted through the Scale I → Scale II interface – ground-state adsorption energies ΔE_{ads} and activation barriers E_a – is conventionally condensed into classical Arrhenius- or Eyring-type rate expressions. In a fully fault-tolerant multiscale paradigm, the structural question becomes whether this information must be first extracted *via* classical measurement or can be utilized directly in Scale II.

The dissipative and coherent implementations of this interface were made explicit in the case study: projective measurement of the entangled register of eqn (12) collapses it to classical scalars, whereas the conditional Gibbs-state construction of eqn (14) transmits it coherently at a quantifiable post-selection cost.

Scale II – Molecular dynamics

At the atomistic scale, the focus shifts from static ground-state energies to the real-time evolution of nuclear positions and thermal sampling: how do adsorbates diffuse, how does the surface reconstruct under reaction conditions, and what are the coverage-dependent lateral interactions? Classical *ab initio* molecular dynamics (AIMD) resolves this but is strictly limited to picosecond-to-nanosecond trajectories on systems of hundreds of atoms by the $\mathcal{O}(N^3)$ cost of each DFT force evaluation.

© 2026 The Author(s). Published by the Royal Society of Chemistry

Table 3 Representative fault-tolerant resource estimates across quantum multiscale modelling. Reported resources are problem- and algorithm-dependent

Scale	Representative example and fault-tolerant (FT) resource status
I. Electronic	*FeMoco ground-state energy*:[52] 1.1 to 1.5 × 10^3 logical qubits and 10^8–10^9 Toffoli gates, with ∼3.2 to 4.5 × 10^6 physical qubits reported for the corresponding surface-code estimates
II. Atomistic	*Nuclear motion (30D)*:[59] <300 logical qubits and ∼3 months on a hypothetical 1 MHz FT processor. FT-MD resources additionally scale with trajectory length, sampling and repeated force evaluations
III. Mesoscopic	*Reaction networks*:[32,60] 4–5 qubit demonstrations. No end-to-end FT estimate is currently available for a chemically realistic CRN; resources depend on network size, oracle cost, mixing/sampling and precision
IV. Continuum	*Advection-diffusion*:[61,62] 14-qubit demonstration; 10^6 grid points require at least 20 data qubits for index encoding. Total FT resources depend on sparsity, condition number, block encoding and precision
Integrated	*Full multiscale workflow*: $C_{\mathrm{FT}}^{\mathrm{total}} \simeq \sum_{i=1}^{\mathrm{IV}} N_i C_i$, where C_i is the cost per quantum subroutine and N_i its number of invocations. The current ∼10^3-logical-qubit, 10^8–10^9-Toffoli estimates therefore represent the cost of an electronic-structure call rather than the complete multiscale workflow

To capture the finite-temperature behavior that defines true catalyst performance, a zero-temperature ($T = 0$) ground-state Hamiltonian is insufficient. Instead, the fault-tolerant analogue requires explicit quantum Gibbs state preparation ($\rho \propto \mathrm{e}^{-\beta H}$) to establish thermal equilibrium within the canonical ensemble (NVT). Once the thermal mixed state is prepared *via* block-encodings, coherent real-time Hamiltonian simulation samples and evolves observables associated with this thermal ensemble. Qubitization and linear-combination-of-unitaries (LCU) methods achieve a near-optimal query complexity of $\mathcal{O}(t\cdot\|H\|\cdot\mathrm{poly}\log(1/\varepsilon))$, yielding a near-linear time dependence and exponential precision tracking.[53]

Concretely, these fault-tolerant algorithms bridge the gap to higher scales by translating abstract wavefunctions into physical observables that govern long-timescale kinetics. First, rather than relying on classical functionals, the quantum computer evaluates analytical atomic forces $\mathbf{F}_I = -\nabla_{\mathbf{R}I}E(\mathbf{R})$ *via* quantum gradient estimation algorithms with polynomial scaling. Second, by preparing Gibbs states across varying reaction coordinates, the framework utilizes thermodynamic integration to directly compute Helmholtz free energy barriers:

$$\Delta A = -\frac{1}{\beta}\ln\left(\frac{Z_1}{Z_0}\right)$$

Third, real-time Hamiltonian simulation allows the direct extraction of dynamical correlation functions, such as the velocity autocorrelation function $\langle \mathbf{v}(t)\cdot\mathbf{v}(0)\rangle = \mathrm{Tr}(\rho\, \mathrm{e}^{iHt/\hbar}\mathbf{v}\, e^{-iHt/\hbar}\mathbf{v})$, yielding transport coefficients whose accuracy is limited only by the simulated Hamiltonian and the estimation precision, rather than by force-field or mean-field approximations.

The resulting temperature-dependent free-energy diffusion barriers $\Delta G_{\mathrm{diff}}^{\ddagger}(T)$ (obtained from free-energy differences ΔA; for adsorbed phases the $p\Delta V$ distinction between Helmholtz and Gibbs barriers is negligible) and lateral interactions ω_{ij} (mapped *via* cluster expansions of co-adsorbed thermal states) leave Scale II to compile the coverage-dependent rate matrix of Scale III. Here, local hopping rates are parameterized *via* transition state theory:

$$k_{i\to j} = \frac{k_{\mathrm{B}}T}{h}\exp\left(-\frac{\Delta A^{\ddagger}(\omega_{ij})}{k_{\mathrm{B}}T}\right)$$

If the underlying Born–Oppenheimer surface is itself evaluated *via* QPE, the full Scale I/II pipeline becomes a coherent quantum computation. Yet even in this best case, the boundary itself remains classical: the exponentially complex correlations sampled within Scale II are compressed into a handful of scalar rates $k_{i\to j}$ before crossing to Scale III. Whether this compression is a physical necessity of coarse-graining or an artifact of the classical kMC target and what a coherent alternative would cost, is precisely the trade-off formalized in Q1.

Scale III – Mesoscopic kinetics

The mesoscopic kinetic baseline evolves the master equation of a CRN:

$$\frac{\mathrm{d}\mathbf{P}}{\mathrm{d}t} = \mathbf{W}\mathbf{P}, \tag{19}$$

where $\mathbf{P}(t)$ is the probability vector over all microstates of the surface and $\mathbf{W}$ is the transition rate matrix whose off-diagonal elements are proportional to the rate constants k_j derived from Scale I/II.[12,13] For a surface with N_{sites} sites, the state space grows exponentially, making direct classical integration intractable. As shown in the CO oxidation example, classical kMC faces severe bottlenecks when rate constants span many orders of magnitude.

Alternative quantum approaches to solving eqn (19) generally rely on mapping the stochastic master equation to a continuous-variable system *via* Schrödingerization or quantum linear systems algorithms (QLSA),[54,55] which require expanding non-Hermitian operators into higher-dimensional dilated Hamiltonians. While these approaches offer polynomial advantages for deterministic differential integration, they struggle with the intrinsically non-linear structures generated by macroscopic mass-action constraints.

To overcome this, Scale III is replaced by a multidimensional QRW mapped directly onto the non-linear mass-action system graph using circuit-theoretic equivalents.[32] Rather than relying

© 2026 The Author(s). Published by the Royal Society of Chemistry

on generic, unstructured graph-walk speedups which fail to track steady-state distributions, this framework maps the CRN onto multidimensional electrical networks utilizing alternative Kirchhoff's and Ohm's laws *via* quantum random access memory (QRAM).[32]

It is critical to acknowledge the hardware overhead this architecture imposes. The proposed query speedups rely strictly on the availability of QRAM for efficient species reachability queries and state sampling. Fault-tolerant QRAM architectures pose physical and engineering challenges. If QRAM cannot be realized with overheads that scale sub-linearly with the CRN state space, the kinetic sampling bottleneck will remain unbroken. This highlights a critical hardware–software co-design constraint: mesoscopic quantum speedups are as dependent on memory-access architecture as they are on algorithmic query complexity.

The general CRN topology is encoded by the stoichiometric matrix $\boldsymbol{U} = [\nu'_{ij} - \nu_{ij}]_{N\times M}$ across the reaction pathways:

$$\sum_i \nu_{ij} S_i \xrightarrow{k_j} \sum_i \nu'_{ij} S_i, \quad j = 1, \ \ldots, \ M, \tag{20}$$

This specific multidimensional architecture yields provable quantum query speedups to decide species reachability, sample highly probable reachable states under network perturbations, and approximate steady-state fluxes over classical kMC.[32]

To transition this mesoscopic picture to the macroscopic continuum description of Scale IV, the steady-state surface coverages θ_i^* and reaction fluxes J_j^* are extracted from the final quantum state preparation. The net reaction rate R_j for each pathway is computed directly from these quantum-accelerated steady-state fluxes. These rates are then contracted with the stoichiometric matrix to yield the local volumetric net production rates:

$$\dot{\omega}_i = \sum_{j=1}^{M} \left(\nu'_{ij} - \nu_{ij}\right) R_j \tag{21}$$

These scalar source terms $\dot{\omega}_i$ successfully compress the discrete, exponentially scaled microstate probabilities $\mathbf{P}(t)$ of the surface down to local, deterministic chemical source profiles. By passing $\dot{\omega}_i$ forward, Scale III provides the mathematical boundary conditions required to couple the microscopic quantum kinetics directly into the fluid dynamics and transport equations of the macroscopic reactor scale.

Scale IV – Reactor scale

At the macroscale, species mass and energy conservation yield a system of coupled, non-linear macroscopic partial differential equations (PDEs):

$$\frac{\partial c_i}{\partial t} + \nabla \cdot \mathbf{J}_i = \dot{\omega}_i\left(\left\{\theta_j^*\right\}, T, \{c_k\}\right), \tag{22}$$

where c_i is the gas-phase concentration of species i, $\mathbf{J}_i$ is its spatial transport flux, and $\dot{\omega}_i$ is the volumetric net production rate explicitly compiled and passed forward from the Scale III CRN.

When discretized on a spatial mesh of n nodes, the full reactor model becomes a highly non-linear algebraic system. To port this to a quantum architecture, two distinct paths exist. The first is a hybrid framework: linearizing the non-linear source terms about the current operating point *via* a classical outer Newton loop, reducing the inner problem to a large, sparse linear system of size $n \times n$. QSVT sparse linear system solvers map this inner task into the fault-tolerant domain,[34] executing the matrix inversion $A\mathbf{x} = \mathbf{b}$ with a gate complexity of $\mathcal{O}(\text{poly}\log(n)\cdot\kappa\cdot\log(1/\varepsilon))$, where κ is the condition number. Alternatively, fully quantum non-linear PDE algorithms utilize Carleman linearization or non-linear Schrödingerization to map the non-linearities directly into an infinite-dimensional linear Bose–Einstein or dilated state space.[35,56] Carleman-based approaches, however, carry a strict validity condition: efficient truncation of the linearized hierarchy is guaranteed only when dissipation dominates the non-linearity,[57] a condition that strongly non-linear Arrhenius source terms will generally violate and that must be verified for any given reactor regime. Both pathways can yield an exponential reduction in scaling with respect to the spatial mesh size n compared to classical direct solvers. Quantum ODE/PDE integrators expand this advantage cleanly to time-dependent propagation.[58]

Crucially, this algorithmic advantage remains strictly conditional on three major bottlenecks. First, for the hybrid Newton-QSVT approach, the classical computer must re-compile and re-load the block-encoding of the Jacobian matrix A at every iteration, presenting a severe data-loading overhead. Second, the speedup assumes the condition number κ remains well-behaved and does not scale exponentially with the mesh resolution n or fluid turbulence. Third, and most fundamentally, the entire framework is bound by the quantum input/output (I/O) problem. While the quantum solver efficiently prepares a state $|x\rangle$ proportional to the spatial solution vector, extracting the full macroscopic concentration field $c_i(\mathbf{x},t)$ *via* quantum state tomography requires $\mathcal{O}(n)$ measurements, which completely nullifies the exponential algorithmic speedup.

True quantum advantage at Scale IV is therefore strictly governed by a hardware–software-engineering co-design principle: the target engineering objective must not require full field visualization. Instead, it must target a global scalar property – such as integrated terminal reactor yield, total reactant conversion, or global entropy production – that can be extracted efficiently *via* quantum expectation value estimation ($\langle x|\hat{O}|x\rangle$) without collapsing the full spatial state vector.

State of the art and resource requirements

The current state of the art is considerably more mature for the individual quantum algorithmic components than for their integration into a complete multiscale workflow. Resource estimates have been reported for electronic-structure calculations and related quantum simulation tasks, whereas fault-tolerant resource requirements for chemically realistic finite-

© 2026 The Author(s). Published by the Royal Society of Chemistry

Table 4 Indicative logical-resource requirements across the four scales of the quantum multiscale workflow. The resource ranges are representative rather than universal and are intended to indicate the order of magnitude of logical-qubit requirements under fault-tolerant assumptions. Actual requirements depend strongly on target precision, Hamiltonian simulation method, state preparation, oracle and data-access costs, and the structure of the multiscale coupling

Scale	Indicative logical resources	Main resource drivers
Electronic	10^2–10^4	Number of orbitals, target precision, Hamiltonian simulation, ansatz or state-preparation requirements, and phase-estimation precision
Atomistic	10^2–10^5	System size, simulation time, force evaluation, quantum state preparation, sampling requirements, and accuracy of dynamical observables
Mesoscopic	10^3–10^5	Network size, number of reaction channels, oracle or data-access requirements, state preparation, and target precision
Continuum	10^3–10^6	Discretization size, conditioning of the discretized equations, linear-system or differentialequation solver requirements, state preparation, and readout
Multiscale workflow	10^3–10^6	Inter-scale coupling, repeated calls between scales, precision requirements, data transfer, state preparation, and the extent to which intermediate quantities can be compressed

temperature dynamics, reaction-network kinetics, and continuum-scale calculations remain less developed. Moreover, existing estimates generally treat the inputs from preceding scales as externally supplied. Thus, an end-to-end resource estimate for a chemically consistent quantum multiscale workflow, including the transfer of information between scales, has not yet been established. Table 3 summarizes representative demonstrations, reported resource estimates, and the remaining requirements at each scale.

A quantitative estimate of the hardware requirements for this complete multiscale workflow is necessarily uncertain because the resources depend on algorithmic choices, target precision, state preparation, data access, and the transfer of information between scales. Table 4 therefore provides indicative logical-qubit requirements for each scale and for the combined CO/Pt(111) workflow, based on currently available resource estimates and algorithmic scaling.

The individual estimates should be regarded as order-of-magnitude indicators rather than definitive requirements. In particular, the resources of the four scales cannot simply be added, since calculations at one scale may need to be repeated as inputs from another scale are updated. The overall requirement is therefore strongly dependent on the frequency of inter-scale communication and the extent to which intermediate quantities can be compressed. For the complete workflow, we estimate an indicative requirement of approximately 10^3–10^6 logical qubits. This range should be viewed as a projection based on the current state of the art rather than a prediction of future hardware requirements.

Although CO oxidation on Pt(111) provides a concrete illustrative example, the proposed hierarchy is applicable to a broader range of catalytic and electrochemical systems, including ammonia synthesis, C–H activation, CO_2 conversion, hydrogen and oxygen evolution and reduction reactions, and catalysis on transition-metal alloys, single-atom catalysts, and magnetic surfaces. Related multiscale structures also arise in molecular-interaction problems relevant to drug discovery, where electronic structure, molecular dynamics, free-energy sampling, and kinetic models can be coupled. The specific quantum algorithms and inter-scale interfaces will depend on the chemistry and its dominant computational bottlenecks; the framework therefore provides a strategy for identifying where quantum subroutines could be incorporated, rather than implying quantum advantage across all scales or application domains.

Open questions in quantum multiscale modeling

Having established that a rigorous quantum analogue exists within each distinct scale, the overarching scientific problem shifts from individual algorithmic isolation to systemic architectural coupling. Table 5 distills this challenge into six interconnected open questions. Q1 and Q2 concern the boundaries themselves–the mathematical structure of the inter-scale channel and whether all boundaries in the hierarchy are equally difficult to implement; Q3 evaluates whether a unified algebraic framework can compose the quantum subroutines across these boundaries; Q4 addresses how uncertainty propagates through them; and Q5 and Q6 confront the end-to-end question: the total query complexity of the composed pipeline and the structural conditions under which it retains a net quantum advantage. Rather than treating these interfaces as mere data-passing steps, we state each problem with mathematical precision as a necessary first step toward formalizing a true calculus of quantum multiscale simulation.

More formally, we define the inter-scale interface $\mathcal{I}_{k\to k+1}$ as a transformation acting on a quantum state (or density operator) ρ_k, associated uncertainty budget ε_k, and admissible observable set $\mathcal{O}_k$, producing inputs for the downstream algorithm. Classical multiscale methods correspond to projective measurement followed by parameter compression, whereas coherent interfaces preserve quantum information through CPTP maps or generalized quantum instruments.

Questions Q1–Q6 should be interpreted as open theoretical problems rather than established properties of quantum multiscale architectures. Wherever examples are provided below,

© 2026 The Author(s). Published by the Royal Society of Chemistry

Table 5 Six open questions defining the quantum multiscale composition problem. $\mathcal{I}_{k\to k+1}$ denotes the inter-scale quantum or classical information channel from Scale k to Scale $k+1$

Q1	*Mathematical structure of the inter-scale channel*. What is the minimal mathematical description of $\mathcal{I}_{k\to k+1}$: a memoryless CPTP map, a generalized quantum instrument, or a non-Markovian channel carrying temporal correlations injected by cross-scale coarse-graining? When adjacent scales both execute fault-tolerant quantum algorithms, must the channel terminate in a projective measurement that collapses the state to classical parameters, or can quantum information be transmitted coherently *via* density-matrix mappings–and what is the query complexity cost of each paradigm?
Q2	*Heterogeneity of inter-scale boundaries*. Are all channels $\mathcal{I}_{k\to k+1}$ equally difficult to implement? The lower boundaries transfer microscopic, Hamiltonian-level information between algorithms sharing a common physical representation, whereas the upper boundaries exchange averaged statistical and continuum-level quantities; the central boundary must bridge these two regimes. Does this representational divide render some channels naturally amenable to coherent composition while others remain irreducibly dissipative – and do hybrid quantum–classical pipelines inherit the same asymmetry at their classical–quantum boundaries?
Q3	*QSVT as a unifying computational primitive*. Can the QSVT serve as a single unifying primitive to implement all core quantum subroutines – QPE, Hamiltonian simulation, quantum walks, and linear system inversions? If so, does this algebraic unification reduce scale-bridging to the clean composition of block-encoded operators within hybrid or fully coherent architectures?
Q4	*Uncertainty propagation across quantum channels*. Quantum amplitude estimation offers a generic quadratic speedup for uncertainty quantification (UQ) across individual scales. However, does the inter-scale quantum channel $\mathcal{I}_{k\to k+1}$ inherently amplify, preserve, or dampen irreducible quantum measurement uncertainty under cross-scale coarse-graining?
Q5	*End-to-end query complexity of the global pipeline*. What is the total query complexity of the integrated four-scale quantum pipeline as a function of target global precision ε, global system size N, and total simulation time T? Is this multi-layer complexity provably sub-polynomial in the corresponding classical multiscale baseline, and which scale boundaries dictate the dominant bottleneck?
Q6	*Conditions for end-to-end quantum advantage*. If physical or informational constraints force $\mathcal{I}_{k\to k+1}$ to be implemented *via* classical measurement and re-preparation, what structural conditions on the individual subroutines guarantee that the composed pipeline retains a net quantum advantage? Under what conditions does the sampling overhead of classical readout at the boundary destroy the speedup achieved within individual layers?

they serve as illustrative constructions rather than proofs of sufficiency or necessity.

Q1 – Mathematical structure of the inter-scale channel

The physical act of cross-scale coarse-graining discards microscopic degrees of freedom, transforming a closed, unitary quantum system at Scale k into an open quantum system at the boundary interface. This transition requires a precise description of the resulting informational channel $\mathcal{I}_{k\to k+1}$. Classical multiscale protocols universally rely on a memoryless Markovian assumption: information is compressed into static scalar parameters (*e.g.*, an energy or a rate constant) under the premise that fast environmental correlations relax instantaneously. In the fault-tolerant quantum computing domain, however, it remains open whether a memoryless CPTP map or a generalized quantum instrument is mathematically sufficient to represent this interface. If the quantum correlations generated within strongly correlated adsorbate systems decay on a timescale comparable to or longer than the downstream coarse-grained dynamics, the interface must instead be modeled as a non-Markovian quantum channel possessing temporal memory. Defining the boundary between Markovian decay and structured quantum memory is the foundational step toward a rigorous scale-crossing calculus.

Whatever its mathematical form, the channel must also be assigned a physical implementation, and here the classical template for multiscale modeling is unambiguous: every classical scale boundary realizes $\mathcal{I}_{k\to k+1}$ as a lossy, dissipative transduction, in which a projective measurement extracts a classical scalar that is injected manually into the next layer. While measurement is computationally free in standard workflows, this design can introduce bottlenecks on a quantum architecture. If Scale k produces a multi-particle state $|\psi\rangle$ encoding a probability distribution over surface configurations, projective measurement of configuration probabilities irreversibly destroys phase information that cannot be reconstructed at Scale $k+1$ without re-running Scale k, while re-preparing an equivalent state downstream may, depending on the problem, incur a state-preparation cost that offsets the localized quantum speedup. The alternative, transmitting the density matrix directly to Scale $k+1$ as a coherent input, demands that the downstream algorithm be explicitly engineered to accept unmeasured quantum inputs. Tracking the query complexity tradeoff between these two paradigms, state-collapsing measurement *versus* coherent channel transmission, is the second half of the structural question that Q1 poses.

Q2 – Heterogeneity of inter-scale boundaries

The four scale boundaries of the hierarchy are not equal. At the lower boundary ($\mathcal{I}_{\mathrm{I}\to\mathrm{II}}$), both algorithms operate on microscopic, Hamiltonian-level representations of the same physical system: eigenstates, forces, and thermal density operators share a common mathematical language, which is precisely why a coherent construction such as the conditional Gibbs-state preparation of the CO/Pt(111) case study is conceivable at all. At the upper boundary ($\mathcal{I}_{\mathrm{III}\to\mathrm{IV}}$), both sides already have averaged, statistical quantities, such as coverages, fluxes, and source terms, whose transfer as classical scalars sacrifices little, since the coarse-graining has largely been paid for upstream. The

© 2026 The Author(s). Published by the Royal Society of Chemistry

central boundary ($\mathcal{I}_{\mathrm{II}\to\mathrm{III}}$) is qualitatively different: it must convert microscopic dynamical information into the statistical parameterization of a master equation, crossing the representational divide between mechanics and kinetics. We conjecture that this central boundary is where coherent transfer is simultaneously most valuable, where the exponentially large configuration space has not yet been compressed, and most difficult to engineer. The open question is whether this asymmetry is intrinsic to the representational divide or merely an artifact of the currently available algorithmic toolbox, and whether hybrid quantum–classical pipelines, in which some layers remain classical, inherit the same asymmetry at their classical–quantum boundaries.

Q3 – QSVT as a unifying computational primitive

The QSVT establishes that the seemingly disparate algorithms defining our multiscale tiers – QPE, Hamiltonian simulation, quantum walks, and linear system solvers – are structurally unified under the algebraic umbrella of polynomial transformations of block-encoded matrices.[34,63] Given a standard block-encoding of a matrix A within a unitary U_A such that $(\langle 0|\otimes I)U_A(|0\rangle\otimes I) = A/\alpha$, QSVT can execute any chosen polynomial transformation p on the singular values of A using only $\mathcal{O}(\deg(p))$ calls to U_A. Within this context, QPE represents phase estimation on a target block-encoding; Hamiltonian simulation applies a Chebyshev approximation of e^{-iHt}; the mesoscopic walk executes a polynomial that amplifies the graph's target subspace; and the inner reactor PDE step evaluates a polynomial that inverts the singular values of the sparse linear subsystem.

If these core algorithmic blocks can be compiled as specialized QSVT polynomials over respective block-encoded operators, the scale-bridging problem mathematically simplifies to the arithmetic composition of block-encodings.[34] However, the massive structural and algebraic disparity between a highly correlated molecular electronic Hamiltonian (Scale I) and a sparse continuum transport matrix (Scale IV)–alongside the requirement for an outer classical Newton loop at the reactor scale–means a single, monolithic end-to-end quantum circuit is structurally infeasible. Instead, a viable unified QSVT architecture requires a sequence of distinct, scale-specific block-encodings interconnected by coherent quantum routing and hybrid classical data-loading subroutines.

Q4 – Uncertainty propagation across quantum channels

In classical multiscale modeling, parameter uncertainties–such as functional-induced errors in DFT activation barriers–propagate nonlinearly through the hierarchy, driving exponential variance in predicted macroscopic rates. While quantum amplitude estimation provides a generic quadratic speedup for uncertainty quantification (UQ) within any individual layer, reducing the classical sampling overhead from $\mathcal{O}(\varepsilon^{-2})$ to $\mathcal{O}(\varepsilon^{-1})$, composing these bounds across scales poses a fundamentally different mathematical challenge. The critical open question is how irreducible quantum measurement uncertainty propagates when transmitted through the coarse-grained inter-scale boundary $\mathcal{I}_{k\to k+1}$, and whether this open quantum channel inherently amplifies, preserves, or dampens errors compared to classical non-linear propagation.

More generally, if inter-scale channels satisfy $\varepsilon_{k+1} = f_k(\varepsilon_k)$, where f_k denotes the uncertainty transformation induced by the inter-scale channel $\mathcal{I}_{k\to k+1}$, then quantum multiscale uncertainty propagation reduces to characterizing the composed map:

$$\varepsilon_{\mathrm{global}} = f_n \circ f_{n-1} \circ \cdots \circ f_1(\varepsilon_0). \tag{23}$$

The mathematical properties of the transformations f_k therefore determine whether uncertainty is amplified, preserved, or damped under successive coarse-graining across scales.

Q5 – End-to-end query complexity of the global pipeline

The individual quantum speedups distributed across the multiscale layers span highly diverse functional dependencies: polynomial scaling in basis functions ($\mathcal{O}(\|H\|/\varepsilon)$), real-time propagation bounds ($\mathcal{O}(t\cdot\|H\|)$), and polylogarithmic matrix inversion scales ($\mathcal{O}(\mathrm{poly}\log(n)\cdot\kappa)$). In classical engineering, the dominant computational bottleneck resides at the mesoscopic step (Scale III), where the exponential state space of the master equation forces stochastic sampling *via* kMC.

A closed-form bound on the total query complexity of the composed four-scale pipeline does not yet exist, and deriving one is non-trivial: the scaling parameters that dominate each layer, such as $\|H\|$, spectral gap Δ, and condition number κ, are not independent, and their joint dependence on global system size N and target precision ε remains unknown. Establishing even a partial bound – for example, at the Scale I/II interface alone – would constitute meaningful progress toward a full pipeline complexity theorem.

Resolving this overarching complexity bound fundamentally requires characterizing the spectral properties of the specific CRN at Scale III, as these govern how the localized quantum walk speedups behave when coupled to adjacent scales. Whether the multidimensional quantum walk framework on mass-action system graphs can decisively break the mesoscopic bottleneck under chemically realistic, sparse topologies with deeply heterogeneous rate hierarchies remains to be demonstrated.[32] Evaluating the total query complexity of the integrated pipeline requires solving a dual theoretical and algorithmic problem: deriving the strict graph properties and spectral gaps governing mass-action chemical networks under steady-state perturbations, and proving that the compiled quantum walk preserves its query advantages when bound to the data inputs and outputs of its adjacent scales.

Q6 – Conditions for end-to-end quantum advantage

Beyond the aggregate complexity of Q5, the end-to-end advantage of the composed pipeline remains highly sensitive to boundary degradation. If physical or informational constraints force an inter-scale interface $\mathcal{I}_{k\to k+1}$ to remain a classical intermediary, error propagation cascades nonlinearly. For

© 2026 The Author(s). Published by the Royal Society of Chemistry

example, if a Scale I QPE routine requires $\mathrm{Time_I}$ gate operations to isolate an activation barrier to precision ε_I, this value is classically processed *via* a highly non-linear, ill-conditioned Arrhenius expression ($k_j \propto \mathrm{e}^{-E_a/k_\mathrm{B}T}$) before parameterizing the Scale III quantum walk.[12,13]

Small variations in ε_I propagate exponentially through the pipeline, meaning the localized benefit of a low $\mathrm{Time_I}$ only accrues if downstream sensitivities are tightly bounded. Composing multiple quantum layers under a global target precision constraint $\varepsilon = \sum_k \varepsilon_k$ thus transforms into a non-trivial, multi-variable error allocation problem. To prevent order-of-magnitude downstream rate distortions at typical catalytic reaction temperatures ($T \sim 500$ K), this error budgeting requires constraining the Scale I algorithmic precision to chemical accuracy ($\varepsilon_\mathrm{I} \leq 1.6$ mHa ≈ 0.043 eV $\approx k_\mathrm{B}T$ at $T = 500$ K). This formalizes a rigorous mathematical boundary where the maximum permissible global pipeline error ε dictates strict upper bounds on the acceptable logical qubit overhead and runtime thresholds for the underlying electronic oracles.

Whether coherent inter-scale coupling provides an accuracy advantage over measurement followed by re-preparation remains an open question. Formal frameworks for composing quantum subroutines, including on superpositions of inputs,[64,65] and quantum algorithms coupling microscopic and macroscopic reaction–diffusion dynamics[66] provide a basis for addressing this question. To date, there is no chemically realistic numerical demonstration showing that retaining coherence across scale boundaries improves the accuracy of a chemically relevant observable. The potential advantage would arise if measurement discards correlations or other information that is relevant to downstream calculations; if the measured quantities form a sufficient representation of the information required at the next scale, measurement and re-preparation may be equally accurate. Establishing when coherent coupling provides a measurable benefit, and whether that benefit justifies its additional resource requirements, is therefore an important direction for future research.

Outlook

The framework developed in this Perspective identifies four fault-tolerant quantum algorithms – quantum phase estimation, coherent Hamiltonian simulation with Gibbs state preparation, multidimensional quantum walks, and QSVT-based sparse linear solvers – as the natural mathematical counterparts of the classical multiscale hierarchy for complex chemical systems including heterogeneous catalysis. We have formalized six open questions whose resolution will determine whether these algorithms can be systematically composed into a unified pipeline that preserves individual localized advantages to deliver a net quantum speedup over state-of-the-art classical multiscale implementations. While this framework is explicitly fault-tolerant in scope, near-term hybrid approaches – such as variational quantum eigensolvers at Scale I or quantum–classical surrogate models at Scale III – may serve as stepping stones that stress-test the inter-scale composition interfaces identified here before fault-tolerant hardware matures.

The central mathematical challenge, crystallized by Questions Q1 and Q2, is the rigorous characterization of the inter-scale information channel $\mathcal{I}_{k\to k+1}$ when adjacent layers operate fully within the quantum domain. The algebraic machinery of QSVT provides a highly promising blueprint for this challenge; by defining an arithmetic for combining block-encoded operators, it offers a path to express inter-scale boundaries as composed block-encodings rather than informationally destructive classical measurements. Concurrently, formal frameworks for the composition of quantum algorithms and nested subroutines, such as those developed by Jeffery provide the foundational query complexity machinery required to analyze how errors and oracle calls propagate across nested quantum structures.[64,65] Whether these algorithmic tools can yield a universal composition theorem, establishing the necessary and sufficient conditions under which cross-scale quantum information transfer remains lossless with respect to end-to-end query complexity, remains the foundational open problem this Perspective defines.

A critical dimension of this composition problem is *uncertainty quantification* (UQ), which must be evaluated not as an isolated problem, but as an intrinsic feature of open quantum channels. In classical multiscale modeling, uncertainties in DFT activation barriers (typically ± 0.1 eV) propagate non-linearly through Arrhenius expressions, resulting in exponential sensitivity in downstream kMC rate constants and predicted turnover frequencies. Quantum amplitude estimation offers a generic quadratic speedup for UQ across every individual scale, reducing the classical sampling cost from $\mathcal{O}(\varepsilon^{-2})$ to $\mathcal{O}(\varepsilon^{-1})$.[67] This suggests a future analogue of classical sensitivity analysis in which quantum multiscale pipelines are optimized by allocating precision budgets ε_k across scales under fixed logical-resource constraints. However, the deeper question highlighted by our framework (Question Q4) is how irreducible quantum measurement uncertainty propagates across a boundary that is itself an open quantum channel. Determining whether the channel $\mathcal{I}_{k\to k+1}$ inherently amplifies, preserves, or dampens uncertainty under cross-scale coarse-graining represents a vital mathematical constraint that directly shapes the answers to Questions Q1, Q2, and Q6.

Two additional technical qualifications on the pipeline deserve explicit statement. First, the multidimensional quantum walk at Scale III operates specifically on the non-linear mass-action system graph of the CRN *via* circuit theory analogies.[32] The multi-dimensional product structure of this state space is highly non-generic; it possesses explicit algebraic features that distinguish it from unstructured random walks on arbitrary graphs. Consequently, evaluating the global query complexity in Question Q5 requires exploiting this specific topological structure alongside QRAM data access architectures.

Second, the QSVT-based solver at Scale IV applies strictly to the *linearized* reactor PDE system arising at each Newton iteration of a global non-linear solver. Because the full macroscopic continuum reactor model is non-linear, the quantum advantage

© 2026 The Author(s). Published by the Royal Society of Chemistry

claimed at Scale IV is strictly conditional: it assumes the global classical non-linear solver spends its dominant computational cost in the repeated solution of large, sparse linear subsystems. This condition generally holds for ultra-fine, well-resolved reactor meshes, but must be verified case-by-case.

Throughout this Perspective, all quoted quantum advantages are conditional rather than universal: preserving asymptotic improvements requires assumptions regarding state preparation, memory access, spectral properties, conditioning, and observable extraction that must ultimately be validated for chemically realistic systems. The value of the framework presented here is to explicitly define the fault-tolerant research agenda for *quantum multiscale modeling*. The core hurdle facing quantum multiscale modeling is not demonstrating that a quantum analogues exist at each scale; they do, though improvements in those single-scale algorithms are both needed and expected. The true challenge is demonstrating that the interfaces between those scales can be engineered to carry information without causing an informational collapse that completely destroys the quantum advantage individual layers offer. Posing that question with mathematical precision is the first necessary step toward its resolution.

Author contributions

S. H. conceived the idea of quantum multiscale modeling and initiated the project (conceptualization) and prepared the initial draft of the manuscript (writing – original draft). K. S., S. K., and L. V. contributed to the development and refinement of the manuscript through critical discussion, revision, and co-writing (writing – review & editing). K. S. obtained and managed the project funding (funding acquisition, project administration). S. K. provided project funding from Toyota Motor Europe and contributed to project administration (funding acquisition, project administration). L. V. provided supervision and contributed to project administration (supervision, project administration). All authors discussed the concepts presented in the manuscript and approved the final version.

Conflicts of interest

The authors declare no conflict of interest.

Data availability

No primary research data were generated or analyzed in this Perspective; all results discussed are available in the cited literature.

References

1 A. Bruix, J. T. Margraf, M. Andersen and K. Reuter, First-Principles-Based Multiscale Modelling of Heterogeneous Catalysis, *Nat. Catal.*, 2019, **2**, 659–670, DOI: 10.1038/s41929-019-0298-3.

2 C. E. Clancy, G. An, W. R. Cannon, Y. Liu, E. E. May, P. Ortoleva, A. S. Popel, J. P. Sluka, J. Su, P. Vicini, *et al.*, Multiscale Modeling in the Clinic: Drug Design and Development, *Ann. Biomed. Eng.*, 2016, **44**, 2591–2610, DOI: 10.1007/s10439-016-1563-0.

3 J. Fish, G. J. Wagner and S. Keten, Mesoscopic and Multiscale Modelling in Materials, *Nat. Mater.*, 2021, **20**, 774–786, DOI: 10.1038/s41563-020-00913-0.

4 R. Radhakrishnan, A Survey of Multiscale Modeling: Foundations, Historical Milestones, Current Status, and Future Prospects, *AIChE J.*, 2021, **67**, e17026, DOI: 10.1002/aic.17026.

5 C. M. Friend and B. Xu, Heterogeneous Catalysis: A Central Science for a Sustainable Future, *Acc. Chem. Res.*, 2017, **50**, 517–521, DOI: 10.1021/acs.accounts.6b00510.

6 A. Hoekstra, B. Chopard and P. Coveney, Multiscale Modelling and Simulation: A Position Paper, *Philos. Trans. R. Soc., A*, 2014, **372**, 20130377, DOI: 10.1098/rsta.2013.0377.

7 H. Ma, J. Liu, H. Shang, Y. Fan, Z. Li and J. Yang, Multiscale Quantum Algorithms for Quantum Chemistry, *Chem. Sci.*, 2023, **14**, 3190–3205, DOI: 10.1039/D2SC06875C.

8 M. Capone, M. Romanelli, D. Castaldo, G. Parolin, A. Bello, G. Gil and M. Vanzan, A Vision for the Future of Multiscale Modeling, *ACS Phys. Chem. Au*, 2024, **4**, 202–225, DOI: 10.1021/acsphyschemau.3c00080.

9 T. M. Bickley, A. Mingare, T. Weaving, M. Williams de la Bastida, S. Wan, M. Nibbi, P. Seitz, A. Ralli, P. J. Love, M. Chung, *et al.*, Extending quantum computing through subspace, embedding and classical molecular dynamics techniques, *Digital Discovery*, 2025, **4**, 3427–3444, DOI: 10.1039/d5dd00225g.

10 J. K. Nørskov, T. Bligaard, J. Rossmeisl and C. H. Christensen, Towards the Computational Design of Solid Catalysts, *Nat. Chem.*, 2009, **1**, 37–46, DOI: 10.1038/nchem.121.

11 D. Marx and J. Hutter, Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods, *Modern methods and algorithms of quantum chemistry*, 2000, vol. 1, p. 141, DOI: 10.1017/CBO9780511609633.

12 M. Pineda and M. Stamatakis, Kinetic Monte Carlo Simulations for Heterogeneous catalysis: Fundamentals, Current Status, and Challenges, *J. Chem. Phys.*, 2022, **156**, 120902, DOI: 10.1063/5.0083251.

13 K. Reuter, in *Modeling and Simulation of Heterogeneous Catalytic Reactions*, John Wiley & Sons, Ltd, 2011, ch. 3, pp. 71–111, DOI: 10.1002/9783527639878.ch3.

14 O. Deutschmann, Modeling of the Interactions Between Catalytic Surfaces and Gas-phase, *Catal. Lett.*, 2015, **145**, 272–289, DOI: 10.1007/s10562-014-1431-1.

15 A. J. Cohen, P. Mori-Sánchez and W. Yang, Insights into Current Limitations of Density Functional Theory, *Science*, 2008, **321**, 792–794, DOI: 10.1126/science.1158722.

16 R. J. Bartlett and M. Musiał, Coupled-cluster Theory in Quantum Chemistry, *Rev. Mod. Phys.*, 2007, **79**, 291–352, DOI: 10.1103/RevModPhys.79.291.

17 W. Jia, H. Wang, M. Chen, D. Lu, L. Lin, R. Car, E. Weinan and L. Zhang, in *SC20: International conference for high performance computing, networking, storage and analysis*, 2020, pp. 1–14, DOI: 10.5555/3433701.3433707.

© 2026 The Author(s). Published by the Royal Society of Chemistry

18 A. Chatterjee and D. G. Vlachos, An Overview of Spatial Microscopic and Accelerated Kinetic Monte Carlo Methods, *J. Comput.-Aided Mater. Des.*, 2007, **14**, 253–308, DOI: 10.1007/s10820-006-9042-9.
19 H. A. Jakobsen, Chemical Reactor Modeling, *Multiphase Reactive Flows*, 2008, DOI: 10.1007/978-3-319-05092-8.
20 A. Omranpour, J. Elsner, K. N. Lausch and J. Behler, Machine Learning Potentials for Heterogeneous Catalysis, *ACS Catal.*, 2025, **15**, 1616–1634, DOI: 10.1021/acscatal.4c06717.
21 I. Batatia, P. Benner, Y. Chiang, A. M. Elena, D. P. Kovács, J. Riebesell, X. R. Advincula, M. Asta, M. Avaylon, W. J. Baldwin, *et al.*, A Foundation Model for Atomistic Materials Chemistry, *J. Chem. Phys.*, 2025, **163**, 184110, DOI: 10.1063/5.0297006.
22 J. T. Margraf and K. Reuter, Systematic Enumeration of Elementary Reaction Steps in Surface Catalysis, *ACS Omega*, 2019, **4**, 3370–3379, DOI: 10.1021/acsomega.8b03200.
23 J. T. Margraf, H. Jung, C. Scheurer and K. Reuter, Exploring Catalytic Reaction Networks with Machine Learning, *Nat. Catal.*, 2023, **6**, 112–121, DOI: 10.1038/s41929-022-00896-y.
24 L. H. Kempen, R. Cheula and M. Andersen, How Accurate are Foundational Machine Learning Interatomic Potentials for Heterogeneous Catalysis?, *J. Chem. Phys.*, 2026, **164**, 194119, DOI: 10.1063/5.0317672.
25 S. Hariharan, S. Kinge and L. Visscher, Modeling Heterogeneous Catalysis using Quantum Computers: An Academic and Industry Perspective, *J. Chem. Inf. Model.*, 2024, **65**, 472–511, DOI: 10.1021/acs.jcim.4c01212.
26 R. Babbush, N. Wiebe, J. McClean, J. McClain, H. Neven and G. K.-L. Chan, Low-Depth Quantum Simulation of Materials, *Phys. Rev. X*, 2018, **8**, 011044, DOI: 10.1103/PhysRevX.8.011044.
27 J. Lee, D. W. Berry, C. Gidney, W. J. Huggins, J. R. McClean, N. Wiebe and R. Babbush, Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction, *PRX Quantum*, 2021, **2**, 030305, DOI: 10.1103/PRXQuantum.2.030305.
28 M. Reiher, N. Wiebe, K. M. Svore, D. Wecker and M. Troyer, Elucidating Reaction Mechanisms on Quantum Computers, *Proc. Natl. Acad. Sci. U. S. A.*, 2017, **114**, 7555–7560, DOI: 10.1073/pnas.1619152114.
29 A. M. Childs, Y. Su, M. C. Tran, N. Wiebe and S. Zhu, Theory of Trotter Error with Commutator Scaling, *Phys. Rev. X*, 2021, **11**, 011020, DOI: 10.1103/PhysRevX.11.011020.
30 G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.*, 2017, **118**, 010501, DOI: 10.1103/PhysRevLett.118.010501.
31 S. Jeffery and S. Zur, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, Association for Computing Machinery, Orlando, FL, USA, 2023, pp. 1125–1130, DOI: 10.1145/3564246.3585158.
32 S. Hariharan, S. Zur, S. Kinge, L. Visscher, K. Schoutens and S. Jeffery, Quantum Walks for Chemical Reaction Networks, *arXiv*, 2026, preprint, arXiv:2509.07890, DOI: 10.48550/arXiv.2509.07890, https://arxiv.org/abs/2509.07890.
33 A. W. Harrow, A. Hassidim and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.*, 2009, **103**, 150502, DOI: 10.1103/PhysRevLett.103.150502.
34 A. Gilyén, Y. Su, G. H. Low and N. Wiebe, in *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, 2019, pp. 193–204, DOI: 10.1145/3313276.3316366.
35 D. An, J.-P. Liu and L. Lin, Linear combination of Hamiltonian simulation for nonunitary dynamics with optimal state preparation cost, *Phys. Rev. Lett.*, 2023, **131**, 150603, DOI: 10.1103/PhysRevLett.131.150603.
36 D. Castaldo and M. Reiher, Utility-scale Quantum Computational Chemistry, *arXiv*, 2026, preprint, arXiv:2603.19081, DOI: 10.48550/arXiv.2603.1908, https://arxiv.org/abs/2603.19081.
37 F. Jensen, *Introduction to computational chemistry*, John Wiley & Sons, 3rd. edn, 2017.
38 T. B. Pedersen, S. Lehtola, I. Fdez Galván and R. Lindh, The Versatility of the Cholesky Decomposition in Electronic Structure Theory, *Wiley Interdiscip. Rev.:Comput. Mol. Sci.*, 2024, **14**, e1692, DOI: 10.1002/wcms.1692.
39 P. R. Nagy, State-of-the-Art Local Correlation Methods Enable Affordable Gold Standard Quantum Chemistry for up to Hundreds of Atoms, *Chem. Sci.*, 2024, **15**, 14556–14584, DOI: 10.1039/D4SC04755A.
40 N. Glaser and M. Reiher, Solving Intractable Chemical Problems by Tensor Decomposition, *Chimia*, 2024, **78**, 215–221, DOI: 10.2533/chimia.2024.215.
41 A. George, S. Mondal, M. Purnaprajna and P. Athri, Review of electrostatic force calculation methods and their acceleration in molecular dynamics packages using graphics processors, *ACS Omega*, 2022, **7**, 32877, DOI: 10.1021/acsomega.2c03189.
42 Y. T. Lin, S. Feng and W. S. Hlavacek, Scaling methods for accelerating kinetic Monte Carlo simulations of chemical reaction networks, *J. Chem. Phys.*, 2019, **150**, 244101, DOI: 10.1063/1.5096774.
43 S. Raimondeau and D. G. Vlachos, Recent developments on multiscale, hierarchical modeling of chemical reactors, *Chem. Eng. J.*, 2002, **90**, 3–23, DOI: 10.1016/S1385-8947(02)00065-7.
44 Y. Saad, *Iterative methods for sparse linear systems*, SIAM, 2003.
45 P. J. Feibelman, B. Hammer, J. K. Nørskov, F. Wagner, M. Scheffler, R. Stumpf, R. Watwe and J. Dumesic, The CO/Pt (111) Puzzle, *J. Phys. Chem. B*, 2001, **105**, 4018–4025, DOI: 10.1021/jp002302t.
46 I. Kassal, S. P. Jordan, P. J. Love, M. Mohseni and A. Aspuru-Guzik, Polynomial-time Quantum Algorithm for the Simulation of Chemical Dynamics, *Proc. Natl. Acad. Sci. U. S. A.*, 2008, **105**, 18681–18686, DOI: 10.1073/pnas.0808245105.
47 N. J. Ward, I. Kassal and A. Aspuru-Guzik, Preparation of Many-body States for Quantum Simulation, *J. Chem. Phys.*, 2009, **130**, 194105, DOI: 10.1063/1.3115177.
48 T. E. O'Brien, B. Tarasinski and B. M. Terhal, Quantum Phase Estimation of Multiple Eigenvalues for Small-scale

 © 2026 The Author(s). Published by the Royal Society of Chemistry

(noisy) Experiments, *New J. Phys.*, 2019, **21**, 023022, DOI: 10.1088/1367-2630/aafb8e.
49 A. N. Chowdhury and R. D. Somma, Quantum Algorithms for Gibbs Sampling and Hitting-time Estimation, *arXiv*, 2016, preprint, arXiv:1603.02940, DOI: 10.48550/arXiv.1603.02940.
50 S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, *et al.*, Evaluating the Evidence for Exponential Quantum Advantage in Ground-state Quantum Chemistry, *Nat. Commun.*, 2023, **14**, 1952, DOI: 10.1038/s41467-023-37587-6.
51 J. J. Goings, A. White, J. Lee, C. S. Tautermann, M. Degroote, C. Gidney, T. Shiozaki, R. Babbush and N. C. Rubin, 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**, e2203533119, DOI: 10.1073/pnas.2203533119.
52 G. H. Low, R. King, D. W. Berry, Q. Han, A. E. DePrince, A. F. White, R. Babbush, R. D. Somma and N. C. Rubin, Fast Quantum Simulation of Electronic Structure by Spectral Amplification, *Phys. Rev. X*, 2025, **15**, 041016, DOI: 10.1103/pb2g-j9cw.
53 A. M. Childs and N. Wiebe, Hamiltonian Simulation using Linear Combinations of Unitary Operations, *arXiv*, 2012, preprint, arXiv:1202.5822, DOI: 10.26421/QIC12.11-12.
54 S. Jin, N. Liu and Y. Yu, Quantum Simulation of Partial Differential Equations: Applications and Detailed Analysis, *Phys. Rev. A*, 2023, **108**, 032603, DOI: 10.1103/PhysRevA.108.032603.
55 D. W. Berry, A. M. Childs, A. Ostrander and G. Wang, Quantum Algorithm for Linear Differential Equations with Exponentially Improved Dependence on Precision, *Commun. Math. Phys.*, 2017, **356**, 1057–1081, DOI: 10.1007/s00220-017-3002-y.
56 H. Krovi, Improved Quantum Algorithms for Linear and Nonlinear Differential Equations, *Quantum*, 2023, **7**, 913, DOI: 10.22331/q-2023-02-02-913.
57 J.-P. Liu, H. Ø. Kolden, H. K. Krovi, N. F. Loureiro, K. Trivisa and A. M. Childs, Efficient Quantum Algorithm for Dissipative Nonlinear Differential Equations, *Proc. Natl. Acad. Sci. U. S. A.*, 2021, **118**, e2026805118, DOI: 10.1073/pnas.2026805118.
58 D. W. Berry, High-order Quantum Algorithm for Solving Linear Differential Equations, *J. Phys. A: Math. Theor.*, 2014, **47**, 105301, DOI: 10.1088/1751-8113/47/10/105301.
59 M. Szczepanik, Á. Nagy and E. Zak, Simulating high-accuracy nuclear motion Hamiltonians using discrete variable representation and Walsh–Hadamard QROM on fault-tolerant quantum computers, *Quantum Sci. Technol.*, 2026, **11**, 045039, DOI: 10.1088/2058-9565/ae9d1b.
60 T. Kabengele, Y. M. Lokare, J. B. Marston and B. M. Rubenstein, Modeling stochastic chemical kinetics on quantum computers, *J. Chem. Phys.*, 2026, **164**, 194117, DOI: 10.1063/5.0332629.
61 J. Ingelmann, S. S. Bharadwaj, P. Pfeffer, K. R. Sreenivasan and J. Schumacher, Two quantum algorithms for solving the one-dimensional advection–diffusion equation, *Comput. Fluids*, 2024, **281**, 106369, DOI: 10.1016/j.compfluid.2024.106369.
62 I. Novikau and I. Joseph, Quantum algorithm for the advection-diffusion equation and the Koopman-von Neumann approach to nonlinear dynamical systems, *Comput. Phys. Commun.*, 2025, **309**, 109498, DOI: 10.1016/j.cpc.2025.109498.
63 J. M. Martyn, Z. M. Rossi, A. K. Tan and I. L. Chuang, Grand Unification of Quantum Algorithms, *PRX Quantum*, 2021, **2**, 040203, DOI: 10.1103/PRXQuantum.2.040203.
64 S. Jeffery, Quantum Subroutine Composition, *arXiv*, 2022, preprint, arXiv:2209.14146, DOI: 10.48550/arXiv.2209.14146.
65 S. Jeffery, Composing Quantum Algorithms, *arXiv*, 2025, preprint, arXiv:2502.09240, DOI: 10.48550/arXiv.2502.09240.
66 M. Lockwood, N. Wiebe, C. Johnson, J. Mülmenstädt, J. Chun, G. Schenter, M. S. Cheung and X. Li, Quantum Algorithm for Subcellular Multiscale Reaction-Diffusion Systems, *arXiv*, 2025, preprint, arXiv:2509.20668, DOI: 10.48550/arXiv.2509.20668.
67 G. Brassard, P. Høyer, M. Mosca and A. Tapp, in *Quantum Computation and Information; Contemporary Mathematics*, American Mathematical Society, vol. 305, 2002, pp. 53–74, DOI: 10.48550/arXiv.quant-ph/0005055.

© 2026 The Author(s). Published by the Royal Society of Chemistry