

High-Performance And Quantum Computing In Cancer Modeling: A Review And Hybrid Hpc-Quantum Approach

Vaibhav Banga¹, Anay Sinhal², Amit Hirawat³, Amit Sinhal⁴, Aruna Sengar⁵, Arpana Sinhal^{6*}

¹ Nokia America Corporation, Sunnyvale, California, USA. vaibhav.banga@nokia.com [0009-0003-2831-1169],

² University of Florida, Gainesville, USA. sinhal.anay@ufl.edu [0009-0008-8328-2336],

³ Poornima University, Jaipur, India. amit.hirawat@poornima.edu.in [0000-0002-1411-7598],

⁴ JK Lakshmipat University, Jaipur, India. amit.sinhal@jkl.edu.in [0000-0001-6697-6995],

⁵ Sagar Institute of Research and Technology, Bhopal. Sengar.aruna1@gmail.com

^{6*} Manipal University Jaipur, Jaipur, India. arpana.sinhal@jaipur.manipal.edu [0009-0008-2058-7685]

Corresponding Author: Arpana Sinhal, Manipal University Jaipur, Jaipur, India. arpana.sinhal@jaipur.manipal.edu

Abstract: High-Performance Computing (HPC) and quantum computing are increasingly applied to model complex diseases such as cancer. This paper surveys recent research on the independent contributions of these technologies to cancer modeling, examines integration efforts, and proposes a novel hybrid approach. HPC has enabled large-scale, high-resolution simulations, including centimeter-scale tumor growth models, with significant speedups using GPU acceleration and distributed computing via the Message Passing Interface (MPI). Quantum computing, while nascent, shows promise in drug discovery and has improved predictive modeling with quantum machine learning, achieving up to a 14% higher Area Under the Receiver Operating Characteristic curve (AUROC) in mortality prediction. From this synthesis, we hypothesize that integrating quantum solvers as accelerators within HPC cancer simulations can reduce overall simulation time by $\geq 30\%$ while maintaining sub-1% error margins. We propose a practical hybrid refinement strategy, where a quantum solver provides an initial guess to a classical iterative solver, a method well-suited for noisy intermediate-scale quantum (NISQ) hardware. To ground this, we propose a novel hybrid approach and outline a long-term experimental roadmap to validate its potential. An experimental design is detailed to validate the hypothesis by coupling a quantum module with an existing HPC cancer simulator to measure speedup, accuracy, and scalability. This work provides a comprehensive perspective on how HPC and quantum computing can advance cancer modeling.

Keywords: High-performance computing (HPC); quantum computing; hybrid HPC-quantum integration; cancer disease modeling; tumor-growth simulation; quantum linear system solver; computational oncology; digital twins; multiscale simulation; accelerators.

1. Introduction

Cancer disease modeling, which encompasses tumor growth, treatment response, and metastatic progression, demands extreme computational power. Researchers have long employed High-Performance Computing (HPC) to simulate these complex processes, leveraging massive parallelism on CPUs and GPUs to approach realistic scales [1]. For instance, HPC-enabled simulations have achieved centimeter-scale 3D tumor growth with sub-millimeter resolution using GPU acceleration and Message Passing Interface (MPI) across cluster nodes. Such implementations allow the exploration of patient-specific tumor scenarios that were previously computationally prohibitive [2].

In parallel, recent advances in quantum computing have shown potential implications for oncology. Quantum computers leverage superposition and entanglement to explore solution spaces in fundamentally new ways, which could address key computational bottlenecks. Early quantum algorithms have been applied to drug discovery; notably, a

hybrid quantum-classical generative model designed novel KRAS inhibitor molecules, two of which showed promise as cancer therapeutics [3]. Quantum machine learning has also emerged, with a quantum support vector machine improving the prediction of mortality in early-onset colorectal cancer patients (AUROC ≈ 0.86) compared to its classical counterpart (AUROC ≈ 0.72) [4].

Despite these parallel tracks of progress, integrating HPC and quantum computing remains a frontier. The vision of hybrid HPC-quantum computing is to treat quantum processing units (QPUs) as specialized accelerators within HPC architectures. This concept, termed "quantum-centric supercomputing," combines the strengths of both platforms: HPC manages large-scale data and the bulk of a simulation, while a quantum co-processor tackles a particularly hard sub-problem [5].

Early demonstrations provide proof-of-concept. Microsoft and Quantinuum demonstrated an end-to-end chemistry simulation using HPC resources integrated with 12 logical qubits, achieving chemical accuracy in a catalysis problem. This work, detailed in a pre-print publication, combined HPC tools for reaction path analysis (AutoRXN and AutoCAS) with a logical-qubit quantum computation for the ground state energy of the active site, showcasing a complete hybrid workflow [6]. Similarly, as reported by the company, a hybrid HPC-quantum workflow at Qubit Pharmaceuticals reduced the required qubits for a molecular property calculation from over 250 to 32 by partitioning the problem between classical and quantum computations [7,13]. While the specific technical method is not detailed in a peer-reviewed paper, such a reduction is plausibly achieved through techniques like active space selection, where only the most critical molecular orbitals are handled by the QPU.

A key insight for near-term applications is the use of quantum algorithms to augment, rather than replace, classical solvers. Given the noisy nature of current quantum hardware, a quantum-computed solution may contain errors. However, this approximate solution can serve as a high-quality initial guess or a "preconditioner" for a classical iterative solver. This "hybrid refinement" strategy leverages the potential speed of a quantum algorithm to get close to the solution quickly, while relying on the robustness of a classical algorithm to refine the result to the required accuracy. This approach is particularly pragmatic for the current Noisy Intermediate-Scale Quantum (NISQ) era.

This paper's objective is threefold. First, we critically assess how HPC and quantum computing have independently advanced cancer modeling. Second, we evaluate attempts at combining both technologies. Third, we formulate an integrated hypothesis for a hybrid HPC-quantum approach in cancer modeling and propose a strategy to test it. By surveying state-of-the-art examples (summarized in Table 1) and synthesizing their findings, we ground our proposed approach in evidence, aiming to spur the development of hybrid methods that can handle the growing complexity of cancer simulations more efficiently.

2. Background

Cancer Modeling and Computational Demands

Cancer is a multi-scale process, spanning molecular interactions to macroscopic phenomena like tumor morphology. Mathematical models often involve systems of partial differential equations (PDEs) for tissue-scale behavior or agent-based models for cell-level interactions [8]. Simulating 3D tumor growth with sub-millimeter resolution involves solving PDEs on millions of grid points over thousands of time steps [1]. Computational cost scales rapidly with model fidelity; for example, a recent study reported that running 300 instances of a complex metastatic cancer model sequentially would take ≈ 764 days, a time reduced to 2.5 days using HPC, illustrating a $>300\times$ speedup [8]. While solving large linear systems from discretized PDEs is a primary bottleneck in such con-

tinuum models, other paradigms, such as large-scale agent-based models, may face different constraints, for instance, in neighbor-finding algorithms or handling massive event queues. Our work focuses on the former.

High-Performance Computing (HPC) in Cancer Research

HPC uses powerful processors (CPUs, GPUs) in parallel to achieve performance far beyond a single machine. Software frameworks like MPI and OpenMP enable distributed and shared-memory parallelism, respectively. HPC has made substantial contributions to cancer modeling. Goodin et al. (2021) employed an MPI+CUDA implementation to simulate a 3D continuum tumor of volume $\approx 2.66 \text{ cm}^3$ at $31.25 \text{ }\mu\text{m}$ resolution, a scale unattainable without HPC. By parallelizing across 64 GPUs, they achieved a $6.7\times$ speedup over a single-GPU run [1]. HPC also enables ensemble simulations; Rocha et al. (2024) ran 300 stochastic realizations of a digital twin model in parallel on an HPC cluster, completing in 2.5 days what would have taken over two years sequentially [8]. However, classical HPC methods face challenges, including diminishing returns with scaling due to communication overhead.

Fig. 1. Schematic comparison of classical HPC vs. quantum computing platforms. Classical HPC systems (left) use a network of CPU and GPU nodes for parallelized simulations. Quantum computers (right) use qubits and quantum gates to perform computations. In hybrid architectures, QPUs are interfaced with HPC clusters to offload specific computations [5,6].





Fig. 2. Performance scaling comparisons. (Left) HPC strong-scaling: Speedup versus GPU nodes for a parallel tumor simulation. Ideal linear scaling is approached at low node counts, but communication overhead causes sub-linear speedup at higher counts. (Right) Theoretical complexity scaling for solving a linear system of size N . Classical iterative solvers can scale as $O(N^k)$, while quantum algorithms may have asymptotically better polynomial scaling (e.g., $O(\sqrt{N})$ or $\text{polylog}(N)$) [9].




Quantum Computing Principles and Relevance to Cancer

Quantum computing uses quantum bits (qubits) to process information in fundamentally new ways, exploiting superposition and entanglement. In cancer research, two areas stand out: quantum simulation of chemistry and quantum machine learning. Many processes in cancer, like drug-target interactions, are quantum mechanical. Simulating these is costly for classical computers. Quantum computers offer a more direct route. For instance, Kassal & Tan (2025) performed the first quantum simulation of real-time molecular dynamics, simulating how a molecule behaves when excited by light, a process pertinent to skin cancer research. Published in the *Journal of the American Chemical Society*, their work used a trapped-ion quantum computer to track the ultrafast dynamics of photo-excited molecules like allene and pyrazine, a task that is extremely challenging for classical HPC [10].

Current Limitations of Quantum Computing

Today's quantum hardware is in the Noisy Intermediate-Scale Quantum (NISQ) era, with devices that cannot yet outperform supercomputers on most practical tasks. Realizing theoretical speedups often requires error-corrected logical qubits, which may need thousands of physical qubits per logical one [6,11]. A key bottleneck in many cancer models is solving large linear systems. Quantum linear system algorithms (QLSAs) like HHL offer exponential speedup in theory but have strict requirements. Recent work by Golden et al. (2022) addressed one such issue by introducing a quantum preconditioner for ill-conditioned systems, improving the condition number scaling from $O(N)$ to $O(\sqrt{N})$ [9]. This suggests a path toward making quantum solvers practical for scientific simulations by reducing the number of required iterations. Given these limitations, hybrid quantum-classical approaches are actively being explored, where a quantum computer tackles a specific sub-problem where it has an advantage, while classical HPC handles the rest [11].

Fig. 3. Workflow diagram of a hybrid HPC-quantum computing pipeline for cancer simulation. Classical HPC handles the main simulation loop. At each time step, a challenging sub-problem (e.g., solving a linear system $Ax=b$) is offloaded to a quantum solver. The QPU computes an approximate solution, which is returned to the classical simulation, for instance, as an initial guess for a classical solver.



3. Literature Review

To assess the state of the art, we surveyed key publications applying HPC or quantum computing to cancer modeling. Table 1 summarizes representative peer-reviewed works and significant pre-prints, highlighting their approach, outcomes, and limitations.

Table 1. Comparison of Approaches in Cancer Modeling Studies

S . N o	Au- thors (Year)	Cancer Modeling Task	Computa- tional Ap- proach	Key Re- sults	Contributions / Limitations
1	Vakili et al. (2025) [3]	Drug design – KRAS inhibitors	Quantum-classical generative modeling (QCBM)	Designed 15 novel molecules; 2 showed strong binding to KRAS.	Contribution: First demonstration of a quantum algorithm leading to experimental drug candidates in oncology. Limitation: Did not show a clear speedup; value was in the quality/novelty of results.
2	Ito, T. et al. (2025) [14] (arXiv pre-print)	Anti-cancer drug response prediction	Quantum-Classical Hybrid Machine Learning (QHML)	Found data normalization significantly enhances QHML model performance.	Contribution: Highlights a critical pre-processing step for improving QHML efficacy. Limitation: Based on simulations, scalability needs further benchmarking.
3	Kassal & Tan (2025) [10]	Quantum dynamics simulation (photochemistry relevant to skin cancer)	Trapped-ion quantum computer simulating real-time molecular dynamics	Simulated a molecule's bond breaking under UV excitation, a first for a real molecule on a quantum device.	Contribution: Showed NISQ devices can tackle time-dependent quantum chemistry problems challenging for HPC. Limitation: Only very small molecules simulated; offers new capability, not yet speed or scale.

4	Khan et al. (2024) [15] (arXiv pre-print)	Brain tumor classification from MRI data	Quantum Convolutional Neural Networks (QCNNs)	Achieved 99.67% classification accuracy on a test dataset.	Contribution: Demonstrated potential of QCNNs in medical imaging. Limitation: Study on a limited dataset; requires further validation.
5	Streviniotis et al. (2025) [12]	In-silico drug trial simulations	Bayesian optimization to allocate HPC resources for 2500+ tumor simulations	Optimized core-hour usage, reducing overall HPC cost by ~20–30%.	Contribution: Introduced an HPC workflow optimizer for large-scale virtual trials. Limitation: Optimizes scheduling, does not speed up individual simulations.
6	Yu et al. (2024) [4]	Mortality prediction in colorectal cancer	Quantum Support Vector Machine (QSVM) vs. classical SVM	QSVM achieved AUROC 0.863, outperforming classical SVM's 0.723 (~14% improvement).	Contribution: Showed QML can improve prediction on limited clinical data. Limitation: Used a quantum-inspired simulator (no physical speedup); computational cost was higher.
7	van Dam, W. et al. (2024) [6] (arXiv pre-print)	Integrated HPC+Quantum workflow (catalyst simulation)	Hybrid workflow: HPC (Auto-RXN) + 12 logical qubit quantum computation + AI	Achieved chemical accuracy in estimating catalyst reaction energy.	Contribution: Landmark hybrid demonstration orchestrating HPC, AI, and quantum computing. Limitation: A proof-of-concept on a problem that fits in 12 logical qubits; focus was on accuracy, not speed.
8	Rocha et al. (2024) [8]	Multiscale immune-tumor interactions (digital twin)	HPC-based ensemble (300 agent-based simulations in parallel)	Completed ensemble in 2.5 days vs. ~764 days serially (~305× speedup).	Contribution: Showcased ensemble HPC to quantify uncertainty in cancer progression. Limitation: Extremely compute-intensive, highlighting classical scalability limits.
9	Golden et al. (2022) [9]	Solver sub-routine acceleration (linear system for fracture networks)	Quantum inverse Laplacian preconditioner for QLSA	Improved condition number scaling from $O(N)$ to $O(\sqrt{N})$ for	Contribution: Provided a theoretical path to quantum-accelerate a core HPC task (iterative linear

				sparse systems.	solvers). Limitation: Theoretical/simulation-based; requires fault-tolerant hardware to realize advantage.
10	Rhodes, C. (2022) [13] (Industry Blog)	Drug discovery for cancer	Hybrid HPC + AI + Quantum approach (NVIDIA QODA)	Reduced required qubits from 250 to 32 for a molecular calculation.	Contribution: Pioneered a pragmatic hybrid pipeline demonstrating resource reduction. Limitation: Results from company announcements, not a peer-reviewed publication. The technical method for qubit reduction is not detailed, but is likely based on techniques such as active space selection.
11	Goodin & Frieboes (2021) [1]	Continuum 3D tumor growth simulation	MPI + CUDA parallelization (multigrid solver on GPUs)	Achieved 6.7× speedup using 64 GPUs vs. 1 GPU.	Contribution: Enabled patient-scale tumor simulations in reasonable time. Limitation: Scaling efficiency drops at high node counts due to communication bottlenecks.

Analysis of the Literature

The surveyed studies reveal distinct but complementary roles for advanced computing in oncology. HPC excels at brute-force simulation and parallel exploration, enabling breakthroughs in scale and statistical power [1, 8, 12]. However, its efficiency is often limited by algorithmic complexity and communication bottlenecks, especially as problem sizes grow.

Quantum computing's value, thus far, has been more qualitative. Studies in drug design and molecular modeling demonstrate that quantum methods can explore chemical space in novel ways or simulate phenomena that classical methods can only approximate. The accuracy improvements seen in quantum machine learning also underscore this point [4, 15, 16, 17]. These are not yet runtime speedups but enhancements in capability or result quality.

Integration efforts are nascent but promising. The Microsoft catalyst simulation and the workflow described by Qubit Pharmaceuticals are concrete steps toward HPC-quantum integration. In both cases, classical computing performed the bulk of the work, but a critical piece was offloaded to a quantum approach. This either reduced the overall resources needed or achieved a higher level of accuracy [4, 6]. These examples guide a pragmatic integration strategy: identify a sub-problem that is a bottleneck for HPC and deploy a targeted quantum algorithm. Literature clearly points to solving large, sparse

linear systems in continuum models as a prime candidate, justifying our subsequent hypothesis.

4. *Proposed Hypothesis*

Drawing from the literature, we propose the following hypothesis:

Integrating a quantum linear solver as a real-time accelerator into an MPI-based finite element tumor growth simulation will reduce the overall computation time by at least 30%, while maintaining accuracy within 1% of the classical solver.

By offloading the computationally intensive step of solving large sparse linear systems to a quantum algorithm, the simulation can run significantly faster without compromising fidelity.

Rationale

In large-scale tumor simulations, a substantial fraction of runtime is spent solving linear equations of the form $Ax=b$ that arise from discretizing PDEs. Goodin et al. reported that even with GPU acceleration, their multigrid linear solves were a major bottleneck. Quantum algorithms, such as the preconditioned QLSA described by Golden et al., could theoretically solve these systems with better asymptotic scaling [9]. For a very large system, this offers dramatic improvement.

While near-term quantum solvers will not achieve full theoretical speedups due to overheads, even a partial acceleration is valuable. We conservatively target a $\geq 30\%$ reduction in total runtime. Our hypothesis also emphasizes maintaining accuracy. A quantum solver's result can be used as an initial guess for a few iterations of a classical solver, ensuring the final solution meets a strict error tolerance ($\leq 1\%$) while still benefiting from the initial quantum speedup. This hybrid refinement approach is robust against noise in current quantum hardware.

Performance Modeling

To provide a quantitative basis for our hypothesis, we model the potential speedup. Consider a tumor simulation with $N=10^6$ degrees of freedom. A classical iterative solver like Conjugate Gradient may require a number of iterations proportional to the condition number κ , with each iteration costing $O(N)$ for a sparse matrix. A quantum algorithm enhanced with a preconditioner, such as the one proposed by Golden et al., could improve the convergence, effectively reducing the number of iterations to scale as $O(\sqrt{\kappa})$. More generally, quantum linear solvers like HHL have a complexity of $O(\kappa \text{ polylog}(N))$, an exponential improvement in N over classical methods, though with a polynomial dependence on κ [9].

Let us assume the linear solve constitutes $P=0.5$ (50%) of the total classical runtime. If a quantum approach provides a modest speedup factor of $S_{\text{portion}}=5\times$ for this part of the calculation (including all overheads), Amdahl's Law gives the total potential speedup S_{total} as:

$$S_{\text{total}} = \frac{1}{(1-P) + \frac{P}{S_{\text{portion}}}} = \frac{1}{(1-0.5) + \frac{0.5}{5}} \approx 1.67\times$$

This corresponds to a percentage time reduction of $(1 - 1/S_{\text{total}}) \times 100\% \approx 40\%$. This simple model demonstrates that a $\geq 30\%$ overall speedup is plausible even if the quantum module provides only a moderate acceleration for a significant portion of the workload.

5. *Experimental Design*

To evaluate the proposed hybrid approach, we outline a stepwise experimental study. The experiment will compare a classical HPC simulation (baseline) with a hybrid simulation where the linear solver is offloaded to a quantum co-processor.

Communication Overhead Analysis

A critical factor for tightly-coupled hybrid algorithms is the communication overhead between the HPC system and the QPU. The total time for a quantum call, $T_{\text{quantum_call}}$, can be modeled as:

$$T_{\text{quantum_call}} = T_{\text{comm_send}} + T_{\text{queue}} + T_{\text{exec}} + T_{\text{comm_recv}}$$

where T_{comm} is the network communication time, T_{queue} is the job queuing time on the QPU, and T_{exec} is the quantum execution time. The communication time itself is $T_{\text{comm}} = T_{\text{latency}} + (\text{DataSize}/\text{Bandwidth})$.

For an iterative solver, this overhead is incurred at every step. Real-world network performance varies dramatically with architecture. For co-located systems within the same data center using high-speed interconnects like InfiniBand, latency can be as low as $T_{\text{latency}} \approx 2 \mu\text{s}$ with a bandwidth of up to 400 Gbps. In contrast, remote, cloud-based QPUs accessed over a Wide Area Network (WAN) can have latencies of 2–50 ms or more.

The data size for sending the problem $Ax=b$ to the QPU scales with the problem size. For an $N \times N$ sparse matrix with an average of s non-zero entries per row, the data size to describe the matrix scales as $O(sN)$. For an iterative simulation where each step may take only milliseconds on the HPC side, a round-trip latency of tens of milliseconds for a remote QPU call would completely negate any potential quantum speedup. This analysis indicates that for real-time accelerator use cases, physical co-location of HPC and quantum resources with low-latency interconnects is a prerequisite for achieving a performance advantage [5].

Matrix Representation for Quantum Processing

To solve the linear system $Ax=b$ on a quantum computer, the classical matrix A must be encoded into a quantum state or operation. Since the matrix A arising from PDE discretizations is generally not unitary, it cannot be directly implemented as a quantum gate. Instead, it must be embedded into a larger unitary operator U using a technique known as block encoding [11].

A block encoding represents A (scaled by a subnormalization factor) as the top-left block of a larger unitary matrix U that acts on a combined system of primary and ancillary qubits. This unitary U can then be used as a subroutine within higher-level quantum algorithms. For example, in QLSAs like HHL or approaches based on the Quantum Singular Value Transformation (QSVT), the block-encoded U is used to perform Hamiltonian simulation of an operator related to A , allowing the algorithm to access and manipulate its spectral properties (eigenvalues or singular values) to construct the solution state $|x\rangle$. [11] This approach effectively transforms the classical matrix problem into a form amenable to quantum processing.

Implementation and Validation

We will implement a representative 3D reaction-diffusion tumor growth model, discretized using finite elements. The initial small-scale (e.g., $8 \times 8 \times 8$ grid) experiment is not intended to demonstrate a quantum speedup, which is infeasible on current hardware for such problems. Instead, its objectives are to: (1) validate the end-to-end software integration between the HPC simulation and the quantum hardware API; (2) obtain empirical measurements of the full communication round-trip time, including network la-

tency and QPU job queuing time; and (3) test and calibrate error mitigation and hybrid refinement strategies on a physical NISQ device.

In this small-scale test, the quantum solution $|x_q\rangle$ will be used as an initial guess for a classical iterative solver (e.g., GMRES or Conjugate Gradient). This hybrid refinement strategy ensures the final solution meets the required accuracy tolerance, combining potential quantum acceleration with classical reliability.

Metrics for validation will include:

Total simulation time: Comparing the classical baseline against the hybrid approach.

Time breakdown: Profiling communication overhead versus computation time.

Accuracy: Comparing tumor growth curves and concentration fields between baseline and hybrid runs, quantifying differences to ensure they are $<1\%$.

Scalability: Testing how the hybrid approach scales with the number of HPC nodes.

The hypothesis will be confirmed if, for a sufficiently large problem size where a quantum advantage is theoretically plausible, the hybrid simulation runs at least 30% faster than the classical counterpart with results differing by less than 1%.

Fig. 4. Proposed experimental procedure for the hybrid HPC-quantum approach.





Fig. 5. Hypothetical target outcome of HPC-quantum integration.



This chart illustrates the target performance gain outlined in our hypothesis, not a projection from existing data. A baseline HPC simulation taking 100 hours (blue bar) might

be reduced to ~70 hours (orange bar) with the addition of a quantum solver, representing a 30% reduction in time-to-solution.

6. Conclusion

HPC has become a cornerstone of cancer modeling, enabling simulations at unprecedented scales. Concurrently, quantum computing is emerging as a powerful tool for specific computational challenges, from drug design to predictive modeling. This review has shown that while HPC provides scale through parallelism, it is constrained by classical algorithmic complexity. Quantum computing, though currently limited by hardware, offers new pathways to efficiency for certain problem classes.

Our assessment indicates that a hybrid HPC-quantum paradigm is not only feasible but necessary to continue accelerating progress. Early demonstrations in related fields validate that end-to-end workflows employing both resources can work in practice [6, 7]. Inspired by these, we have proposed a specific, testable hypothesis focused on speeding up high-fidelity tumor simulations. By integrating a quantum solver into the simulation loop, we aim to reduce time-to-solution by at least 30% without sacrificing accuracy.

A key takeaway is the practicality of a hybrid refinement strategy, where a noisy quantum solution serves as a high-quality starting point for a classical solver. This approach mitigates the limitations of current NISQ hardware while still capturing a potential quantum advantage. Our analysis also underscores the critical importance of low-latency interconnects, suggesting that co-location of HPC and quantum resources is essential for tightly-coupled accelerator models [5].

If our hypothesis is validated, the implications are significant. A 30% simulation speedup could shorten the time for personalized treatment planning from days to overnight. It would enable more robust uncertainty analyses within fixed time budgets, improving the reliability of model predictions. Success in this integration would motivate hybrid approaches for other areas of cancer research, such as genomics and pathology image analysis. The convergence of HPC and quantum computing represents a promising frontier for computational oncology. By building on the incremental successes in the literature, we can chart a path toward synergistic platforms that keep pace with the growing complexity of cancer research, ultimately translating computational advancements into faster scientific discoveries and improved patient outcomes.

7. Declarations

Funding: No external funding was received for this study.

Competing interests: The authors declare no competing interests.

Author contributions: Anay Sinhal performed the literature survey, developed the hybrid HPC-quantum hypothesis and drafted the manuscript. Dr. Arpana Sinhal supervised the overall study, provided the oncology perspective and critically reviewed the draft. Dr. Amit Sinhal refined the HPC methodology, shaped the experimental design and revised the manuscript. All authors read and approved the final manuscript.

Data availability: No new datasets were generated or analysed. All information discussed is taken from the published sources cited in the References.

References

1. Goodin, D. A., & Frieboes, H. B. (2021). Simulation of 3D centimeter-scale continuum tumor growth at sub-millimeter resolution via distributed computing. *Computers in Biology and Medicine*, 134, 104507.
2. High-Performance Computing Takes Aim at Cancer. (2016). Science and Technology Review. Lawrence Livermore National Laboratory.

3. Ghazi Vakili, M., Gorgulla, C. et al. (2025). Quantum-computing-enhanced algorithm unveils potential KRAS inhibitors. *Nature Biotechnology*, 1-6.
4. Yu, J. Y., Sim et al. (2024). Evaluation of conventional and quantum computing for predicting mortality based on small early-onset colorectal cancer data. *Applied Soft Computing*, 162, 111781.
5. Alexeev, Y., et al. (2024). Quantum-centric supercomputing for materials science: A perspective on challenges and future directions. *Future Generation Computer Systems*, 160, 666-710.
6. van Dam, W., Liu, H., Low, G. H., et al. (2024). End-to-end integration of HPC, reliable quantum computing, and AI for chemical reactions. *arXiv preprint arXiv:2409.05835*.
7. Qubit Pharma & Sorbonne Cut Qubit Needs for Molecular Simulation. (2024). Qubit Pharmaceuticals Blog.
8. Rocha, A., Aguilar, B., Getz, M., Shmulevich, I., & Macklin, P. (2024). A multiscale model of immune surveillance in micrometastases gives insights on cancer patient digital twins. *npj Systems Biology and Applications*, 10(1).
9. Golden, J., & Viswanathan, H. (2022). Quantum computing and preconditioners for hydrological linear systems. *Scientific Reports*, 12(1), 1-10.
10. Navickas, T., Tan, T., Chan, H., et al. (2025). Experimental quantum simulation of chemical dynamics. *Journal of the American Chemical Society*.
11. Chong, F. T., Franklin, D., & Martonosi, M. (2017). Programming languages and compiler design for realistic quantum hardware. *Nature*, 549(7671), 180-187.
12. Streviniotis, E., Giatrakos, N., Kotidis, Y., Ntiniakou, T., & Ponce de León, M. (2025). RATS: A Resource Allocator for Optimizing the Execution of Tumor Simulations over HPC Infrastructures. *Information Systems*.
13. Rhodes, C. (2022). Qubit Pharmaceuticals Accelerates Drug Discovery With Hybrid Quantum Computing. *NVIDIA Blog*.
14. Ito, T., Artem, L., & Tsunoda, T. (2025). Optimal normalization in quantum-classical hybrid models for anti-cancer drug response prediction. *arXiv preprint arXiv:2505.10037*.
15. Khan, M., Innan, N., Galib, A., & Bennai, M. (2024). Brain Tumor Diagnosis Using Quantum Convolutional Neural Networks. *arXiv preprint arXiv:2401.15804*.
16. Repetto, V., Ceroni, E. G., Buonaiuto, G., & D'Aurizio, R. (2024). Quantum enhanced stratification of Breast Cancer: exploring quantum expressivity for real omics data. *arXiv preprint arXiv:2409.14089*.
17. Saggi, M. K., Bhatia, A. S., Isaiah, M., Gowher, H., & Kais, S. (2024). Multi-Omic and Quantum Machine Learning Integration for Lung Subtypes Classification. *arXiv preprint arXiv:2410.02085*.