Quantum Machine Learning and Radiation Oncology

Sakshi Kumar, Mrittunjoy Guha Majumdar

Quantum Machine Learning and Radiation Oncology

Sakshi [1] and Mrittunjoy Guha Majumdar [1,2,]

March 11, 2026

[1] School of Artificial Intelligence, Amrita Vishwa Vidyapeetham, Delhi NCR
[2] School of Natural Sciences and Engineering, National Institute of Advanced Studies, Bengaluru

Abstract

Quantum machine learning (QML) leverages quantum phenomena such as superposition, entanglement, and interference to address computational challenges that are difficult for classical systems. In radiation oncology, where treatment planning, image guidance, adaptive replanning, and multi-modal data integration require high-dimensional optimisation and complex pattern recognition, QML may offer advantages in computational efficiency and representational capacity. This review systematically maps quantum algorithms to the clinical radiation oncology workflow. We outline foundational concepts in quantum computing, including complexity insights from the BQP class and practical constraints of the noisy intermediate-scale quantum (NISQ) era. The radiation oncology pipeline—from consultation and simulation to treatment planning, delivery, and follow-up—is analysed to identify computational bottlenecks where quantum methods could provide benefit. Current developments in quantum-enabled diagnostic modelling, quantum-inspired clinical decision support for precision radiotherapy, quantum reinforcement learning for adaptive treatment policies, and emerging concepts such as quantum digital twins are reviewed. Key challenges—including hardware limitations, barren plateaus, data encoding constraints, limited datasets, regulatory considerations, and the gap between theoretical speedups and clinical implementation—are critically examined. Finally, a three-horizon roadmap is proposed outlining near-term proof-of-concept research, mid-term integration with treatment planning systems, and long-term prospects for fault-tolerant quantum simulation in oncology.

Keywords. Quantum machine learning; Radiation oncology; Quantum algorithms; Treatment planning optimisation; Adaptive radiotherapy

1 Introduction

Radiation oncology is a data-intensive discipline that relies on complex, multi-step workflows encompassing patient simulation, treatment planning, image guidance, and adaptive replanning [1, 2]. Each step involves computationally demanding tasks: high-resolution image registration, multi-objective optimisation under uncertainty, Monte Carlo dose calculation, and extraction of high-dimensional radiomic features [3, 4, 5]. Classical machine learning (ML) has made significant inroads into automating and improving these tasks, yet it faces fundamental limitations when confronted with the scale, dimensionality, and inherent quantum nature of the underlying physical processes [6]. Quantum computing offers a fundamentally different paradigm, leveraging superposition, entanglement, and interference to explore solution spaces that are inaccessible to classical architectures [7]. In the near term, noisy intermediate-scale quantum (NISQ) devices [8] are being combined with classical algorithms in hybrid quantum-classical frameworks that have already shown promise for optimisation, simulation, and machine learning tasks [9]. The field of quantum machine learning (QML) aims to harness these advantages by embedding classical data into exponentially large Hilbert spaces, potentially capturing correlations that escape classical models [10, 11].

In radiation oncology, early theoretical and experimental work has explored QML for beam angle optimisation [12], toxicity prediction [13], adaptive treatment planning [14], and multi-omics integration [15]. However, a comprehensive review that systematically maps the landscape of quantum algorithms to the specific computational challenges of the radiation oncology workflow is still lacking. This article fills that gap by:

- Providing a rigorous taxonomy of quantum algorithms relevant to oncology, distinguishing between fault-tolerant and NISQ-era approaches.
- Detailing the radiation oncology workflow and pinpointing where quantum methods could offer transformative advantages.
- Reviewing the state of the art in quantum-enhanced diagnostic modelling, clinical decision support, adaptive radiotherapy, and digital twin concepts.
- Critically assessing the current challenges, limitations, and future directions for translating QML into clinical practice.

By synthesising knowledge from quantum information science and radiation oncology, we aim to provide a foundation for interdisciplinary research and to guide the community toward realistic, high-impact applications of quantum computing in cancer care.

2 Background and Foundations

Radiation oncology requires complex treatment planning that balances tumor control with normal tissue protection. Classical multivariate optimization, such as Pareto-based convex planning, has improved planning time and quality while providing a baseline to evaluate emerging computational approaches [16]. Classical Machine Learning (ML) methods, including neural networks, SVMs, have been applied to integrate high-dimensional imaging and clinical data but face limitations due to small sample sizes and non-linear interactions. Quantum Machine Learning (QML) leverages superposition, entanglement, and probabilistic measurement to represent complex data more efficiently. Quantum-enhanced feature spaces can embed classical data into exponentially large Hilbert spaces, enabling better pattern recognition and potentially better handling of high-dimensional clinical datasets [11, 17].

2.1 Quantum Computing and Machine Learning in Healthcare

QML operates with the formalism of quantum mechanics, encoding information in qubits that evolve in Hilbert spaces under unitary transformations [7]. Unlike classical bits, qubits can exist in superposition, representing multiple states simultaneously, and can exhibit entanglement, which introduces non-classical correlations essential for enhancing representational power [18]. Classical data first must be embedded into quantum states via angle or amplitude encoding, creatung quantum featue maps that oreserve information in the Hilbert space. Once encoded Variational Quantum Circuits (VQCs) are constructed from parameterized unitary gates process the data, with model outputs obtained through projective measuremntsproducing probabilistic predictions that requre repeat sampling for stability [7]. Practical implementations are constrained by noise and decoherence inherent in current quantum hardware. In the Noisy Intermediate-Scale Quantum (NISQ) era, these limitations make hybrid quantum-classical approaches, such as variational classifiers adn quantum kernels, the most feasible frameworks for early applications in healthcare [8].

QML offers several paradigms to exploit the representational and optimization advantages of quantum computing. One approach involves quantum kernels where classical data is embedded via parameterized quantum circuits into exponentially large Hilbert spaces which creates implicit feature maps that can enhance classification performance on high dimensional datasets [10]. Variational Quantum Circuits (VQCs) represent a near-term QML framework consisting of parameterized quantum circuits trained using classical optimization algorithms. These models can learn complex decision boundaries, making them suitable for supervised tasks such as disease classification or patient stratification in healthcare [19]. Quantum Convolutional Neural Networks (QCNN) extends this paradigm by leveraging hierarchical entanglement and parameter sharing to mimic classical CNN feature extraction, offering a powerful tool for structured

healthcare data such as high-resolution medical images [20]. Finally, hybrid optimization strategies inspired by the Quantum Approximate Optimization Algorithm (QAOA) demonstrate how parametrized quantum circuits combined with classical optimization can tackle combinatorial and learning problems, providing a flexible framework for quantum-enhanced predictive modeling in healthcare [9].

Healthcare data is inherently high-dimensional, heterogeneous, and multi-modal, presenting challenges that classical machine learning often struggles to address. Radiomics which extracts quantitative image features to capture intra-tumor heterogeneity that represent the type of complex data that can benefit from enhanced representational frameworks [21]. Integrating imaging with genomic and clinical variables has shown improved predictive modeling in oncology, but this requires methods capable of capturing intricate multi-modal relationships [22]. Classical models in radiation oncology such as neural networks and other supervised methods, often face overfitting and poor generalization due to limited patient data and high dimensional features [23]. Oncology prediction is further complicated by uncertainty and variability in multi-omics datasets [14]. Deep learning has advanced biomedical imaging and genomics by extracting hierarchical non-linear feature, setting a benchmark for representational power [6]. QML provides a potentially transformation alternative which offers access to exponentially larger Hilbert spaces, entanglement-enabled correlations and hybrid quantum-classical optimization , all of which may improve representation, learning and prediction in high -dimensional multi-modal healthcare datasets.

2.1.1 Complexity Theoretic Understanding

To rigorously assess the potential of quantum machine learning (QML) in radiation oncology, we must ground the discussion in quantum computational complexity theory. This framework provides formal language to describe when quantum advantage is possible, the conditions under which it can be realized, and the fundamental obstacles that remain even with ideal hardware. The class BQP (Bounded-error Quantum Polynomial time) comprises decision problems solvable by a uniform family of polynomial-size quantum circuits with error probability at most 1/3 for all instances. Formally, a language ($L \subseteq \{0,1\}^)$ is in BQP if there exists a polynomial-time uniform family of quantum circuits $\{Q_n\}_{n\in\mathbb{N}}$ such that for all inputs x:

$$x \in L \implies \Pr(Q_{|x|} \text{ accepts } x) \geq \frac{2}{3}, \quad x \notin L \implies \Pr(Q_{|x|} \text{ accepts } x) \leq \frac{1}{3}.$$

The relationship between BQP and classical complexity classes is characterized by the inclusions BPP $\subseteq$ BQP $\subseteq$ PP $\subseteq$ PSPACE, where BPP is classical bounded-error probabilistic polynomial time, PP is probabilistic polynomial time with unbounded error, and PSPACE is polynomial space. It is widely conjectured that BPP $\subsetneq$ BQP, implying the existence of problems solvable in polynomial time on quantum computers that require superpolynomial time on

classical machines [24]. However, it remains open whether BQP $\subseteq$ NP or vice versa. For radiation oncology, this formalism is relevant insofar as clinical tasks can be reduced to decision problems within BQP. For instance, inverse treatment planning often requires solving large systems of linear equations $A\mathbf{x} = \mathbf{b}$ where A is the dose-influence matrix. If A is sparse and well-conditioned, the Harrow-Hassidim-Lloyd (HHL) algorithm solves this system in $O(\log(N)\kappa^2/\epsilon)$ time [25], where N is the matrix dimension, κ the condition number, and ϵ the precision. Classical algorithms scale as $O(N\log(N))$ under the same assumptions, so HHL offers an exponential speedup in N. However, this advantage is contingent on the ability to load $\mathbf{b}$ efficiently (a nontrivial input bottleneck) and to extract the solution vector $\mathbf{x}$ without exponential overhead. Moreover, for dense or ill-conditioned matrices, the speedup vanishes.

A more subtle potential advantage lies in quantum kernel methods. Consider a quantum feature map $\phi : \mathcal{X} \to \mathcal{H}$ that embeds classical data into an n-qubit Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$. The quantum kernel is defined as

$$k(x, x') = |\langle\phi(x)|\phi(x')\rangle|^2 = |\mathrm{Tr}\left[|\phi(x)\rangle\langle\phi(x)| \cdot |\phi(x')\rangle\langle\phi(x')|\right]|^2 .$$

For certain feature maps (e.g., those based on instantaneous quantum polynomial time circuits), estimating $k(x, x')$ to within additive error ϵ is #P-hard under polynomial-time reductions [17]. Formally, if there existed a classical randomized algorithm that, for any x, x', outputs an estimate $\tilde{k}$ such that $|\tilde{k} - k(x, x')| \leq \epsilon$ in time $\mathrm{poly}(n, 1/\epsilon)$, then BQP $\subseteq$ BPP would follow, collapsing the polynomial hierarchy. This provides a theoretical guarantee that quantum kernels can, in principle, capture correlations that are classically intractable to compute. In radiomics, where one extracts thousands of quantitative features from medical images, the ability to compute such kernels could reveal nonlinear interactions among features that are invisible to classical kernel methods. However, this advantage is only meaningful if the clinical prediction task requires the evaluation of such hard kernels. Moreover, on NISQ devices, the estimation of these kernels is subject to shot noise: with M measurements, the empirical kernel $\hat{k}$ satisfies $|\hat{k} - k| \leq O(1/\sqrt{M})$ with high probability, so achieving precision ϵ requires $M = \Omega(1/\epsilon^2)$ circuit repetitions. Thus, the asymptotic advantage may be offset by sampling overhead in practice.

Variational quantum algorithms, which form the backbone of most near-term QML proposals, optimize a parameterized quantum state $|\psi(\boldsymbol{\theta})\rangle$ to minimize a cost function $C(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|H|\psi(\boldsymbol{\theta})\rangle$. The trainability of such circuits is governed by the variance of the gradient. For a sufficiently expressive ansatz that forms a unitary 2-design, the partial derivative with respect to any parameter θ_μ satisfies

$$\mathbb{E}_{\boldsymbol{\theta}}\left[\frac{\partial C}{\partial \theta_\mu}\right] = 0, \quad \mathrm{Var}_{\boldsymbol{\theta}}\left[\frac{\partial C}{\partial \theta_\mu}\right] \leq O\left(2^{-n}\right),$$

provided that H is a global observable (i.e., acts nontrivially on all qubits) [26]. This exponential decay of gradient variance—the *barren plateau*—implies that

for problems requiring many qubits, gradient-based optimization becomes infeasible, as the number of shots needed to estimate a nonzero gradient grows as $\Omega(2^n)$. More recently, the analysis has been extended to local observables and shallow circuits. For a local Hamiltonian $H = \sum_i H_i$ where each H_i acts on at most k qubits, the gradient variance can scale polynomially in n if the circuit depth is $O(\log n)$ and the ansatz respects locality [27]. This suggests that QML models for medical imaging, which must process high-dimensional data, should be designed with local cost functions and shallow, problem-inspired architectures to avoid barren plateaus. For example, in dose optimization, the cost function naturally decomposes into local terms (e.g., dose-volume constraints for individual voxels), which may preserve trainability.

From the perspective of statistical learning theory, quantum models do not inherently require fewer training examples than classical models. In the PAC (Probably Approximately Correct) framework, a concept class $\mathcal{C}$ is learnable by a quantum learner if there exists a quantum algorithm that, with probability at least $1 - \delta$, outputs a hypothesis h such that $\text{error}(h) \leq \epsilon$ using m examples. Arunachalam and de Wolf [28] showed that any concept class learnable by a quantum learner with m examples is also learnable by a classical learner with $O(m \log m)$ examples, up to constant factors. Formally, if a concept class has quantum sample complexity $m_Q(\epsilon, \delta)$, then its classical sample complexity $m_C(\epsilon, \delta)$ satisfies $m_C = O(m_Q \log m_Q)$. Thus, quantum advantage cannot arise from superior sample efficiency; it must come from representational power (i.e., the ability to express functions that classical circuits cannot approximate with a given number of parameters) or from computational speedups in training. This result has profound implications for radiation oncology, where datasets are often small ($N \sim 10^2 - 10^3$). QML cannot circumvent the fundamental need for high-quality, well-annotated data; any claim of improved accuracy on small datasets must be scrutinized for overfitting and should be accompanied by rigorous cross-validation and comparison to classical baselines.

Where quantum models may excel is in reducing the time required for training or inference. Consider a classical neural network with P parameters; training via backpropagation scales as $O(P^2)$ per epoch. A variational quantum circuit with P parameters may, in principle, be trained using the parameter-shift rule, which requires $O(P)$ circuit evaluations per gradient step, each evaluation requiring $O(1)$ shots. If each circuit evaluation is exponentially faster than the corresponding classical computation (e.g., because it leverages quantum parallelism to compute a kernel that would take classical $O(2^n)$ time), then a time advantage emerges. However, this is heavily contingent on the absence of barren plateaus, the efficiency of data encoding, and the overhead of measurement. In the noisy intermediate-scale quantum (NISQ) regime, additional complexity-theoretic considerations apply. The effective depth of implementable circuits is limited by decoherence: for a circuit of depth D with gate error rate η, the overall fidelity scales as $(1 - \eta)^D$. To maintain a non-negligible success probability, we require $D \ll 1/\eta$. This restricts the class of problems that can be

solved on NISQ devices to those with shallow circuits. Moreover, the overhead of quantum error correction, which would enable deeper circuits, is currently prohibitive for the problem sizes of interest in oncology.

Data loading presents another fundamental bottleneck. Amplitude encoding of an N-dimensional vector into an $n = \lceil \log_2 N \rceil$-qubit state requires coherent arithmetic operations that themselves scale as $\text{poly}(n)$ in depth, but may be impractical for $N \sim 10^6$ (e.g., a 1024×1024 image). Angle encoding, while simpler, uses one qubit per feature, limiting the dimensionality of data that can be processed with available qubits. Thus, even if a quantum algorithm offers an asymptotic advantage in problem size N, the constant factors and practical constraints of data encoding may dominate for clinically relevant scales.

Table 1: Complexity-Theoretic Summary for Key Clinical Tasks

Clinical Task	*Quantum Algorithm*	*Theoretical Advantage*	*Conditions f*
Inverse planning (linear systems)	HHL	Exponential in matrix dimension N	Sparse, well-
Deformable image registration	HHL / QPE	Exponential in grid size	Sparse defor
Dose calculation (Monte Carlo)	QAE	Quadratic in precision ϵ	Low-noise, d
Radiomic feature selection	Grover	Quadratic in number of features F	Oracle for fe
Multi-omics classification	Quantum kernel	Potential P-hard kernel evaluation	Kernel maps

Complexity theory provides both a rigorous foundation for optimism and a sobering framework for understanding limitations. The promise of exponential speedups for certain linear algebra and Monte Carlo tasks is real but contingent on problem structure, encoding efficiency, and hardware capabilities. The absence of universal advantage in sample complexity or PAC learning reminds us that QML is not a panacea for small datasets. Barren plateaus impose architectural constraints that must be respected in any practical QML model. For radiation oncology, the path forward lies in identifying subproblems that satisfy these stringent conditions—such as sparse, well-conditioned linear systems from deformable registration, or low-dimensional feature selection tasks—and in developing hybrid quantum-classical workflows that insulate the quantum core from the data loading and output extraction bottlenecks. Only through such targeted, complexity-aware design can the theoretical advantages of quantum computing be translated into tangible clinical impact.

2.2 The Radiation Oncology Workflow

The radiation oncology workflow represents a meticulously orchestrated, multi-step process that transforms a patient's initial consultation into a precisely delivered course of treatment. This workflow, standardized over the past three decades, comprises several key phases: patient selection and consultation, simulation, treatment planning, quality assurance, treatment delivery, and follow-up

care [1, 29]. Each phase involves distinct clinical expertise, specialized technologies, and computational methods, creating a complex pipeline where accuracy and precision are paramount. Understanding this workflow is essential for identifying where emerging computational approaches, such as quantum machine learning, might offer transformative potential.

2.2.1 Patient Selection and Consultation

The workflow begins with a comprehensive consultation, where the radiation oncologist reviews the patient's diagnostic imaging, pathology results, and overall clinical status to determine whether radiation therapy is an appropriate treatment option [30]. This multidisciplinary evaluation considers tumour characteristics, prior treatments, patient performance status, and potential synergies with surgery or systemic therapies. For patients with a history of prior radiation, this consultation becomes particularly complex, requiring careful assessment of cumulative dose risks and consideration of reirradiation as a specialized pathway [31]. The decision to proceed involves shared decision-making with the patient, including discussion of treatment goals, potential side effects, and the expected duration of therapy.

2.2.2 Simulation and Immobilisation

Once the decision to treat is made, the patient undergoes simulation—a critical planning procedure that establishes the geometric foundation for all subsequent steps [32]. During simulation, the patient is positioned on the CT table in the exact posture that will be reproduced for each daily treatment fraction. Immobilisation devices are custom-fabricated to ensure reproducibility: thermoplastic masks for head and neck patients, vacuum-locked body bags for pelvic or thoracic treatments, and knee supports or bite blocks as needed [33]. These devices restrict patient movement and enable millimetre-level positioning accuracy across the entire treatment course. A planning CT scan is acquired with the patient in this treatment position, providing a volumetric dataset that captures both the tumour region and surrounding normal anatomy [2]. Unlike diagnostic scans, this planning CT is optimised for dose calculation, with the CT table being flat to replicate treatment conditions and lasers used to establish reference coordinates. For many patients, additional imaging modalities are co-registered with the planning CT to enhance target definition: MRI fusion provides superior soft-tissue contrast for brain and prostate tumours, while PET-CT adds metabolic information that can identify biologically active subregions requiring higher radiation doses [34]. Four-dimensional CT (4D-CT) may be employed for tumours affected by respiratory motion, capturing images throughout the breathing cycle to quantify tumour movement and inform motion management strategies [35].

2.2.3 Treatment Planning

Following simulation, the acquired images and contours are transferred to a treatment planning system (TPS)—specialised software that serves as the computational engine of modern radiation oncology [36]. The TPSperforms several essential functions: it imports three-dimensional image datasets, enables contouring of target volumes and organs at risk, computes expected dose distributions based on physical beam models, and optimises treatment parameters to achieve clinical objectives.

2.2.4 Contouring and Volume Definition

The first planning step involves delineating the structures to be treated and protected. The radiation oncologist contours the gross tumour volume (GTV), clinical target volume (CTV)—which includes microscopic disease extension—and planning target volume (PTV), which adds a margin to account for setup uncertainties and organ motion [37]. Simultaneously, organs at risk (OARs) such as the spinal cord, lungs, heart, and parotid glands are outlined, each with specific dose constraints that must be respected to minimise toxicity [38]. This contouring process, historically performed manually, is increasingly augmented by auto-segmentation algorithms that use deep learning to identify anatomical structures, reducing both the time required and inter-physician variability [39].

2.2.5 Dose Prescription and Optimisation

With contours complete, the medical dosimetrist and physicist, working under the radiation oncologist's guidance, design the beam arrangement and calculate the treatment plan. Modern radiation therapy employs inverse planning, particularly for intensity-modulated radiation therapy (IMRT) and volumetric modulated arc therapy (VMAT) [40]. In inverse planning, the physician specifies desired dose constraints—for example, "deliver 70 Gy to the PTVwhile keeping spinal cord dose below 45 Gy"—and the optimisation algorithm computes the beam intensities and multileaf collimator (MLC) sequences that best achieve these objectives [41]. For stereotactic treatments delivering extremely high doses per fraction, planning incorporates additional complexity, including 4D datasets that account for tumour motion over time [42]. The TPScalculates dose distributions using sophisticated algorithms that model radiation transport through tissues, accounting for tissue heterogeneity based on CT Hounsfield units [3]. The resulting dose distribution is evaluated using dose-volume histograms, which quantify the percentage of each structure receiving specific dose levels, and by visual inspection of isodose curves overlaid on the planning CT. For reirradiation patients, this planning phase requires additional complexity: previous treatment doses must be retrieved, converted to equieffective doses using appropriate radiobiological models, and accumulated—often through deformable image registration—to estimate the new cumulative dose to organs at risk [43].

2.2.6 Quality Assurance

Before the plan can be delivered to the patient, it undergoes rigorous quality assurance (QA) testing. The physicist verifies that the calculated plan can be accurately delivered by the specific treatment machine for which it was designed [44]. This typically involves delivering the plan to a phantom—a plastic device containing dosimeters—and comparing the measured dose distribution with the calculated one [45]. Discrepancies trigger investigation and potential plan modification. This QAstep ensures that the complex, computer-optimised plan translates into safe and accurate treatment delivery.

2.2.7 Treatment Delivery and Image Guidance

With QAcomplete, the patient begins their course of treatment, which typically spans one to eight weeks with daily fractions delivered Monday through Friday [1]. Each daily session follows a consistent routine: the therapist positions the patient on the treatment couch using the immobilisation devices and skin marks or tattoos created during simulation. Lasers in the treatment room align with these marks to establish the initial position. Modern treatment delivery relies heavily on image guidance (IGRT) to ensure millimetre accuracy. Before delivering the radiation beam, the therapist acquires images of the patient in treatment position—typically cone-beam CT (CBCT) for volumetric verification or planar X-rays for orthogonal alignment [46]. These images are compared with the planning CT, and the treatment couch is automatically adjusted to correct any misalignments [47]. This daily imaging accounts for both setup errors and internal anatomical changes, such as bladder filling or tumour shrinkage during the treatment course. For tumours affected by respiratory motion, real-time motion management may be employed, with the beam gated to fire only during specific phases of the breathing cycle or the treatment couch tracking tumour movement [35]. The actual radiation delivery is painless and lasts only a few minutes, though the entire session—including positioning, imaging, and delivery—typically occupies approximately 30–60 minutes. During beam-on time, the linear accelerator gantry may rotate around the patient, delivering radiation from multiple angles to concentrate dose at the tumour while spreading entrance dose across healthy tissues. Throughout delivery, therapists monitor the patient via cameras and intercom from an adjacent control room.

2.2.8 On-Treatment Monitoring and Adaptation

Throughout the treatment course, patients are monitored for both treatment response and side effects. The radiation oncologist conducts weekly on-treatment visits to assess the patient's clinical status, manage any emerging toxicity, and answer questions. From a technical perspective, the weekly imaging acquired for positioning is also reviewed to detect anatomical changes that might necessitate plan adaptation [48]. If significant changes occur—such as tumour shrinkage, weight loss altering external contour, or pleural effusion developing—the original plan may no longer be optimal. In such cases, adaptive replanning may be

triggered, ranging from simple recalculations on new anatomy to complete re-simulation and re-planning [49]. Emerging online adaptive workflows, enabled by MR-linacs or CT-linacs with artificial intelligence assistance, can even perform real-time replanning with the patient on the treatment couch, creating a new plan adapted to the day's anatomy within a single session [50].

2.2.9 Post-Treatment Follow-Up

Following treatment completion, patients enter a structured follow-up program. An end-of-treatment visit reviews the completed course, discusses expected side effects and their management, and outlines the follow-up schedule. Subsequent follow-up appointments, typically at 4–12 weeks post-treatment and then at regular intervals, assess treatment response, manage late-emerging side effects, and coordinate surveillance imaging [1]. This longitudinal follow-up generates critical data on treatment outcomes and toxicity that, when systematically collected, can inform future treatment decisions and refine planning objectives.

2.2.10 Computational Challenges Across the Workflow

Throughout this workflow, numerous computational challenges emerge that represent potential opportunities for advanced computational methods. Image registration—aligning diagnostic MRI or PET with planning CT, or daily CBCTwith the reference plan—requires robust algorithms capable of handling non-rigid deformations [4]. Dose calculation demands accurate physical modelling with ever-increasing speed, particularly for online adaptive workflows. Optimisation algorithms must navigate complex, multi-objective trade-offs between tumour coverage and organ sparing, often with dozens of conflicting constraints [5]. Auto-segmentation must achieve clinically acceptable accuracy across diverse anatomies and imaging conditions. Cumulative dose assessment for reirradiation requires deformable dose accumulation with rigorous quality assurance. These computational tasks—image analysis, optimisation, physical simulation, and pattern recognition—each present characteristics where quantum machine learning might offer advantages over classical approaches, whether through more efficient optimisation, enhanced feature extraction, or accelerated simulation. Understanding the clinical workflow thus provides the essential context for evaluating where quantum computing could address real clinical needs rather than merely offering theoretical improvements.

2.3 Algorithmic Frontier in Quantum Information and Intelligence

The computational challenges inherent in modern radiation oncology—multi-objective optimization under uncertainty, high-dimensional feature extraction, deformable image registration, and stochastic simulation—demand algorithmic innovations that transcend classical limitations. Quantum computing offers a

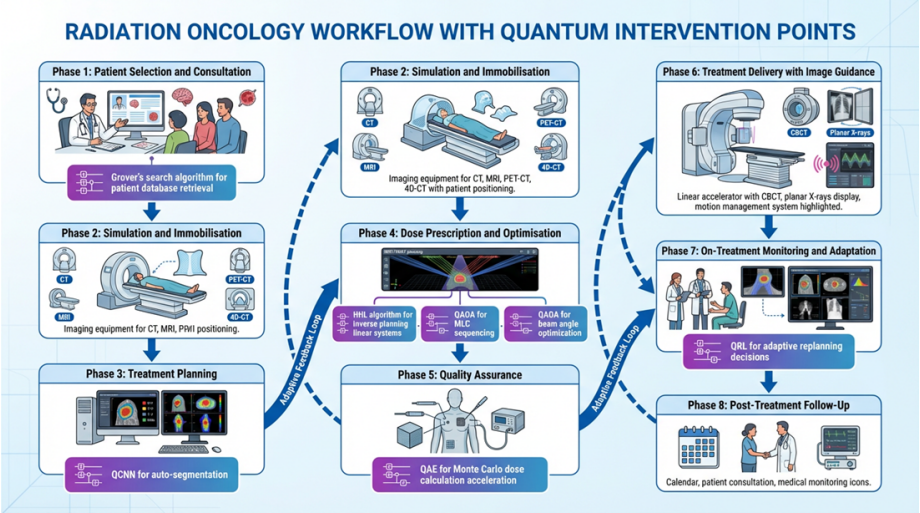

Figure 1: The radiation oncology workflow with quantum intervention points. The flowchart illustrates the sequential phases of clinical radiation therapy: patient selection and consultation, simulation and immobilisation (with multimodal imaging), treatment planning (contouring of GTV, CTV, PTV, and OARs), dose prescription and optimisation (inverse planning for IMRT/VMAT), quality assurance (phantom delivery), treatment delivery with image guidance (CBCT, motion management), on-treatment monitoring and adaptation (online/offline replanning), and post-treatment follow-up. Quantum algorithms that can address computational bottlenecks at each stage are highlighted: Grover's search for patient database retrieval, QAE for Monte Carlo dose calculation acceleration, HHL for inverse planning linear systems, QAOA for MLC sequencing and beam angle optimization, QCNN for auto-segmentation, and QRL for adaptive replanning decisions. Feedback loops between stages, particularly for adaptive radiotherapy, are indicated by arrows.

fundamentally different computational paradigm, leveraging superposition, entanglement, and interference to explore solution spaces inaccessible to classical architectures. This section provides a rigorous taxonomy and critical examination of quantum algorithms relevant to radiation oncology, distinguishing between fault-tolerant algorithms requiring error-corrected logical qubits and near-term variational algorithms suitable for noisy intermediate-scale quantum (NISQ) devices. We emphasize the theoretical foundations of each approach, their computational complexity advantages, and their potential translational pathways to clinical deployment.

2.3.1 Foundational Concepts in Quantum Computation

Before examining specific algorithms, we establish the relevant quantum mechanical principles that confer computational advantage. A quantum state resides in a Hilbert space spanned by computational basis states $\{|0\rangle, |1\rangle\}^{\otimes n}$, enabling superposition: $|\psi\rangle = \sum_{i=0}^{2^n-1} \alpha_i |i\rangle$ with $\sum |\alpha_i|^2 = 1$. Entanglement creates non-classical correlations between qubits that cannot be factorized into product states, while interference allows quantum amplitudes to combine constructively or destructively, amplifying desired computational paths. The challenge of extracting classical information from quantum states—the measurement postulate—imposes fundamental limits on quantum speedups and necessitates repeated measurements or sophisticated estimation techniques. The current quantum computing landscape encompasses two regimes: (i) *fault-tolerant quantum computing* (FTQC), which requires error correction codes with substantial qubit overhead but enables deep circuits with arbitrarily high precision; and (ii) *noisy intermediate-scale quantum* (NISQ) computing, characterized by limited qubit counts (∼50–1000), finite coherence times, and absence of error correction. Most near-term medical applications will likely employ hybrid quantum-classical variational algorithms within the NISQ paradigm, while long-term transformative advances await fault-tolerant implementations.

2.3.2 Fault-Tolerant Quantum Algorithms

Fault-tolerant algorithms achieve exponential or polynomial speedups over classical methods but require logical qubits with error rates below the fault-tolerance threshold. While their clinical implementation remains distant, they establish the theoretical upper bounds of quantum advantage and motivate continued hardware development. To begin with, we have the Shor's algorithm, which solves the hidden subgroup problem for abelian groups, achieving exponential speedup over the best classical factoring algorithms through quantum phase estimation on the modular exponentiation function $f(x) = a^x \mod N$. Despite its revolutionary implications for cryptography, direct applications in radiation oncology are limited. However, Shor's algorithm serves as a existence proof that quantum computers can solve certain problems exponentially faster than classical machines, justifying investment in quantum infrastructure for medical applications. Grover's algorithm performs unstructured search over a space

of size N with $O(\sqrt{N})$ queries, providing a quadratic speedup over classical $O(N)$ scaling. The algorithm iteratively applies the Grover diffusion operator $G = (2\,|\psi\rangle\,\langle\psi| - I)O$, where O is an oracle marking target states. In radiation oncology contexts, this quadratic speedup can accelerate: (i) nearest-neighbor search in large patient databases for case-based reasoning; (ii) combinatorial optimization for beam angle selection in treatment planning; and (iii) feature subset selection in high-dimensional radiomic spaces, where exhaustive search is intractable. The optimal number of iterations $k \approx \frac{\pi}{4}\sqrt{N}$ ensures maximum probability of measuring a marked state.

Quantum Phase Estimation (QPE) forms the computational kernel of numerous quantum algorithms, estimating eigenvalues $e^{2\pi i\phi}$ of unitary operators U with precision 2^{-m} using the inverse quantum Fourier transform. The HHL algorithm leverages QPE to solve systems of linear equations $A\vec{x} = \vec{b}$ with complexity $O(\log(N)\kappa^2/\epsilon)$ for sparse matrices, where κ is the condition number and ϵ the precision—an exponential improvement over classical $O(N\log(N))$ solvers under specific conditions. For radiation oncology, HHL could revolutionize: (i) inverse treatment planning, where optimization reduces to solving large linear systems; (ii) deformable image registration, which requires solving partial differential equations; and (iii) dose calculation kernels that involve transport equations. However, practical implementation requires efficient block-encoding of A, well-conditioned matrices, and amplitude amplification for, with caveats regarding sparsity and condition number that limit applicability to certain problem instances. Quantum Amplitude Estimation (QAE) generalizes Grover's algorithm to estimate the amplitude a of a target state $|\Psi_1\rangle$ in a superposition $|\Psi\rangle = \sqrt{1-a}\,|\Psi_0\rangle + \sqrt{a}\,|\Psi_1\rangle$. Using m ancilla qubits and quantum phase estimation, QAE achieves quadratic speedup over classical Monte Carlo, requiring $O(1/\epsilon)$ samples versus classical $O(1/\epsilon^2)$. In radiation oncology, QAE can accelerate: (i) Monte Carlo dose calculation, where statistical uncertainty scales inversely with the square root of particle histories; (ii) radiobiological endpoint estimation (tumour control probability, normal tissue complication probability); (iii) uncertainty quantification in adaptive radiotherapy; and (iv) probabilistic risk assessment for reirradiation scenarios.

The Quantum Fourier Transform (QFT) implements the discrete Fourier transform on quantum amplitudes with $O(n^2)$ gates, achieving exponential speedup over the classical fast Fourier transform's $O(n2^n)$. While direct QFT on classical data requires amplitude encoding—a nontrivial preprocessing step—its applications include: (i) image filtering in frequency space for CBCT or MRI denoising; (ii) feature extraction in radiomics; and (iii) convolution operations in quantum machine learning architectures. The QFT also underpins QPE and Shor's algorithm, making it a fundamental building block. Quantum walks generalize classical random walks by evolving a walker's amplitude coherently through position and coin spaces. The hitting time for marked vertices can exhibit quadratic or exponential speedups depending on graph structure. In radiation oncology, quantum walks could: (i) model particle transport in Monte Carlo

dose calculations, potentially capturing quantum interference effects in coherent scattering; (ii) optimize brachytherapy seed placement through graph-based search; and (iii) analyze connectivity in radiomic feature networks for outcome prediction.

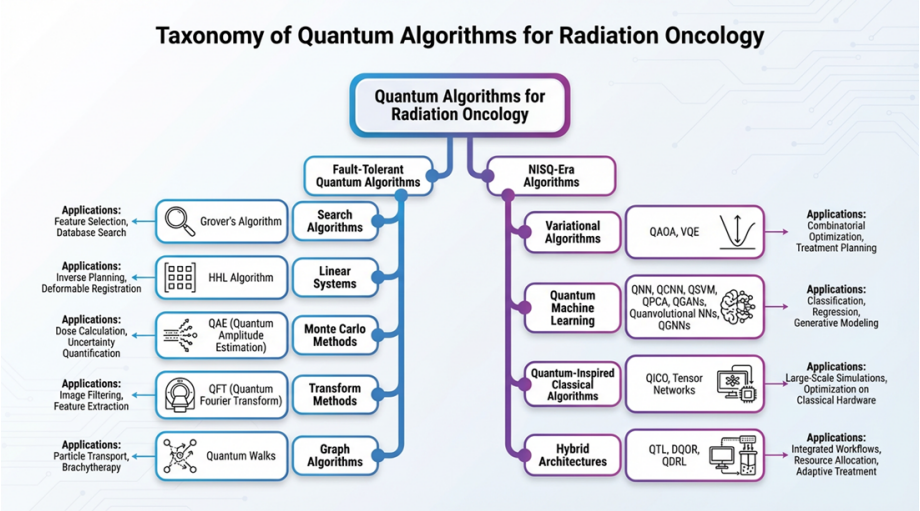

Figure 2: Taxonomy of quantum algorithms relevant to radiation oncology. The tree diagram distinguishes between fault-tolerant algorithms (requiring error-corrected logical qubits) and NISQ-era algorithms (suitable for near-term hardware). Fault-tolerant branch includes search (Grover), linear systems (HHL), Monte Carlo (QAE), transform (QFT), and graph (quantum walks) methods. NISQ-era branch divides into variational algorithms (QAOA, VQE), quantum machine learning (QNN, QCNN, QSVM, QPCA, QGANs, quanvolutional NNs, QGNNs), quantum-inspired classical algorithms (QICO, tensor networks), and hybrid architectures (QTL, DQOR, QDRL). Icons indicate clinical applications: beam angle optimisation, genomics, imaging, etc.

2.3.3 Hybrid Variational Algorithms and Quantum-Enhanced Machine Learning Architectures for the NISQ Era

Given current hardware constraints, variational quantum algorithms (VQAs) represent the most viable pathway for near-term quantum advantage. These algorithms optimize parameterized quantum circuits $U(\vec{\theta})$ using classical optimizers, minimizing cost functions $C(\vec{\theta})$ derived from measurement outcomes. To begin with, Quantum Approximate Optimization Algorithm (QAOA) approximates solutions to combinatorial optimization problems by alternating between problem Hamiltonian H_C and mixer Hamiltonian H_B evolution:

$$U(\beta,\gamma) = e^{-i\beta_p H_B} e^{-i\gamma_p H_C} \cdots e^{-i\beta_1 H_B} e^{-i\gamma_1 H_C} H^{\otimes n} \tag{1}$$

where p is the circuit depth, and (β_i, γ_i) are variational parameters. QAOA has been applied to: (i) multi-leaf collimator (MLC) sequencing in IMRT/VMAT, a constrained combinatorial problem; (ii) beam angle optimization, where the search space grows factorially with candidate angles; (iii) treatment scheduling under resource constraints; and (iv) inverse planning with multiple conflicting objectives. The approximation ratio and depth requirements for clinically relevant problem sizes remain active research areas. Quantum-Inspired Immune Clone Optimization (QICO) synthesizes quantum computing principles (superposition, measurement collapse) with artificial immune systems (clonal selection, affinity maturation). Each antibody is represented as a quantum superposition state, and clonal expansion is modulated by measurement probabilities. In IMRT inverse planning, QICO has demonstrated improved convergence to Pareto-optimal solutions compared to classical genetic algorithms and particle swarm optimization, particularly for non-convex objective landscapes. Quantum Measurement Regression (QMR) frames regression as parameter estimation from quantum measurements, where uncertainties are naturally encoded in measurement statistics. Deep Quantum Ordinal Regressor (DQOR) extends this by integrating convolutional neural networks (CNNs) with QMR layers to perform ordinal regression—predicting ordered categorical outcomes such as toxicity grades (0–5) or tumour response categories. The quantum layer processes CNN-extracted features through a variational circuit, with measurements yielding probabilities over ordinal classes. This architecture offers: (i) inherent uncertainty quantification; (ii) ability to capture correlations between adjacent classes; and (iii) potential resistance to overfitting on small medical datasets.

Quantum machine learning (QML) integrates parameterized quantum circuits into classical learning architectures, leveraging exponentially large Hilbert spaces as feature spaces. The expressibility and entangling capacity of quantum circuits can capture correlations inaccessible to classical neural networks, though trainability issues (barren plateaus) and finite sampling noise pose significant challenges. Quantum Neural Networks (QNNs) consist of layered variational circuits $U(\vec{\theta})$ acting on input-encoded quantum states $\rho(x)$, with measurement outcomes $f(x;\vec{\theta}) = \mathrm{Tr}[OU(\vec{\theta})\rho(x)U^\dagger(\vec{\theta})]$ serving as predictions. For classification tasks, the output is compared to true labels via a cost function, and gradients $\partial_\theta f$ are estimated using the parameter-shift rule. QNNs have been demonstrated for medical image classification on downsampled datasets, but scalability to clinically relevant resolutions (512×512 pixels) requires advanced encoding schemes such as amplitude encoding or quantum random access memory (QRAM). Quantum Convolutional Neural Networks (QCNNs) mimic classical CNNs by applying translationally invariant quantum circuits that preserve spatial structure. Each layer performs convolution-like operations through entangling gates acting on local patches, followed by pooling via partial trace or mid-circuit measurements. QCNNs offer: (i) exponential reduction in parameters compared to classical CNNs; (ii) ability to process quantum data directly (e.g., from quantum sensors); and (iii) potential for enhanced gener-

alization. Improved QCNN (IQCNN) architectures incorporate better entanglement schemes (e.g., hierarchical entanglement), adaptive pooling strategies, and error mitigation techniques to address NISQ limitations. Applications in radiation oncology include: (i) tumour segmentation on multi-modal imaging; (ii) detection of metastatic lesions; and (iii) radiomic feature extraction for outcome prediction.

Quantum Support Vector Machines (QSVMs) map classical data x to quantum states $|\phi(x)\rangle$ via a feature map $U_\phi(x)$, then compute kernel entries $K(x_i, x_j) = |\langle\phi(x_i)\rangle \phi(x_j)|^2$ through swap tests or fidelity estimation. Quantum kernels can exploit classically intractable feature spaces, potentially achieving better separation for complex datasets. In radiation oncology, QSVMs have been applied to: (i) prostate cancer recurrence prediction from multi-parametric MRI; (ii) lung nodule malignancy classification on CT; (iii) toxicity prediction from dosimetric and patient data. The main challenge lies in designing feature maps that are both expressive and efficiently implementable on NISQ hardware. Quantum Principal Component Analysis (QPCA) uses quantum phase estimation to perform PCA on density matrices ρ, exponentiating ρ via the trick $e^{-i\rho t} \approx \prod_k e^{-i\rho_k \Delta t}$ for ensembles $\rho = \sum_k \rho_k$. The algorithm extracts eigenvectors corresponding to largest eigenvalues with exponential speedup when ρ is low-rank. QPCA could accelerate dimensionality reduction in: (i) radiomics, where thousands of features are extracted per patient; (ii) genomics, for identifying biomarkers of radiation response; (iii) dosimetry, for discovering latent patterns in dose distributions predictive of toxicity. Quantum Reinforcement Learning (QRL) incorporates quantum circuits into reinforcement learning agents, representing policies or value functions as parameterized quantum states. The quantum advantage arises from: (i) efficient exploration through superposition of actions; (ii) compressed representation of large state spaces; (iii) quadratic speedups in policy evaluation via amplitude amplification. Quantum Deep Reinforcement Learning (QDRL) extends this to deep architectures, where classical neural networks process raw observations and quantum layers implement the RL core. Potential applications include: (i) adaptive radiotherapy policy optimization, where the agent learns when to replan based on anatomical changes; (ii) personalized fractionation scheduling; (iii) real-time motion management in stereotactic treatments.

Quantum Transfer Learning (QTL) combines pre-trained classical networks (e.g., ResNet, DenseNet) with variational quantum circuits for task-specific fine-tuning. The classical layers extract general features from medical images, while the quantum layer captures dataset-specific correlations that may be subtle but clinically significant. QTL is particularly attractive when medical datasets are small, as the quantum component operates on a reduced-dimensional representation, mitigating barren plateau issues and improving sample efficiency. Quantum Wavelet Transform with Attention Pyramid CNN (QWT-APCNN) integrates quantum wavelet transforms for multi-scale image decomposition with attention mechanisms in a pyramid CNN. The quantum wavelet layer performs

simultaneous spatial-frequency analysis, decomposing images into approximation and detail coefficients through controlled rotations and QFT. The attention pyramid CNN then fuses multi-scale features, emphasizing clinically relevant regions (e.g., tumour boundaries, infiltrative edges). QWT-APCNN could enhance: (i) delineation of gross tumour volumes (GTV) in head and neck cancer; (ii) detection of microscopic disease extension; (iii) fusion of PET/CT/MRI for target definition. While originally developed for quantum chemistry, Variational Quantum Eigensolver (VQE) has emerged as a versatile hybrid algorithm that can be adapted for optimisation problems in radiation oncology. VQE prepares a parameterised quantum state $|\psi(\vec{\theta})\rangle$ and measures the expectation value of a problem Hamiltonian H, minimising $\langle\psi(\vec{\theta})|H|\psi(\vec{\theta})\rangle$ via a classical optimizer. In treatment planning, one could map the dose-delivery constraints to a Hamiltonian and use VQE to search for near-optimal beam configurations. Its scalability and robustness to noise make it a prime candidate for NISQ-enabled radiotherapy optimisation.

Quantum Generative Adversarial Networks (QGANs) extend the classical GAN framework by employing quantum circuits for either the generator, the discriminator, or both. The quantum generator can produce synthetic medical images or radiomic profiles that faithfully represent the training data distribution, even with limited samples. This is particularly valuable for augmenting small oncology datasets, simulating rare toxicity patterns, or generating realistic anatomical changes for adaptive radiotherapy training. Early proof-of-concept studies have demonstrated QGANs for medical image synthesis, although clinical validation remains pending. Quanvolutional neural networks replace the initial convolutional layers of a classical CNN with random or trained quantum circuits that act on local image patches. Each patch is encoded into a quantum state, processed by a parameterised circuit, and measured to produce a feature map. This approach has been shown to achieve competitive accuracy on medical image classification tasks (e.g., breast cancer detection) while using far fewer trainable parameters than classical CNNs. The inherent randomness of the quantum convolution can also act as a regulariser, reducing overfitting on small datasets. Many oncology data types are naturally graph-structured: gene regulatory networks, tumour-atlas connectivity, or patient similarity networks. Quantum Graph Neural Networks (QGNNs) generalise graph neural networks by representing node features as quantum states and employing parameterised quantum circuits for message passing. The ability to encode and process high-dimensional node attributes in superposition may reveal higher-order interactions that are invisible to classical graph models. Potential applications include predicting radiation response from multi-omics networks and modelling tumour evolution during therapy. Although not strictly quantum algorithms, tensor networks (particularly matrix product states) are inspired by the efficient representation of quantum many-body states. They have been successfully applied to compress high-dimensional medical data, perform supervised learning, and solve optimisation problems. In radiomics, where thousands of correlated

features are extracted per patient, tensor networks can uncover low-rank structures and reduce dimensionality while preserving predictive power. Their classical implementation on conventional hardware makes them an immediately accessible bridge between quantum concepts and clinical applications.

2.3.4 Theoretical Considerations and Limitations

Despite the promise of quantum algorithms, several fundamental challenges temper expectations for near-term clinical impact. The *barren plateau* phenomenon causes gradients in variational quantum circuits to vanish exponentially with qubit count, rendering training intractable for large circuits unless carefully initialized or constrained. *Finite sampling noise* limits the precision of gradient estimates, requiring many circuit repetitions that erode quantum advantage. *Quantum decoherence* and gate errors restrict circuit depth on NISQ devices, constraining the complexity of implementable algorithms. *Input/output bottlenecks*—the need to encode classical data into quantum states and extract classical results—can dominate runtime, particularly for dense data such as medical images. Finally, *provable speedups* often require stringent conditions (sparsity, well-conditioned matrices, structured data) that may not hold for real-world oncology problems. These limitations motivate a pragmatic research agenda: (i) developing problem-specific quantum algorithms that exploit the structure of radiotherapy optimization; (ii) designing hardware-efficient ansätze with built-in symmetries; (iii) integrating quantum and classical components to leverage the strengths of both paradigms; and (iv) benchmarking quantum methods against state-of-the-art classical algorithms on clinically relevant problem instances. Only through rigorous, domain-aware evaluation can the true potential of quantum computing in radiation oncology be assessed.

Table 2: Summary of Quantum Algorithms and Their Potential Applications in Radiation Oncology

Algorithm Class	*Specific Algorithm*	*Complexity Advantage*
Search	Grover's Algorithm	$O(\sqrt{N})$ vs. $O(N)$
Linear systems	HHL Algorithm	Exponential for sparse, well-conditioned
Monte Carlo	QAE	Quadratic in precision
Optimization	QAOA	Heuristic for combinatorial problems
Optimization	QICO	Heuristic for non-convex landscapes
Machine learning	QSVM, QCNN, QNN	Expressive quantum feature spaces
Dimensionality reduction	QPCA	Exponential for low-rank matrices
Reinforcement learning	QRL/QDRL	Quadratic in exploration, compressed representation
Signal processing	QFT, QWT	Exponential for Fourier transforms

The algorithmic frontier described herein represents a rapidly evolving landscape. Many approaches remain theoretical or have been demonstrated only on simplified, downsampled problems. However, the convergence of quantum

hardware development, algorithmic innovation, and clinical need creates unprecedented opportunities for transformative advances in radiation oncology. The following sections critically review the existing literature on quantum applications in medical physics and outline a roadmap for future research.

3 Quantum Machine Learning Applications in Radiation Oncology

Having established the foundational principles of quantum computing and the algorithmic landscape relevant to radiation oncology in Section 2, we now turn to the critical question: where and how can quantum machine learning (QML) address real clinical needs? This section provides a comprehensive review of existing and proposed applications of QML in radiation oncology, organized by clinical task and computational paradigm. We examine four major areas: (i) quantum-enhanced diagnostic and predictive modelling, which leverages quantum feature spaces and variational circuits to improve tumour characterization and outcome prediction; (ii) quantum-inspired clinical decision support systems (CDSS) that integrate multi-omics data for precision radiotherapy; (iii) quantum reinforcement learning (QRL) for adaptive treatment planning and dose fractionation; and (iv) the emerging concept of quantum digital twins for system-level oncology integration. For each application area, we describe the quantum algorithms employed, the clinical problem addressed, the current state of evidence from proof-of-concept studies, and the challenges that remain before clinical deployment. Throughout, we emphasize not only the potential advantages of quantum approaches but also the critical evaluation of their performance against classical baselines and the practical constraints of near-term quantum hardware.

3.1 Quantum Enhanced Diagnostic and Predictive Modeling

Quantum-Enhanced ML (QuEML) models show slightly higher accuracy (96% vs 95%) and faster training ($\tilde{1}92.5$ µs gain) [15]. The quantum-enhanced diagnostic workflow extends the classical radiomics pipeline by explicitly incorporating feature harmonization, quantum feature encoding, quantum model training, and quantum-aware optimization, followed by rigorous evaluation and interpretability analysis. The quantum-enhanced diagnostic workflow extends the classical radiomics pipeline by explicitly incorporating feature selection [51]

3.1.1 Data Acquisition and Pre-processing

The initial stages of the radiation oncology workflow—image acquisition, reconstruction, and pre-processing—present several opportunities for quantum and quantum-inspired techniques to enhance data quality and accelerate computationally intensive tasks. These methods operate before any disease-specific

modelling begins, yet their impact propagates through the entire pipeline, influencing the reliability of subsequent feature extraction, segmentation, and predictive modelling. A fundamental task in medical imaging is the reconstruction of clinically usable images from raw sensor data. Whether in CT (sinogram inversion), PET (coincidence detection), or MRI (k-space filling), reconstruction often involves statistical inference problems that require sampling from high-dimensional posterior distributions. Classical Markov chain Monte Carlo (MCMC) methods converge as $O(1/\sqrt{M})$ with the number of samples M, a rate that is often prohibitively slow for real-time applications. Quantum amplitude estimation (QAE) provides a quadratic speedup, achieving root-mean-square error ϵ with $O(1/\epsilon)$ oracle calls rather than the classical $O(1/\epsilon^2)$ [52]. Formally, given a unitary operator $\mathcal{A}$ that acts on a quantum register as

$$\mathcal{A}|0\rangle = \sqrt{a}|\Psi_1\rangle + \sqrt{1-a}|\Psi_0\rangle,$$

QAE estimates the amplitude a to precision ϵ using $O(1/\epsilon)$ applications of $\mathcal{A}$ and the inverse quantum Fourier transform, leveraging quantum phase estimation. In the context of image reconstruction, if one can construct a quantum oracle that encodes the likelihood of a candidate reconstruction, QAE can accelerate the estimation of posterior moments or maximum-a-posteriori (MAP) solutions. Similarly, in histogram-based dose verification, QAE can rapidly estimate dose-volume metrics from simulated particle histories, reducing the computational burden of Monte Carlo dose calculation.

Quantum phase estimation (QPE) serves as a subroutine for estimating eigenvalues of unitary operators, with applications in model parameter estimation from imaging data. For example, in quantitative MRI, parameters such as relaxation times T_1 and T_2 are extracted by fitting signal models to acquired data. This fitting can be formulated as a nonlinear least-squares problem, which is often solved iteratively. QPE offers the potential to directly estimate frequencies (and hence relaxation rates) from the time-domain signal by encoding the signal into the phases of a quantum state [52]. While practical implementations on NISQ hardware remain challenging due to circuit depth requirements, the asymptotic advantage motivates continued investigation. Beyond reconstruction, hybrid quantum-classical algorithms are being explored for image pre-processing tasks that enhance the quality of input data for downstream radiomic analysis. These methods typically use small-scale quantum circuits to perform operations that are computationally expensive classically, while leaving the bulk of the data processing to classical hardware [53]. Quantum denoising algorithms leverage the fact that noise often manifests as high-frequency components in the image's Fourier representation. By encoding the image into a quantum state and applying the quantum Fourier transform (QFT), one can filter out these components through amplitude amplification techniques that preferentially preserve low-frequency coefficients. The QFT operates with $O(n^2)$ gates on an n-qubit register, offering an exponential speedup over the classical fast Fourier transform's $O(n2^n)$ for amplitude-encoded data. However, the need to load classical image data into quantum states (e.g., via amplitude encoding) currently limits

this approach to small image patches or downsampled representations. Quantum clustering algorithms, such as those based on quantum k-means or quantum support vector machines, have been proposed for medical image segmentation. These methods map pixel features (intensity, texture, spatial coordinates) into quantum states and compute distances or kernel values in the exponentially large Hilbert space. For instance, a quantum k-means algorithm can find cluster assignments with quadratic speedup in the number of data points using Grover search [53]. While proof-of-concept studies on small 8×8 image patches have shown feasibility, scaling to clinically relevant image sizes (512×512 or larger) remains an open challenge due to qubit limitations and input/output bottlenecks.

Quantum anomaly detection algorithms can identify regions of interest that deviate from normal tissue appearance. By training a quantum circuit to represent the distribution of healthy tissue features, one can use the circuit's output probability as a measure of abnormality. This approach is conceptually similar to quantum kernel density estimation and may offer advantages in capturing complex, non-linear feature correlations that characterize tumour boundaries. The ultimate goal of these pre-processing techniques is to improve the reliability of downstream radiomic feature extraction and predictive modelling [53]. Radiomics relies on the assumption that quantitative image features (e.g., shape, texture, intensity histograms) are reproducible and reflect underlying tumour biology. Noise, artefacts, and segmentation variability can corrupt these features, leading to non-reproducible biomarkers. By enhancing image quality and segmentation accuracy, quantum-assisted pre-processing can reduce this variability, potentially improving the performance of both classical and quantum predictive models. However, any quantum enhancement at this stage must be rigorously validated against classical state-of-the-art methods to ensure that the added complexity is justified by measurable improvements in downstream task performance. The practical implementation of these techniques hinges on efficient strategies for encoding classical image data into quantum states. Three primary encoding methods are relevant:

- *Amplitude encoding:* Represents an N-dimensional vector as amplitudes of an $n = \lceil \log_2 N \rceil$-qubit state. This is exponentially efficient in qubit count but requires coherent arithmetic to prepare the state from classical data, which is often prohibitive on NISQ hardware.

- *Angle encoding:* Encodes each feature as a rotation on a separate qubit, using N qubits for N features. This is simple to implement but scales linearly in qubit count, limiting applicability to small image patches or feature vectors.

- *QROM and QRAM:* Quantum random access memory (QRAM) and quantum read-only memory (QROM) provide theoretical mechanisms for loading classical data in superposition, but their physical realization remains a major engineering challenge.

Given these constraints, current hybrid approaches typically operate on downsampled images or extracted feature vectors, rather than on full-resolution clinical images. The development of more efficient encoding methods and the maturation of quantum hardware will be essential for translating the theoretical advantages described above into clinically applicable tools. In summary, data acquisition and pre-processing represent the earliest entry point for quantum methods in the radiation oncology workflow. While the field is still in its infancy, with most techniques demonstrated only on simplified problems, the potential for quantum speedups in sampling, parameter estimation, and image enhancement justifies continued exploration. The critical next step is to benchmark these methods against classical alternatives on clinically realistic datasets, under conditions that reflect the noise and scale constraints of actual practice.

3.1.2 Image Processing and Feature Extraction

Medical images constitute the primary data source for radiation oncology, informing target delineation, treatment planning, and response assessment. The extraction of clinically meaningful features from these images—whether for segmentation, tumour characterization, or radiomic analysis—is a computationally intensive task that must capture complex, multi-scale spatial patterns. Quantum machine learning offers several paradigms for enhancing this process, leveraging entanglement, superposition, and interference to represent and extract features that may be inaccessible to classical architectures. Quantum convolutional neural networks (QCNNs) extend the classical CNN paradigm by replacing convolutional filters with parameterized quantum circuits that act on local image patches [15]. In a typical QCNN architecture, a classical image is first encoded into a quantum state using either angle encoding (where pixel intensities determine rotation angles on individual qubits) or amplitude encoding (where the entire image is represented as a superposition state). Parameterized quantum gates—including entangling operations such as CNOT and CZ—then act as quantum filters, creating correlations across qubits that capture complex spatial features. The output is obtained through projective measurements, yielding a high-dimensional feature map that can be used for subsequent classification or regression tasks [15]. The representational power of QCNNs derives from two quantum mechanical properties: (i) entanglement enables the simultaneous encoding of correlations between spatially separated pixels that would require exponentially many classical parameters to represent; and (ii) superposition allows each quantum filter to evaluate multiple feature combinations in parallel. Studies on ultrasound breast images have demonstrated that QCNN-based quantum filters can outperform classical CNN filters of comparable size, and hybrid quantum-classical-quantum (QCQ) architectures further validate the advantage conferred by entanglement-based feature extraction [54]. However, these demonstrations have thus far been limited to downsampled images (typically 8×8 or 16×16 patches) due to qubit constraints and the overhead of classical-quantum data conversion.

Improved QCNNs (IQCNNs) build upon the basic QCNN architecture by incorporating trainable quantum filters, deeper variational circuits, and advanced pre-processing strategies to capture more subtle clinical patterns [53]. The "improved" designation encompasses several architectural innovations:

- *Trainable quantum filters:* Rather than using fixed entangling gates, IQCNNs parameterize the gates within each convolutional layer, allowing the quantum circuit to adapt to the specific feature extraction task during training. This increases the expressivity of the model at the cost of additional trainable parameters and deeper circuits.
- *Hierarchical entanglement schemes:* By structuring entanglement to respect the spatial hierarchy of image features (e.g., local edges, textures, and global shape), IQCNNs can extract multi-scale representations more efficiently than generic all-to-all entanglement.
- *Adaptive pooling strategies:* Quantum pooling operations, implemented via partial trace or mid-circuit measurements, are designed to preserve clinically relevant information while reducing dimensionality. In IQCNNs, these pooling layers may themselves be parameterized and optimized during training.
- *Integration with classical pre-processing:* IQCNNs often incorporate classical pre-processing steps—such as edge detection, wavelet decomposition, or contrast enhancement—to reduce the burden on the quantum circuit and to ensure that the input data is optimally formatted for quantum encoding.

The combination of these techniques enables IQCNNs to model prognosis more accurately by learning deeper representations of tumour heterogeneity, subtle infiltrative patterns, and treatment-induced changes. Nevertheless, the exact mechanisms by which quantum circuits capture these clinical nuances remain an active area of investigation, and systematic ablation studies are needed to disentangle the contributions of quantum entanglement from those of increased parameter counts and architectural complexity.

Beyond feature extraction, the selection of informative features from a high-dimensional pool is critical for building parsimonious and generalizable predictive models. Quantum-inspired immune clone optimization (QICO) is a hybrid metaheuristic that combines principles from quantum computing—specifically superposition and measurement collapse—with artificial immune systems (clonal selection, affinity maturation) to perform feature selection [53]. In QICO, each candidate feature subset is represented as a quantum superposition state, and the population of subsets evolves through clonal expansion, where the probability of cloning a given subset is proportional to its measurement probability (i.e., its fitness). Affinity maturation is implemented through quantum-inspired mutation operators that explore the feature space while preserving promising substructures. The algorithm terminates when measurement probabilities concen-

trate on a near-optimal subset. Studies in cancer classification have shown that QICO can outperform discovery-based feature selection methods (e.g., forward selection, genetic algorithms) in both accuracy and computational efficiency, and the approach is conceptually extendable to dose optimization problems where the feature space comprises beamlet weights or other treatment parameters [53]. However, as with other quantum-inspired methods, QICO runs on classical hardware and does not require a physical quantum computer; its advantage stems from algorithmic innovation rather than quantum parallelism per se. Hybrid quantum-classical image processing algorithms have been developed to enhance multi-scale feature extraction and to model ordered clinical imaging outcomes. The quantum wavelet transform (QWT) with attention pyramid CNN (QWT-APCNN) is a prominent example [53]. The QWT decomposes an image into approximation and detail coefficients at multiple scales using controlled rotations and the quantum Fourier transform, achieving simultaneous spatial-frequency analysis that is computationally more efficient than classical wavelet transforms for amplitude-encoded data. The resulting multi-scale representation is then fed into a pyramid CNN equipped with attention mechanisms that emphasize clinically relevant regions (e.g., tumour boundaries, areas of suspected infiltration). This architecture has been applied to tasks such as tumour segmentation and response assessment, where the combination of quantum multi-scale analysis and classical attention-based refinement yields improved robustness to variations in image quality and acquisition protocol.

Quantum genetic algorithms (QGAs) leverage quantum principles to enhance the exploration of large combinatorial spaces in image processing. By representing candidate solutions as quantum states and using quantum gates to implement crossover and mutation, QGAs can theoretically achieve quadratic speedups over classical genetic algorithms [53]. In practice, most implementations are quantum-inspired and run on classical hardware, but they have been successfully applied to optimize imaging representations (e.g., selecting optimal filter banks or registration parameters). Deep quantum ordinal regression (DQOR) models extend the QMR framework introduced in Section 2.3.3 to the task of predicting ordered clinical outcomes from imaging data. By integrating CNNs for feature extraction with quantum measurement layers that naturally model uncertainty, DQOR can predict ordinal endpoints such as toxicity grades or tumour response categories while providing calibrated confidence estimates [53]. This is particularly valuable in adaptive radiotherapy, where decisions must be made under uncertainty and where the cost of misclassification varies with the severity of the outcome. The potential advantages of quantum approaches in image processing and feature extraction can be traced to two fundamental quantum phenomena: interference and parallelism. Quantum interference allows amplitudes corresponding to different feature combinations to combine constructively or destructively, effectively amplifying patterns that are indicative of a particular class or outcome while suppressing irrelevant variations. Quantum parallelism enables a single circuit to simultaneously evaluate multiple feature maps, providing an exponential speedup in the number of fea-

ture combinations that can be explored per circuit evaluation [53]. However, these theoretical advantages are contingent on the ability to encode image data efficiently, to design circuits that avoid barren plateaus, and to extract useful classical information from measurement outcomes without excessive sampling overhead.

In summary, image processing and feature extraction constitute a rich domain for quantum and quantum-inspired methods in radiation oncology. From QCNNs that leverage entanglement as a computational resource, to hybrid architectures like QWT-APCNN that combine quantum multi-scale analysis with classical attention mechanisms, these techniques offer the potential to extract more expressive and clinically meaningful representations from medical images. The critical next steps involve systematic benchmarking against classical state-of-the-art methods on large, clinically representative datasets, and the development of encoding strategies that can scale to full-resolution images without prohibitive qubit or circuit depth requirements.

3.1.3 Feature Selection

In radiomics and multi-omics analysis, the dimensionality of extracted features often exceeds the number of available patient samples by orders of magnitude—a regime that is prone to overfitting, multicollinearity, and the curse of dimensionality. Feature selection aims to identify a subset of informative, non-redundant features that are predictive of clinical outcomes while discarding noise and irrelevant measurements. Quantum and quantum-inspired methods offer new approaches to this problem, leveraging principles such as superposition, entanglement, and eigenvalue estimation to explore high-dimensional feature spaces more efficiently than classical techniques. Quantum principal component analysis (qPCA) is a dimensionality reduction technique that identifies the principal components of a dataset by performing eigenvalue decomposition on its covariance matrix using quantum phase estimation [53]. Given a set of M data vectors $\{\mathbf{x}_j\}_{j=1}^{M}$ in $\mathbb{R}^N$, the classical PCA problem involves computing the eigenvectors of the covariance matrix $C = \frac{1}{M}\sum_{j=1}^{M} \mathbf{x}_j \mathbf{x}_j^\top$. For high-dimensional data ($N \gg M$), this computation scales as $O(N^3)$ classically, which is prohibitive for large radiomic feature sets. The qPCA algorithm operates on the density matrix $\rho = \frac{1}{M}\sum_{j=1}^{M} |\mathbf{x}_j\rangle\langle\mathbf{x}_j|$ representing the quantum states corresponding to each data vector. The key insight is that the exponential of the density matrix, $e^{-i\rho t}$, can be applied to an arbitrary quantum state using the trick

$$e^{-i\rho t} \approx \prod_{k=1}^{K} e^{-i\rho_k \Delta t},$$

where $\rho = \sum_k \rho_k$ is decomposed into a sum of easily implementable terms [53]. With this capability, quantum phase estimation can be used to estimate the eigenvalues and prepare the eigenvectors of ρ, which correspond to the principal components of the original data. Under appropriate conditions, qPCA offers an

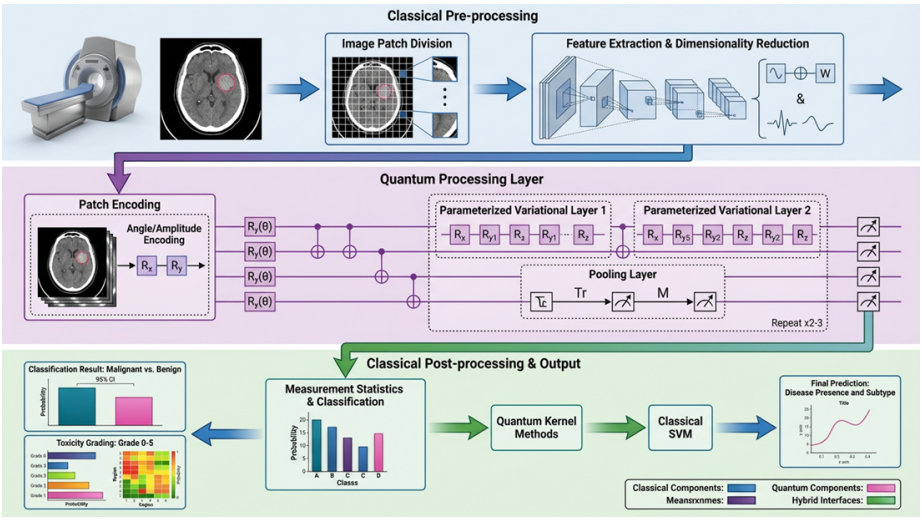

Figure 3: Hybrid quantum-classical architecture for medical image analysis using a quantum convolutional neural network (QCNN). Top (classical pre-processing): a CT or MRI image is divided into patches; feature extraction via classical CNN layers or wavelet transforms reduces dimensionality. Middle (quantum processing): each patch is encoded into a quantum circuit via angle or amplitude encoding. The circuit comprises data encoding layers (e.g., Ry rotations), entangling layers (CNOT), variational layers (trainable quantum filters), and pooling layers (partial trace or mid-circuit measurements). Measurement yields a probability distribution over classes. Bottom (classical post-processing): measurement statistics are converted to classification outputs (e.g., malignant vs. benign, toxicity grade) with confidence intervals. A parallel track shows quantum kernel methods where circuit outputs feed into a classical SVM. Colour coding distinguishes classical (blue), quantum (purple), and hybrid (green) components.

exponential speedup over classical PCA in terms of the dimension N, requiring only $O(\log N)$ qubits and $O(\log N)$ gates per step.

In the context of radiation oncology, qPCA could accelerate dimensionality reduction in several applications:

- *Radiomics:* Thousands of quantitative features (shape, texture, intensity, wavelet) are extracted from medical images. qPCA could identify a low-dimensional subspace that captures the majority of variance in these features, enabling more robust predictive models and revealing latent imaging phenotypes associated with treatment response.

- *Multi-omics integration:* When combining genomics, proteomics, and radiomics data, the total feature count can reach tens of thousands. qPCA offers a pathway to compress these heterogeneous datasets into a unified representation that preserves biologically relevant signal while discarding noise.

- *Dosimetry:* Dose-volume histograms and voxel-level dose distributions can be treated as high-dimensional vectors. qPCA could uncover latent patterns in dose deposition that correlate with toxicity, potentially informing more personalized planning objectives.

Despite its theoretical promise, several practical obstacles currently limit the application of qPCA to clinical problems:

- *State preparation:* The algorithm assumes efficient preparation of the states $|\mathbf{x}_j\rangle$ corresponding to each data vector. For dense, high-dimensional radiomic feature vectors, this requires either amplitude encoding (which demands coherent arithmetic operations that are challenging on NISQ hardware) or access to a quantum random access memory (QRAM) that can load classical data in superposition. Neither is currently available at the scale required for clinical datasets.

- *Low-rank assumption:* The exponential speedup of qPCA is most pronounced when ρ is low-rank, i.e., when most of the variance is captured by a small number of principal components. While this often holds for radiomic features (which are frequently correlated), it is not guaranteed for all datasets, and the algorithm's performance degrades as the rank increases.

- *Measurement overhead:* Extracting the classical principal components from the quantum eigenvectors requires tomographic reconstruction, which scales exponentially with the number of qubits if full state information is needed. In practice, one can only estimate a few features of the eigenvectors (e.g., overlaps with known states), limiting the utility of the output for downstream classical modelling.

- *NISQ-era constraints:* On current hardware, circuit depth and noise limit qPCA to proof-of-concept demonstrations on synthetic, low-dimensional data. Scaling to clinically relevant feature sets will require advances in error correction, coherence times, and qubit counts.

Beyond qPCA, other quantum techniques have been proposed for feature selection. Quantum support vector machines (QSVMs) with explicit feature maps can be used to rank features by their contribution to the decision boundary, effectively performing feature selection as part of model training [15]. Grover's algorithm can accelerate exhaustive search over feature subsets, providing a quadratic speedup over classical branch-and-bound methods when the number of features F is modest (typically $F \lesssim 30$) [52]. However, for the high-dimensional feature spaces typical of radiomics ($F \sim 10^3 - 10^4$), even a quadratic speedup leaves the problem intractable, and heuristic methods remain necessary. Quantum-inspired immune clone optimization (QICO), discussed in Section 3.1.2, has also been applied to feature selection, demonstrating improved performance over classical genetic algorithms in cancer classification tasks [53]. While not requiring quantum hardware, QICO illustrates how quantum principles can inspire classical algorithms that outperform traditional methods.

Feature selection does not occur in isolation; it is intimately connected to the encoding, modelling, and validation steps that precede and follow it. In a quantum-enhanced diagnostic pipeline, feature selection methods such as qPCA or QICO would typically be applied after feature extraction (Section 3.1.2) and before quantum feature encoding (Section 3.1.4). The selected features must be compatible with the encoding scheme used by the downstream quantum model—for instance, if amplitude encoding is used, the number of selected features must be a power of two, and their dynamic range must be suitable for representation as probability amplitudes. This coupling between selection and encoding adds an additional layer of complexity that is absent in classical workflows. In summary, feature selection represents a critical bottleneck in high-dimensional oncology data analysis, and quantum methods offer intriguing theoretical advantages. qPCA, in particular, promises exponential speedups in dimensionality reduction under ideal conditions. However, substantial hardware and algorithmic advances are required before these methods can be applied to clinically realistic datasets. In the near term, quantum-inspired classical algorithms such as QICO may offer a more practical pathway to improved feature selection, providing some of the benefits of quantum principles without the need for fault-tolerant hardware.

3.1.4 Quantum Feature Encoding

The translation of classical clinical data—whether radiomic features, genomic profiles, or imaging-derived biomarkers—into quantum states is a fundamental prerequisite for any quantum machine learning (QML) application. This process, known as quantum feature encoding, determines both the representational

capacity of the quantum model and the feasibility of its implementation on near-term hardware. The choice of encoding strategy directly impacts the number of qubits required, the depth of the quantum circuit, and the resilience to noise, making it a critical design decision in any QML pipeline for radiation oncology. Three primary encoding methods are relevant to clinical data:

- *Basis encoding:* Classical binary data is directly mapped to computational basis states. For a binary string $x = x_1 x_2 \dots x_n$, the quantum state $|x\rangle = |x_1\rangle \otimes |x_2\rangle \otimes \cdots \otimes |x_n\rangle$ is prepared. This method is straightforward and requires only single-qubit gates, but it is inefficient for continuous-valued radiomic features unless they are first discretized, which may lose information.

- *Angle encoding:* Each feature θ_i is encoded as a rotation on a separate qubit: $|\psi(\boldsymbol{\theta})\rangle = \bigotimes_{i=1}^{n} R_Y(\theta_i)|0\rangle$. This requires n qubits for n features and has circuit depth $O(1)$, making it attractive for NISQ implementations. However, the linear scaling in qubits limits its applicability to feature sets of moderate size ($n \lesssim 100$), which is often insufficient for radiomics where thousands of features are extracted.

- *Amplitude encoding:* A normalized 2^n-dimensional classical vector $\mathbf{x} = (x_0, x_1, \dots, x_{2^n-1})$ is encoded as the amplitudes of an n-qubit quantum state: $|\psi_{\mathbf{x}}\rangle = \sum_{i=0}^{2^n-1} x_i |i\rangle$. This is exponentially efficient in qubit count, requiring only $\lceil \log_2 N \rceil$ qubits for an N-dimensional feature vector. However, state preparation typically requires $O(2^n)$ gates and coherent arithmetic operations, which are challenging on current hardware. Moreover, the normalization condition $\sum_i |x_i|^2 = 1$ imposes constraints on the data representation.

Hybrid strategies, such as amplitude encoding of compressed features or angle encoding of selected features, are often employed to balance expressivity with implementability. The choice of encoding must also account for the downstream quantum model: kernel methods require the ability to estimate inner products $\langle \phi(\mathbf{x}) | \phi(\mathbf{x}') \rangle$, which is natural with amplitude encoding, while variational circuits often use angle encoding for simplicity.

A compelling demonstration of quantum feature encoding in clinical practice is provided by Huang et al. [55], who developed an integrated framework for the classification of pure ground-glass nodules (pGGNs) in lung CT. pGGNs present a diagnostic challenge due to their subtle appearance and the need to distinguish indolent from potentially malignant lesions. The workflow comprises several stages that illustrate the principles of quantum feature encoding in a real clinical context. From each CT image, a comprehensive set of radiomic features was extracted, capturing intensity distributions, texture patterns, shape descriptors, and wavelet decompositions. These features constitute a high-dimensional representation of the nodule's phenotype. To reduce dimensionality and mitigate overfitting, mutual information feature selection (MIFS) was employed

[56]. MIFS selects features that have high relevance to the target class (malignant vs. benign) while minimizing redundancy, yielding a compact feature subset that retains predictive power. This classical pre-processing step is essential to make the subsequent quantum encoding tractable, as it reduces the number of features to a size compatible with angle encoding on available qubits. The selected features can be encoded into quantum states using angle encoding, with each feature determining a rotation on a dedicated qubit. This choice reflects the practical constraints of NISQ hardware: angle encoding requires only shallow circuits and is resilient to certain types of noise, making it feasible for implementation on current quantum processors and simulators. Three quantum classifiers can be evaluated on the encoded features:

- *Quantum Support Vector Classifier (QSVC):* This method uses a quantum kernel $k(\mathbf{x}, \mathbf{x}') = |\langle\phi(\mathbf{x})|\phi(\mathbf{x}')\rangle|^2$, where $|\phi(\mathbf{x})\rangle$ is the angle-encoded state. The kernel matrix is estimated on a quantum computer or simulator and then passed to a classical SVM for training. This hybrid approach leverages quantum resources only for the computationally intensive kernel estimation.

- *Pegasos QSVC:* An adaptation of the Pegasos (Primal Estimated sub-GrAdient SOlver for SVM) algorithm that incorporates quantum kernel estimation, enabling more efficient training on large datasets.

- *Quantum Neural Network (QNN):* A variational circuit with trainable parameters, optimized to minimize a classification loss function. The QNN was trained using the parameter-shift rule, with gradients estimated from measurement outcomes.

These quantum models can be compared against a classical Support Vector Machine (SVM) with radial basis function kernel, serving as a baseline.

QML models—particularly the QSVC-based methods—exhibit enhanced accuracy and robustness compared to the classical SVM, especially in the small-sample setting typical of rare nodule types [55]. This suggests that quantum feature maps can capture discriminative patterns that are not easily accessible to classical kernels when data are limited. Crucially, such a framework can incorporate SHapley Additive exPlanations (SHAP) to attribute the model's predictions to individual radiomic features. By analyzing the SHAP values in the context of the quantum feature map, we can see that the quantum models are not only accurate but also interpretable: the features identified as most influential aligned with clinical knowledge of pGGN malignancy risk (e.g., texture heterogeneity, margin irregularity). This integration of explainable AI with QML addresses a major barrier to clinical adoption, as clinicians require transparency in decision support systems. There are several salient points for quantum feature encoding in radiation oncology:

- *Dimensionality reduction is essential:* Direct amplitude encoding of full radiomic feature vectors is currently infeasible. Classical pre-processing

(feature selection, PCA) is necessary to reduce the data to a size compatible with angle encoding on available qubits.

- *Angle encoding is viable for small-to-moderate feature sets:* With current qubit counts (50–100), angle encoding can accommodate up to 100 features, which is sufficient for many radiomic applications after feature selection.
- *Kernel methods benefit from quantum encoding:* The QSVC approach, which uses quantum resources only for kernel estimation, is particularly well-suited to the NISQ era, as it requires only shallow circuits and avoids the trainability issues of variational models.
- *Interpretability must be designed in:* The integration of SHAP analysis with quantum models demonstrates that QML need not be a black box. By attributing predictions to input features, these methods can build clinical trust and provide insights into disease biology.

Despite these promising results, several challenges remain. Angle encoding scales linearly in qubits, limiting its applicability to feature sets that have already been aggressively reduced. Amplitude encoding, while exponentially more efficient in qubit count, requires coherent state preparation that is beyond current NISQ capabilities. Hybrid encoding schemes—for example, using amplitude encoding for groups of related features—may offer a middle ground. Additionally, the sensitivity of quantum feature maps to noise and gate errors means that hardware implementations may not replicate the performance seen in simulators. Robustness studies and error mitigation techniques will be essential as these methods move toward clinical deployment. The integration of quantum feature encoding with classical pre-processing and post-hoc interpretability, as demonstrated in the pGGN classification framework, provides a template for future QML applications in radiation oncology. By carefully matching the encoding strategy to the data characteristics and hardware constraints, and by validating performance against classical baselines, researchers can begin to translate the theoretical advantages of quantum computing into clinically meaningful improvements in diagnostic accuracy and treatment personalization.

3.1.5 Quantum Model Training and Inference

The training and inference phases of quantum machine learning encompass a diverse array of methodologies, each tailored to the specific clinical task at hand—whether classification, regression, optimization, or simulation. The choice of quantum model, optimization strategy, and inference protocol fundamentally determines both the predictive performance and the practical feasibility of deployment in radiation oncology workflows. This subsection provides a comprehensive review of quantum models used for training and inference in oncological applications, organized by model family and clinical context. Quantum models can be broadly categorized by their architectural principles and the nature of their interaction with classical hardware:

- *Variational quantum circuits (VQCs):* Parameterized quantum circuits trained via classical optimization to minimize a cost function. These include Quantum Neural Networks (QNNs), Quantum Convolutional Neural Networks (QCNNs), and variational classifiers.
- *Quantum kernel methods:* Models that estimate classically intractable kernels using quantum circuits, then feed the resulting kernel matrix into a classical support vector machine. Quantum Support Vector Machines (QSVMs) and quantum kernel estimators fall into this category.
- *Quantum generative models:* Models that learn to generate samples from a target distribution, such as Quantum Boltzmann Machines (QBMs) and Quantum Generative Adversarial Networks (QGANs).
- *Quantum reinforcement learning (QRL):* Models that learn optimal policies through interaction with an environment, with policies represented as parameterized quantum circuits.
- *Hybrid quantum-classical models:* Architectures that combine classical neural networks (for feature extraction) with quantum layers (for representation learning), such as Quantum Transfer Learning (QTL) and hybrid QNNs.

Each of these model classes has been explored for applications in oncology, with varying degrees of theoretical justification and empirical validation.

Quantum Neural Networks (QNNs) represent the most direct quantum analogue of classical neural networks. A QNN consists of a parameterized quantum circuit $U(\boldsymbol{\theta})$ that acts on an input-encoded quantum state $\rho_{\text{in}}(\mathbf{x})$, producing an output state from which measurement outcomes $f(\mathbf{x};\boldsymbol{\theta}) = \text{Tr}[OU(\boldsymbol{\theta})\rho_{\text{in}}(\mathbf{x})U^\dagger(\boldsymbol{\theta})]$ are obtained. For classification tasks, the output is compared to the true label via a cost function $C(\boldsymbol{\theta})$, and gradients $\partial C/\partial\theta_\mu$ are estimated using the parameter-shift rule [53]. QNNs have been applied to a range of cancer diagnosis and classification problems, including mammography interpretation, prostate cancer grading, and lung nodule characterization [53]. Their primary advantage lies in the expressivity of the quantum feature space: with n qubits, a QNN can represent functions that would require exponentially many classical parameters. However, this expressivity comes with trainability challenges, as discussed in Section 2.3.5. In practice, QNNs have been demonstrated on downsampled medical images and reduced feature sets, with performance competitive with classical baselines but not yet exceeding them on clinically relevant scales. Quantum Support Vector Machines (QSVMs) offer a different paradigm: rather than training a variational circuit, they use quantum computers to estimate kernels that are classically hard to compute. Given a quantum feature map $\phi : \mathcal{X} \to \mathcal{H}$, the kernel $k(\mathbf{x},\mathbf{x}') = |\langle\phi(\mathbf{x})|\phi(\mathbf{x}')\rangle|^2$ is estimated via a swap test or fidelity estimation on a quantum device. The resulting kernel matrix is then passed to a classical SVM for training [53]. The theoretical advantage of QSVMs is the

potential for exponential speedup in kernel evaluation under certain complexity-theoretic assumptions (Section 2.3.1). Empirically, QSVMs have demonstrated strong performance in cancer diagnostics, including prostate cancer detection from multi-parametric MRI and lung nodule classification from CT [15]. The quantum kernel mapping appears to capture non-linear relationships in the data that are not accessible to classical kernels with limited complexity. However, QSVMs are fundamentally linear classifiers in the quantum feature space; if the data remain linearly inseparable even after the quantum map, performance degrades [53]. Moreover, the requirement to estimate the full kernel matrix scales as $O(M^2)$ in the number of samples M, which can be prohibitive for large datasets.

The Pegasos QSVC variant adapts the primal estimated sub-gradient solver for SVM to incorporate quantum kernel estimation, enabling more efficient training on larger datasets by avoiding the storage of the full kernel matrix [55]. This hybrid approach has shown particular promise in small-sample settings, where the inductive bias of the quantum kernel can compensate for limited data. Quantum Transfer Learning (QTL) addresses one of the fundamental challenges of QML in healthcare: the scarcity of labeled medical data. In QTL, a classical neural network pre-trained on a large natural image corpus (e.g., ImageNet) is used as a fixed feature extractor. The high-dimensional features produced by this network are then passed to a variational quantum circuit for task-specific fine-tuning [53]. This approach offers several advantages:

- *Dimensionality reduction:* The classical network compresses raw images into a compact feature vector (typically 256–1024 dimensions), which can then be encoded into a quantum state using angle encoding on a modest number of qubits.
- *Mitigation of barren plateaus:* By operating on a reduced-dimensional representation, the quantum circuit can be designed with locality-preserving structures that avoid the exponential gradient vanishing that plagues larger circuits.
- *Improved generalization:* The inductive bias of the quantum layer, combined with the rich features from the pre-trained network, often yields better generalization on small medical datasets than either component alone.

QTL has been successfully applied to mammography image classification, demonstrating faster convergence and improved accuracy compared to classical transfer learning baselines [53]. The quantum layer appears to capture subtle, dataset-specific correlations that complement the general features extracted by the classical network. Beyond static classification, hybrid quantum models have been developed for predicting temporal outcomes in radiation oncology. Quantum recurrent neural networks (QRNNs) extend the variational circuit paradigm to sequential data, making them suitable for modeling tumor volume evolution

over the course of treatment and daily setup variations [14]. In these models, the hidden state at each time step is represented as a quantum state, and parameterized circuits implement the recurrence.

At the predictive modeling stage, hybrid QNNs and quantum Bayesian inference models are employed to enhance robustness under data sparsity and uncertainty. By integrating variational quantum circuits with probabilistic reasoning frameworks, these models can:

- *Represent uncertainty explicitly:* Quantum probability distributions naturally capture both aleatoric uncertainty (inherent randomness in outcomes) and epistemic uncertainty (lack of knowledge due to limited data).
- *Propagate uncertainty through time:* In sequential models, quantum states can represent distributions over hidden states, enabling principled handling of noisy observations and incomplete data.
- *Improve small-sample performance:* The inductive bias of the quantum circuit, combined with Bayesian regularization, reduces overfitting when training data are scarce.

Studies have shown that these hybrid quantum-classical models achieve improved uncertainty-aware classification and regression in small-sample oncological datasets compared to classical benchmarks [14]. Classical outcome prediction models (OPMs) estimate tumor control probability (TCP) and normal tissue complication probability (NTCP) using dose-volume histograms combined with machine learning or deep learning frameworks. These models typically treat patient characteristics as fixed inputs and produce point estimates of outcome probabilities, with uncertainty handled through separate confidence intervals or ensemble methods [14]. In contrast, quantum-inspired OPMs represent patient states as quantum probability distributions. In this formalism, the state of a patient is described by a density matrix ρ that encodes not only the most likely values of clinical variables but also the correlations and uncertainties among them. Outcome probabilities are obtained through projective measurements corresponding to clinical endpoints:

$$P(\text{outcome} = i) = \text{Tr}[M_i \rho],$$

where $\{M_i\}$ are positive operator-valued measures (POVMs) representing the clinical events of interest (e.g., tumor control within a specified time, grade ≥ 2 toxicity). This representation naturally captures several sources of uncertainty that are critical in radiation oncology:

- *Noisy imaging:* Uncertainties in target delineation and organ-at-risk segmentation can be encoded as mixed states.
- *Incomplete biomarkers:* Missing genomic or proteomic data correspond to partial traces over unobserved degrees of freedom.

- *Inter-patient variability:* Population-level distributions can be represented as ensembles of pure states, with individual patients as samples from this ensemble.

While these quantum-inspired OPMs are typically implemented on classical hardware (using tensor networks or probabilistic graphical models), they illustrate how quantum concepts can inform the design of more expressive and uncertainty-aware predictive models, even in the absence of a physical quantum computer.

The training of quantum models for oncology applications can follow several distinct paradigms, depending on the model class and the availability of quantum hardware:

- *Classical optimization of variational circuits:* For QNNs, QCNNs, and variational classifiers, training proceeds via iterative updates of the circuit parameters $\boldsymbol{\theta}$ using gradient-based or gradient-free optimizers. Gradients are estimated using the parameter-shift rule, which requires $O(P)$ circuit evaluations per gradient step for P parameters, each evaluation comprising multiple shots to estimate expectation values.
- *Kernel estimation for QSVMs:* Training involves estimating the kernel matrix on a quantum device (or simulator) and then solving the dual SVM problem classically. The quantum computational cost is dominated by the $O(M^2)$ kernel evaluations, each requiring $O(1/\epsilon^2)$ shots for precision ϵ.
- *Hybrid training loops:* In models like QTL, the classical and quantum components are trained jointly: the classical network may be fine-tuned while the quantum circuit parameters are updated, requiring careful coordination to avoid destructive interference between the two optimizers.
- *Quantum natural gradient:* Advanced optimizers incorporate the geometry of the quantum state space via the quantum Fisher information matrix, potentially accelerating convergence but at the cost of additional circuit evaluations.

The choice of optimizer, learning rate, and batch size must be adapted to the noise characteristics of the hardware and the statistical properties of the clinical data. In the NISQ era, simulators often outperform real hardware due to decoherence and gate errors, necessitating error mitigation techniques such as zero-noise extrapolation and measurement error correction.

Once trained, quantum models are used for inference on new patient data. Inference typically involves:

1. *Data encoding:* The new patient's features are encoded into a quantum state using the same encoding scheme employed during training.
2. *Circuit execution:* The trained circuit is applied to the encoded state, either once (for variational models) or multiple times (for kernel methods).

3. *Measurement:* Projective measurements are performed on the output state, yielding a probability distribution over possible outcomes.

4. *Classical post-processing:* The measurement statistics are converted into clinical predictions (e.g., class labels, risk scores, TCP/NTCP estimates) and confidence intervals.

For kernel methods, inference requires evaluating the kernel between the new sample and all training samples, which can be computationally intensive for large datasets but benefits from the same quantum speedups as training. Despite the diversity of quantum models explored for oncology applications, several common challenges emerge:

- *Scalability:* Current demonstrations are limited to small feature sets and downsampled images. Scaling to full-resolution clinical data will require advances in both hardware (more qubits, lower noise) and algorithms (more efficient encoding, hierarchical models).

- *Trainability:* Barren plateaus remain a fundamental obstacle for large variational circuits. Problem-specific ansätze, layer-wise training, and initialization strategies are active areas of research.

- *Verification and validation:* Many published results are based on simulators with idealized noise models. Performance on real hardware often degrades, and systematic comparisons to classical state-of-the-art are sometimes lacking.

- *Interpretability:* While SHAP and similar methods can provide post-hoc explanations, the inner workings of quantum models remain opaque. Developing intrinsically interpretable quantum architectures is an important long-term goal.

Looking forward, the most promising directions for quantum model training and inference in radiation oncology include: (i) hybrid architectures that leverage classical networks for feature extraction and quantum circuits for final classification, (ii) quantum kernel methods for small-to-medium datasets where classical kernels struggle, and (iii) quantum-inspired probabilistic models that capture uncertainty more faithfully than classical alternatives. As hardware improves and algorithms mature, these approaches may gradually transition from proof-of-concept to clinically deployable tools.

3.1.6 Quantum-Aware Optimization and Treatment Planning

Treatment planning in radiation oncology is fundamentally a constrained multi-objective optimization problem. The planner must select beam angles, intensities, and delivery parameters to maximize tumor control probability (TCP) while minimizing normal tissue complication probability (NTCP), subject to hard dose constraints on organs at risk. The resulting search space is vast and combinatorially complex: for intensity-modulated radiotherapy (IMRT), the

Table 3: Summary of Quantum Models for Oncology Applications

Model Class	*Clinical Application*	*Key Advantage*
QNN	Cancer diagnosis, grading	Expressive feature space
QSVM	Prostate cancer, lung nodule detection	Hard-to-classical kernels
QTL	Mammography classification	Data efficiency, faster convergence
Hybrid QNN	Tumor volume prediction, setup variation	Temporal modeling, uncertainty repres
Quantum-inspired OPM	TCP/NTCP estimation	Natural uncertainty capture

number of possible beamlet weight combinations grows exponentially with the number of beamlets; for volumetric modulated arc therapy (VMAT), the continuous space of gantry angles and multileaf collimator (MLC) positions adds further complexity. Classical optimization methods—gradient descent, simulated annealing, genetic algorithms—have been successfully applied but often require extensive computation time and may converge to local optima. Quantum-aware optimization algorithms offer the potential to explore these landscapes more efficiently, leveraging superposition, entanglement, and interference to identify high-quality solutions. The Quantum Approximate Optimization Algorithm (QAOA) is a hybrid quantum-classical algorithm designed to find approximate solutions to combinatorial optimization problems [9]. QAOA operates by alternating between two Hamiltonians: a problem Hamiltonian H_C that encodes the cost function, and a mixer Hamiltonian H_B that drives exploration. Starting from an equal superposition state $|+\rangle^{\otimes n}$, the algorithm applies p layers of alternating unitaries:

$$U(\boldsymbol{\beta}, \boldsymbol{\gamma}) = e^{-i\beta_p H_B} e^{-i\gamma_p H_C} \cdots e^{-i\beta_1 H_B} e^{-i\gamma_1 H_C} H^{\otimes n},$$

where $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)$ and $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_p)$ are variational parameters optimized classically to minimize the expectation value $\langle \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) | H_C | \psi(\boldsymbol{\beta}, \boldsymbol{\gamma}) \rangle$. The final state encodes a superposition of candidate solutions, and measurement yields a sample from this superposition. In radiation oncology, QAOA has been proposed for several combinatorial subproblems:

- *Multileaf collimator (MLC) sequencing:* The problem of translating a desired fluence map into a sequence of MLC apertures is a constrained combinatorial optimization. Each aperture must satisfy leaf motion constraints, and the sequence must minimize delivery time while reproducing the target fluence. QAOA can explore the space of feasible aperture sequences more efficiently than classical search [53].

- *Beam angle optimization:* Selecting an optimal set of beam angles from a discrete set of candidates (typically 5–9 angles from 360°) is a combinatorial problem with $\binom{N}{k}$ possibilities. For $N = 36$ candidate angles and $k = 7$ beams, the search space contains $\sim 10^7$ combinations, which is intractable for exhaustive search. QAOA offers a heuristic for exploring this space, potentially identifying angle sets that achieve better target coverage and organ sparing [12].

- *Treatment scheduling:* When multiple patients compete for limited treatment slots, scheduling becomes a resource allocation problem. QAOA can be adapted to find schedules that minimize wait times while respecting machine availability and patient priorities.

- *Inverse planning with multiple objectives:* The multi-objective nature of treatment planning can be handled by scalarizing the objectives into a single cost function with appropriate weights. QAOA then optimizes this scalarized function, and the process is repeated for different weight vectors to approximate the Pareto front [53].

The approximation ratio—the quality of the solution found relative to the true optimum—improves with circuit depth p, but the scaling of this improvement for clinically relevant problem sizes remains an active area of research. Current implementations on simulators have demonstrated feasibility for small instances (e.g., 10–15 qubits), but scaling to the hundreds of variables typical of clinical planning will require advances in hardware and error mitigation.

Beyond standalone optimization, quantum algorithms serve as optimization engines within larger machine learning architectures. Quantum gradient descent refers to the use of quantum circuits to estimate gradients of cost functions with respect to model parameters, leveraging quantum parallelism to evaluate multiple parameter shifts simultaneously [53]. In the context of quantum neural networks (QNNs) applied to treatment planning, the parameters $\boldsymbol{\theta}$ might represent beamlet weights or MLC positions, and the cost function could encode dosimetric objectives and constraints. The parameter-shift rule allows gradient estimation with $O(P)$ circuit evaluations for P parameters, but quantum gradient descent aims to reduce this further by encoding multiple shifts in superposition. While theoretically promising, practical implementations are challenged by the need for coherent control and the overhead of measurement. Quantum Deep Reinforcement Learning (QDRL) extends the reinforcement learning paradigm by representing policies or value functions as parameterized quantum circuits. In the context of adaptive radiotherapy, QDRL has been applied to decision-making tasks where the agent learns to select dose-fractionation schemes based on the evolving patient state [53]. The quantum advantage arises from:

- *Superposition of actions:* Multiple candidate dose decisions can be represented simultaneously, enabling efficient exploration of the action space.

- *Compressed state representation:* High-dimensional patient states (imaging, dosimetry, biomarkers) can be encoded into a modest number of qubits, mitigating the curse of dimensionality.

- *Amplitude amplification for policy improvement:* Quantum amplitude amplification can preferentially amplify actions that lead to higher rewards, accelerating policy convergence.

Studies comparing QDRL implementations on simulators versus IBMQ hardware have demonstrated promising results, with minor performance gaps attributable to noise and decoherence [53]. The ability to run QDRL policies on near-term hardware suggests that adaptive decision support may be among the first clinical applications of quantum reinforcement learning.

A distinctive feature of quantum optimization is the ability to harness constructive and destructive interference to bias the search toward high-quality solutions. Interference-based optimization algorithms encode the cost function into a phase oracle and use quantum interference to amplify the amplitude of optimal solution states while suppressing suboptimal configurations [52]. Formally, consider a superposition over all candidate solutions: $|\psi\rangle = \sum_x \alpha_x |x\rangle$. Applying a phase oracle U_f that maps $|x\rangle \to e^{i\gamma f(x)}|x\rangle$ based on the cost function $f(x)$, followed by a mixing operation, creates interference patterns that redistribute amplitude. After multiple iterations, the probability of measuring low-cost solutions is enhanced. This mechanism underpins algorithms such as Grover search and amplitude amplification, and it can be integrated into hybrid optimization frameworks for radiotherapy planning. In clinical contexts, interference-based methods could be used to:

- *Amplify promising beam angle combinations:* By encoding the dosimetric cost of an angle set into a phase, interference can preferentially amplify sets that achieve target coverage while sparing organs at risk.
- *Suppress unsafe dose distributions:* Solutions that violate hard dose constraints can be de-amplified, ensuring that only clinically feasible plans survive the optimization process.
- *Explore Pareto-optimal trade-offs:* By varying the phase encoding, different regions of the Pareto front can be amplified, providing the planner with a diverse set of candidate plans.

However, as with other amplitude amplification techniques, the quadratic speedup of interference-based search is contingent on the ability to construct efficient oracles for the cost function, which remains challenging for complex, non-linear objectives. Quantum annealing is a metaheuristic for solving optimization problems by adiabatically evolving a quantum system from an easy-to-prepare ground state to the ground state of a problem Hamiltonian. The D-Wave quantum annealer, with thousands of qubits, has been used to explore optimization problems in various domains, including radiation oncology [14]. In quantum annealing, the system evolves under a time-dependent Hamiltonian:

$$H(t) = A(t)H_{\text{mixer}} + B(t)H_{\text{problem}},$$

where $A(t)$ decreases from a large value to zero while $B(t)$ increases from zero to a large value. If the evolution is sufficiently slow, the system remains in the instantaneous ground state, ending in the ground state of H_{problem}, which encodes the optimal solution.

Quantum tunneling—a phenomenon whereby the system passes through energy barriers rather than climbing over them—enables quantum annealers to escape local minima more effectively than classical simulated annealing, which relies on thermal fluctuations. This advantage is particularly relevant for IMRT optimization, where the cost landscape is rugged due to the interplay of multiple constraints and objectives. Hybrid quantum tunneling annealing approaches combine quantum annealing with classical refinement. In one study, such a hybrid optimizer optimized IMRT beamlet weights approximately 46% faster than classical simulated annealing while achieving comparable or better plan quality [14]. The speedup arises from the quantum annealer's ability to quickly identify promising regions of the search space, which are then fine-tuned using classical gradient-based methods. This hybrid paradigm—using quantum resources for global exploration and classical resources for local refinement—represents a pragmatic pathway to near-term clinical impact. For quantum optimization methods to be clinically useful, they must integrate seamlessly with existing treatment planning systems (TPS). This integration requires:

- *Problem encoding:* Clinical objectives (dose constraints, target coverage) must be mapped to a Hamiltonian or cost function that can be processed by quantum hardware. This often involves translating dose-volume histogram (DVH) criteria into quadratic unconstrained binary optimization (QUBO) formulations or into phase oracles.
- *Interface development:* APIs and middleware are needed to convert TPS data structures into quantum circuits or annealing problems and to decode the results back into clinically interpretable plans.
- *Validation workflows:* Any plan produced by a quantum optimizer must undergo the same quality assurance (QA) checks as a classical plan, including independent dose calculation and phantom measurements.

Early proof-of-concept studies have demonstrated the feasibility of such integration for beam angle optimization, with hybrid quantum-classical workflows achieving improved target conformity and organ-at-risk sparing compared to classical heuristics [12]. As quantum hardware matures, these prototypes could evolve into clinical tools.

Despite the promise of quantum-aware optimization, several challenges must be addressed before widespread clinical adoption:

- *Problem size:* Current quantum annealers have thousands of qubits, but the connectivity and precision are limited. For IMRT optimization with thousands of beamlets, problem decomposition or variable reduction is necessary.
- *Circuit depth and noise:* QAOA requires circuit depths that grow with problem size and desired approximation ratio. On NISQ hardware, noise limits the achievable depth, restricting QAOA to small problem instances.

- *Oracle construction:* Interference-based methods require efficient oracles for the cost function. For complex clinical objectives, constructing these oracles with low gate overhead is non-trivial.
- *Verification:* Quantum optimizers produce stochastic outputs; multiple runs are needed to obtain reliable solutions. The statistical nature of quantum sampling must be accounted for in clinical validation.
- *Hybrid workflow complexity:* Integrating quantum and classical components adds layers of complexity to the planning process, requiring specialized expertise and robust software infrastructure.

Looking forward, the most promising directions include: (i) continued development of hybrid quantum-classical algorithms that leverage the strengths of both paradigms, (ii) problem-specific ansätze and encodings that reduce resource requirements, (iii) integration of quantum optimizers as coprocessors within commercial TPS, and (iv) rigorous clinical trials to demonstrate that quantum-optimized plans translate into improved patient outcomes. The substantial speedups already demonstrated for IMRT beamlet optimization suggest that this application area may be among the first to see clinical benefit from quantum computing.

Table 4: Quantum Optimization Methods for Treatment Planning

Method	*Clinical Application*	*Key Advantage*
QAOA	MLC sequencing, beam angle selection	Heuristic for combinatorial problems
QDRL	Adaptive dose fractionation	Policy learning under uncertainty
Interference-based	Global search in plan space	Amplification of optimal solutions
Quantum annealing	IMRT beamlet optimization	Tunneling through barriers
Hybrid tunneling annealing	IMRT optimization	46% speedup demonstrated

3.1.7 Evaluation, Validation and Interpretability

The translation of quantum machine learning (QML) models from research prototypes to clinical tools requires rigorous evaluation, validation, and interpretability frameworks. Unlike classical models, where decades of methodological development have established standards for performance assessment, QML introduces unique challenges: the stochastic nature of quantum measurements, the sensitivity of results to hardware noise, and the opacity of quantum feature spaces. This subsection reviews methods for evaluating QML models, validating their outputs against clinical ground truth, and interpreting their predictions in a manner that builds clinician trust. Quantum Measurement Regression (QMR) is a physics-based regression framework that frames prediction as parameter estimation from quantum measurements [53]. In QMR, the target variable y (e.g., tumor control probability, toxicity grade) is encoded as the expectation value

of a measurement operator M on a parameterized quantum state $\rho(\mathbf{x})$:

$$\hat{y} = \mathrm{Tr}[M\rho(\mathbf{x})] = \langle M \rangle_{\rho(\mathbf{x})}.$$

The state $\rho(\mathbf{x})$ is prepared by a variational circuit that encodes the input features $\mathbf{x}$, and the parameters are trained to minimize the mean squared error between $\hat{y}$ and the true outcomes. The key advantage of QMR for interpretability lies in the structure of the measurement operator M. By decomposing M into a sum of local observables $M = \sum_i M_i$, one can attribute the prediction to different subsets of features or different regions of the input space. For example, in radiomics, M_i might correspond to a specific image texture feature, and its expectation value quantifies the contribution of that feature to the final prediction. This provides a quantum analogue of feature attribution methods like SHAP, but grounded in the physical process of measurement rather than post-hoc approximation. Furthermore, the variance of the measurement outcome, $\mathrm{Var}[M] = \langle M^2 \rangle - \langle M \rangle^2$, naturally encodes the uncertainty in the prediction. In clinical contexts where decisions must be made despite incomplete information, this uncertainty quantification is as important as the point estimate itself. QMR thus offers a unified framework for prediction and uncertainty estimation, with interpretability built into the measurement calculus.

Quantum computing's ability to explore high-dimensional Hilbert spaces makes it a natural tool for identifying diagnostic signatures in complex, multi-omics data [52]. A diagnostic signature is a pattern of features—genomic mutations, protein expression levels, radiomic texture descriptors—that reliably distinguishes between clinical states (e.g., responder vs. non-responder, early-stage vs. metastatic). Classical signature discovery methods often struggle with the curse of dimensionality and with capturing non-linear interactions among features. Quantum approaches to signature discovery leverage:

- *Quantum kernel methods:* By mapping data into exponentially large feature spaces, quantum kernels can reveal separations that are not linearly realizable in the original space. The support vectors identified by a QSVM correspond to the most informative patients, and the kernel expansion can be analyzed to identify which feature combinations are most discriminative.

- *Quantum principal component analysis:* As discussed in Section 3.1.3, qPCA identifies the directions of maximum variance in the data. The principal components—eigenvectors of the density matrix—can be interpreted as composite signatures, with large coefficients indicating features that contribute strongly to the overall variance.

- *Quantum annealing for feature subset selection:* By formulating feature selection as a quadratic unconstrained binary optimization (QUBO) problem, quantum annealers can search for optimal feature subsets that maximize predictive accuracy while minimizing redundancy. The resulting subset provides a compact, interpretable signature.

These methods do not merely produce black-box classifiers; they generate insights into the underlying biology by highlighting which variables and interactions are most predictive. However, careful validation is required to ensure that discovered signatures generalize across patient populations and are not artifacts of the quantum encoding or optimization process. A comprehensive example of quantum-enabled evaluation and validation is provided by the QProteoML framework, developed specifically to handle high-dimensional, imbalanced, and redundant oncological data [57]. QProteoML integrates multiple quantum techniques into a unified pipeline for biomarker discovery and predictive modeling:

The framework first applies quantum principal component analysis to the high-dimensional feature space (e.g., proteomic expression levels, radiomic features). qPCA identifies a low-dimensional subspace that preserves the majority of the variance, isolating salient biomarkers while discarding noise. This step is crucial for mitigating the curse of dimensionality and for improving the stability of subsequent modeling. In the reduced space, Quantum Support Vector Machines (QSVMs) are employed to model complex decision boundaries. The quantum kernel evaluated by the QSVM captures non-linear relationships among the retained features, enabling accurate prediction of outcomes such as drug sensitivity or treatment response [57]. The kernel matrix itself can be analyzed to understand which patient pairs are most similar in the quantum feature space, providing insights into patient stratification. To further refine the biomarker set, quantum annealing solves a QUBO problem that selects a minimal subset of features with maximal predictive power and minimal redundancy. The annealing process efficiently explores the combinatorial space of feature combinations, identifying signatures that might be missed by greedy classical algorithms. Oncology datasets are often small and imbalanced, with rare but clinically important subgroups (e.g., patients with specific toxicity patterns). Quantum Generative Adversarial Networks (QGANs) address this by training a quantum generator and a classical discriminator adversarially to produce synthetic data that faithfully preserves the statistical structure of the original dataset [57]. The quantum generator can capture correlations that are difficult for classical GANs to reproduce, yielding higher-quality synthetic samples. These augmented data can then be used to train more robust predictive models or to validate the stability of discovered signatures. The QProteoML framework exemplifies how quantum techniques can be combined in a principled manner to address the full pipeline from data preprocessing to model validation. Its success in proteomic biomarker discovery for multiple myeloma demonstrates the practical potential of integrated quantum approaches in precision oncology [57].

Benchmarking QML models against classical alternatives is essential for establishing their clinical utility. Studies comparing QNNs and Deep Quantum Ordinal Regressors (DQORs) to classical machine learning and deep learning models have shown promising results in diagnostic tasks. For prostate cancer grading and diabetic retinopathy classification, QNNs and DQOR models have demonstrated superior accuracy, particularly in small-sample regimes where

classical models tend to overfit [53]. The DQOR architecture, introduced in Section 2.3.3, is especially well-suited to ordinal clinical outcomes (e.g., toxicity grades 0–5, tumor stages I–IV). By combining classical convolutional feature extractors with quantum measurement layers, DQOR naturally models the ordering of classes and provides calibrated probability estimates. In comparative studies, DQOR achieved higher area under the receiver operating characteristic curve (AUC) and better calibration than classical ordinal regression models, suggesting that the quantum inductive bias aligns well with the structure of clinical data [53]. However, these comparisons must be interpreted with caution. Many published benchmarks are performed on downsampled data or on simulators with idealized noise models. As QML models are deployed on actual hardware, performance may degrade, and careful error mitigation is required. Moreover, the computational cost of training and inference—including the overhead of repeated measurements—must be factored into any fair comparison. A QML model that achieves slightly higher accuracy but requires 1000× more computational resources may not be clinically practical. Validating QML models for clinical use requires adherence to established standards for medical AI, with additional considerations specific to quantum computing:

- *Independent test sets:* Models must be evaluated on held-out data from multiple institutions to assess generalizability. Given the scarcity of quantum-ready medical datasets, this may require collaborative efforts to pool resources.
- *Prospective validation:* Simulator-based results must be complemented by prospective studies on actual quantum hardware, accounting for device-specific noise and drift. Performance metrics should be reported with confidence intervals that reflect the stochastic nature of quantum measurement.
- *Comparative effectiveness:* QML models should be compared not only to classical ML baselines but also to current clinical practice (e.g., human expert judgment, standard nomograms). A statistically significant improvement in AUC may not translate into clinically meaningful benefit if the absolute gain is small.
- *Robustness analysis:* Sensitivity to variations in input data (e.g., different imaging protocols, missing biomarkers) must be assessed. Quantum models may be more or less robust than classical models depending on the encoding scheme and circuit architecture.
- *Fairness and bias:* As with any AI system, QML models must be evaluated for performance disparities across demographic subgroups. Quantum encoding could inadvertently amplify biases present in the training data.

The opacity of quantum models poses a significant barrier to clinical adoption. Clinicians are unlikely to act on recommendations they cannot understand, especially when those recommendations involve trade-offs between tumor control

and toxicity. Several approaches to quantum interpretability are under development:

- *Quantum SHAP:* Extending the SHAP (SHapley Additive exPlanations) framework to quantum models by measuring the fidelity contribution of each encoded feature to the final prediction. This requires the ability to perturb the input state and observe changes in measurement statistics.
- *Fidelity-based attribution:* The fidelity $F = |\langle\psi(\mathbf{x})|\psi(\mathbf{x}')\rangle|^2$ between quantum states encodes their similarity. By analyzing how fidelity changes when features are masked, one can infer feature importance.
- *Measurement operator decomposition:* As in QMR, expressing the measurement operator as a sum of local terms provides a natural decomposition of the prediction into contributions from different feature groups.
- *Visualization of quantum states:* Techniques such as quantum circuit diagrams, Bloch sphere representations, and t-SNE projections of quantum states can help clinicians develop intuition for how the model processes information.

Critically, interpretability must be designed into the model from the outset, not added as an afterthought. Hybrid architectures that combine classical feature extractors with quantum layers—where the classical part is already interpretable via established methods—offer a promising path forward.

Table 5: Methods for Evaluation, Validation and Interpretability in QML for Oncology

Method	*Purpose*	*Key Features*
Quantum Measurement Regression (QMR)	Interpretable prediction with uncertainty	Measurement operator decomposition; natural uncertainty quantification
QProteoML Framework	Integrated biomarker discovery	qPCA + QSVM + quantum annealing + QGANs
Diagnostic signature discovery	Identification of predictive feature patterns	Quantum kernels, qPCA, QUBO-based feature selection
Benchmarking protocols	Performance comparison to classical models	Independent test sets, hardware validation, robustness analysis
Quantum SHAP / fidelity attribution	Post-hoc interpretability	Feature importance via perturbation analysis

In summary, evaluation, validation, and interpretability are not afterthoughts but integral components of the QML development cycle for radiation oncology.

The unique characteristics of quantum models—their stochastic outputs, sensitivity to noise, and expressive feature spaces—demand tailored approaches that go beyond classical best practices. As the field matures, the development of standardized benchmarks, open-source validation tools, and interpretability frameworks will be essential for translating quantum innovations into clinically trusted tools.

3.2 A Quantum-Inspired Clinical Decision Support System for Multi-Omics–Based Precision Radiotherapy

Quantum computing handles high-dimensional multi-omic datasets (genomics, ctDNA, proteomics, radiomics) efficiently and improves multi-omic fusion, aiding patient stratification and precision-based treatment planning [15, 53]. Quantum superposition enables concurrent exploration of multi-omic feature combinations that are computationally expensive classically [53]. CDSS integrates OPMs (prediction) and ODMs (dose decisions) for personalized treatment [14]. The flow of a Quantum Inspired Clinical Decision Support System for multi-omics-based precision radiotherapy consists of the Multi-omics data integration and pre-processing, feature selection and quantum encoding, multi-omics fusion, Quantum Inspired Predictive modeling, decision policy and optimization, clinical recommendation and validation & deployment in adaptive radiotherapy (online-during treatment sessions and offline-between treatment sessions).

3.2.1 Multi-omics Data Acquisition, Harmonization and Pre-processing

The foundation of any precision radiotherapy clinical decision support system (CDSS) is the integration of diverse data modalities that collectively capture the patient's unique disease state. Modern oncology generates an unprecedented volume of multi-omics data: genomic sequencing (DNA mutations, copy number variations), transcriptomic profiles (gene expression), proteomic signatures (protein abundance and modifications), metabolomic measurements, and radiomic features extracted from medical images. Each modality offers a distinct lens into tumor biology, but their integration poses formidable challenges: data heterogeneity, differing scales and units, batch effects, missing values, and the need for harmonization before they can be jointly analyzed. Quantum and quantum-inspired methods offer novel approaches to these pre-processing tasks, potentially preserving more of the underlying biological signal while reducing dimensionality for downstream quantum encoding. Before any quantum model can be applied, raw multi-omics data must undergo a series of pre-processing steps:

- *Data acquisition:* Samples must be collected under standardized protocols to minimize technical variability. For genomics, this involves DNA sequencing and variant calling; for proteomics, mass spectrometry or antibody arrays; for radiomics, segmentation and feature extraction from medical images.

- *Harmonization:* Batch effects—systematic technical variations introduced by different processing times, instruments, or laboratories—must be removed to ensure that biological signals are not confounded by experimental artifacts. Classical methods such as ComBat and limma are commonly used, but they may not capture complex, non-linear dependencies.
- *Normalization:* Data from different modalities exist on different scales (e.g., gene expression counts, protein intensities, radiomic texture values). Normalization to a common scale (e.g., z-scores, unit vectors) is required for subsequent integration.
- *Missing data imputation:* Multi-omics datasets are often incomplete, with certain assays missing for some patients. Imputation must respect the correlations both within and across modalities.
- *Dimensionality reduction:* The combined feature space can easily exceed $10^4 - 10^5$ dimensions, far beyond what can be directly encoded into quantum states with current hardware. Feature selection or compression is essential.

Quantum and quantum-inspired techniques can contribute at each of these stages, offering advantages in capturing cross-modal correlations and in handling the inherent uncertainty of biological measurements.

The Quantum Wavelet Transform (QWT) provides a framework for multi-scale analysis of omics data that can be particularly valuable for harmonization [53]. Wavelet transforms decompose signals into approximation and detail coefficients at multiple resolutions, enabling the separation of biological signal from technical noise. The QWT implements this decomposition using controlled rotations and the quantum Fourier transform, achieving exponential speedup over classical wavelet transforms for amplitude-encoded data. In the context of multi-omics harmonization, QWT can be applied to:

- *Denoise individual omics layers:* By thresholding wavelet coefficients corresponding to high-frequency noise, each modality can be cleaned while preserving low-frequency biological structure.
- *Align feature resolutions:* Different omics modalities operate at different biological scales (e.g., genomic variants are discrete, while proteomic measurements are continuous). Wavelet decomposition can bring them to a common multi-scale representation, facilitating integration.
- *Identify batch effect patterns:* Batch effects often manifest as low-frequency artifacts across samples. By analyzing wavelet coefficients, one can identify and remove components that correlate with batch labels rather than biological variables.

While theoretical, these applications illustrate how quantum signal processing could enhance the quality of harmonized multi-omics data before quantum encoding. Quantum Signal Processing (QSP) is a framework for implementing

polynomial transformations of eigenvalues of a unitary operator. In the context of omics data, QSP can be used to filter noise and enhance signals of interest [52]. For a data matrix X encoded in a quantum state ρ, QSP allows the application of spectral filters that amplify components corresponding to large singular values (biological signal) while suppressing those corresponding to small singular values (noise). This is analogous to classical singular value decomposition (SVD) filtering but can be performed more efficiently on quantum hardware for large matrices.

For multi-omics integration, QSP could be applied to the combined data matrix to extract a low-rank approximation that captures the dominant biological variations across modalities. This serves both as a denoising step and as a form of dimensionality reduction, producing a compact representation suitable for downstream quantum modeling. Mutual information (MI) is a fundamental quantity for measuring the statistical dependence between variables. In multi-omics analysis, MI can be used to:

- Identify features that are strongly associated with clinical outcomes.
- Quantify redundancy between features to guide feature selection.
- Assess the degree of integration between different omics layers: high MI between modalities suggests coordinated biological processes.

Classical MI estimation becomes challenging in high dimensions due to the need for density estimation or binning. Quantum Mutual Information (QMI) estimation leverages the ability of quantum computers to represent joint probability distributions as density matrices and to compute entropies via von Neumann entropy:

$$I(X:Y) = S(\rho_X) + S(\rho_Y) - S(\rho_{XY}),$$

where $S(\rho) = -\mathrm{Tr}(\rho \log \rho)$ is the von Neumann entropy. For data encoded into quantum states, QMI can be estimated using quantum phase estimation and tomography, potentially offering more accurate estimates in high-dimensional spaces [15]. In a CDSS pipeline, QMI could be used to:

1. *Guide feature selection:* Features with high QMI relative to the clinical outcome are retained; those with high QMI relative to other features (redundancy) are pruned.
2. *Validate harmonization:* After batch correction, QMI between technical covariates and omics features should be minimized, while QMI between biological variables and outcomes should be preserved or enhanced.
3. *Discover cross-omics interactions:* Pairs of features from different modalities with unexpectedly high QMI may indicate biological pathways that span multiple omics layers.

Once harmonized and pre-processed, the multi-omics data must be encoded into quantum states for input to the CDSS. The choice of encoding depends on the dimensionality after pre-processing and the quantum model to be used:

- *Amplitude encoding:* If the pre-processed feature vector has been reduced to dimension $N = 2^n$, it can be amplitude-encoded into an n-qubit state. This is the most efficient use of qubits but requires that the data be normalized and that the state preparation circuit be implementable.
- *Angle encoding:* For moderate-dimensional data ($n \sim 50 - 100$), angle encoding on individual qubits is simpler and more noise-resilient. This is appropriate if feature selection has reduced the set to a manageable size.
- *Hierarchical encoding:* Different omics layers can be encoded into separate registers and entangled to capture cross-modal correlations. For example, genomic data might be angle-encoded into one register, radiomic features into another, with entangling gates between them to model interactions.

The pre-processing steps described above directly influence the feasibility of these encodings. For instance, QWT-based denoising may enable more faithful amplitude encoding by reducing the dynamic range of the data, while QMI-guided feature selection ensures that only the most informative features are retained for angle encoding. The output of this pre-processing stage is a set of quantum-ready data representations: either quantum states ρ_i for each patient, or a combined density matrix ρ_{cohort} representing the patient population. These states feed directly into the subsequent stages of the quantum-inspired CDSS: feature extraction and selection (Section 3.2.2), multi-omics fusion (Section 3.2.3), and predictive modeling (Section 3.2.4). The quality of the pre-processing determines the upper bound on the performance of all downstream components; noisy, unharmonized data cannot be rescued by even the most sophisticated quantum model.

Table 6: Quantum Techniques for Multi-omics Pre-processing

Technique	*Application*	*Advantage*
Quantum Wavelet Transform (QWT)	Multi-scale harmonization, denoising	Exponential speedup for amplitude-encoded data; separation of signal from noise
Quantum Signal Processing (QSP)	Spectral filtering, noise reduction	Efficient polynomial transformations of singular values
Quantum Mutual Information (QMI)	Dependency analysis, feature selection	Accurate entropy estimation in high dimensions
Hybrid encoding strategies	Quantum state preparation	Balances qubit efficiency with noise resilience

In summary, the data acquisition, harmonization, and pre-processing stage is the critical first step in any quantum-enabled multi-omics CDSS. By leveraging

quantum techniques such as QWT, QSP, and QMI, it is possible to clean, integrate, and compress heterogeneous biological data into forms that are amenable to quantum encoding and analysis. While many of these techniques are still theoretical or demonstrated only on small-scale problems, they point toward a future where quantum computers contribute not only to prediction but also to the foundational data processing that makes prediction possible.

3.2.2 Feature Extraction, Selection and Quantum Encoding

Once multi-omics data have been acquired, harmonized, and pre-processed, the next critical step in a quantum-inspired clinical decision support system (CDSS) is the transformation of these heterogeneous data into a compact, informative representation suitable for quantum modeling. This stage encompasses three interrelated tasks: (i) extracting relevant features from each omics modality, (ii) selecting a subset of features that are most predictive of clinical outcomes while minimizing redundancy, and (iii) encoding the selected features into quantum states for subsequent processing. Quantum and quantum-inspired methods offer novel approaches to each of these tasks, potentially capturing complex, non-linear relationships that classical methods might miss. Before feature selection can begin, the raw omics data often reside in spaces of extremely high dimensionality—gene expression arrays with $20,000+$ features, proteomic mass spectrometry with tens of thousands of peaks, radiomic feature sets with thousands of texture and shape descriptors. Direct application of selection algorithms to such high-dimensional spaces is computationally prohibitive and statistically unstable. Dimensionality reduction is therefore an essential first step. Quantum Principal Component Analysis (qPCA), introduced in Section 3.1.3, provides a framework for identifying the directions of maximum variance in high-dimensional data using quantum phase estimation [53]. For a dataset encoded as a density matrix $\rho = \frac{1}{M}\sum_{j=1}^{M} |\mathbf{x}_j\rangle\langle\mathbf{x}_j|$, qPCA efficiently estimates the eigenvectors corresponding to the largest eigenvalues—the principal components. These components capture the dominant modes of variation across the multi-omics dataset, effectively compressing the information while preserving the structure most relevant for downstream analysis. In the context of a multi-omics CDSS, qPCA can be applied in several ways:

- *Cross-modal compression:* By constructing a combined density matrix that incorporates all omics modalities, qPCA identifies principal components that reflect coordinated variation across genomics, proteomics, and radiomics. These components may correspond to underlying biological processes that manifest across multiple molecular layers.

- *Modality-specific compression:* Alternatively, qPCA can be applied separately to each omics type, producing a reduced representation for each modality that retains the majority of its variance. These reduced representations can then be concatenated or fused in subsequent steps.

- *Noise filtering:* The eigenvectors associated with small eigenvalues cor-

respond to noise or technical artifacts. By projecting the data onto the subspace spanned by the top k principal components, one effectively filters out this noise, enhancing the signal-to-noise ratio for downstream modeling.

The exponential speedup offered by qPCA in terms of the feature dimension N makes it particularly attractive for multi-omics datasets, where N can be in the tens of thousands. However, as noted previously, practical implementation on NISQ hardware faces challenges in state preparation and tomography, so near-term applications may rely on quantum-inspired classical algorithms that mimic qPCA using tensor networks or other techniques.

Following dimensionality reduction, feature selection aims to identify a subset of the most informative features—whether original features or principal components—that maximize predictive performance while minimizing redundancy. This is a combinatorial optimization problem: given N features, there are 2^N possible subsets, and exhaustive search is infeasible for any clinically realistic N. Quantum Genetic Algorithms (QGAs) offer a heuristic approach to this problem, leveraging quantum principles to enhance the exploration of the feature subset space [53]. In a QGA, each candidate feature subset is represented as a quantum state, with qubits indicating whether a feature is included (1) or excluded (0). The population of subsets evolves through quantum analogues of selection, crossover, and mutation:

- *Quantum selection:* Subsets with higher fitness (e.g., higher cross-validated accuracy) are assigned higher amplitudes, increasing their probability of being measured and propagated to the next generation.
- *Quantum crossover:* Entangling operations combine the states of two parent subsets to produce superpositions of offspring, exploring combinations of features that might not be present in either parent alone.
- *Quantum mutation:* Single-qubit rotations introduce small perturbations, allowing the algorithm to escape local optima by exploring nearby subsets.

The quantum advantage in this context arises from superposition: multiple subsets are evaluated in parallel, and the interference patterns that emerge from quantum operations can guide the search more efficiently than classical genetic operators. While most QGA implementations to date have been on classical simulators, they have demonstrated improved convergence rates on benchmark feature selection problems compared to classical genetic algorithms [53]. A more direct quantum approach to feature selection leverages Grover's algorithm, which provides a quadratic speedup for unstructured search. In the context of biomarker selection, one can formulate the problem as searching for a feature subset that maximizes a given objective function (e.g., mutual information with the outcome, classification accuracy). Bisection Grover's Search adapts the standard Grover algorithm to handle the case where the number of marked items (good feature subsets) is unknown [15].

The algorithm proceeds as follows:

1. The feature space is divided into intervals based on subset size or some other heuristic.
2. Grover search is applied within each interval to find subsets that exceed a fitness threshold.
3. The threshold is adjusted iteratively (bisection) to converge on the optimal or near-optimal subset.

Bisection Grover's Search has been proposed for quantum-accelerated biomarker selection in high-dimensional radiomic and multi-omic spaces, where it can efficiently identify optimal feature subsets while reducing combinatorial complexity prior to predictive modeling [15]. The quadratic speedup over classical exhaustive search is particularly valuable when the number of features is moderate ($N \sim 30 - 50$)—a regime that becomes accessible after initial dimensionality reduction via qPCA or similar methods. Quantum-Inspired Immune Clone Optimization (QICO) offers a classical alternative that nevertheless draws inspiration from quantum principles. As described in Section 3.1.2, QICO combines concepts from quantum computing (superposition, measurement collapse) with artificial immune systems (clonal selection, affinity maturation) to perform efficient feature selection [53]. In QICO, each antibody represents a candidate feature subset, and the population of antibodies evolves through:

- *Clonal expansion:* Antibodies with higher affinity (fitness) are cloned with a probability proportional to their measurement probability in a quantum-inspired representation.
- *Hyper-mutation:* Clones undergo mutation at rates that depend on affinity, with higher-affinity clones mutating less to preserve good solutions.
- *Quantum-inspired superposition:* The representation of antibodies as quantum-inspired probability amplitudes allows the algorithm to maintain diversity and explore multiple regions of the feature space simultaneously.

QICO has demonstrated superior performance compared to classical discovery-based feature selection methods in cancer classification tasks, efficiently identifying optimal subsets from large biological feature spaces and supporting precision-oriented oncology modeling [53]. Its advantage lies in balancing exploration (via superposition-inspired diversity) and exploitation (via clonal selection), making it well-suited to the rugged fitness landscapes typical of multi-omics feature selection. An alternative to explicit feature selection is to use quantum feature maps that implicitly weight features through the structure of the quantum circuit. In quantum kernel methods, the feature map $\phi : \mathcal{X} \to \mathcal{H}$ embeds the input data into a quantum Hilbert space, and the kernel $k(\mathbf{x}, \mathbf{x}') = |\langle \phi(\mathbf{x}) | \phi(\mathbf{x}') \rangle|^2$ determines the similarity between samples [15]. The choice of feature map—the specific sequence of parameterized gates and entangling operations—implicitly defines which feature combinations are emphasized and which

are suppressed. In the context of a multi-omics CDSS, one can design quantum feature maps that:

- *Weight features by clinical relevance:* By parameterizing the encoding gates, the model can learn to amplify the contribution of features that are most predictive of the outcome while attenuating irrelevant ones.
- *Capture cross-omics interactions:* Entangling gates between qubits encoding different omics modalities allow the kernel to reflect interactions between, for example, a genomic mutation and a radiomic texture feature.
- *Perform implicit dimensionality reduction:* If the feature map is chosen such that the effective dimension of the quantum feature space is much smaller than the classical input dimension, the kernel implicitly performs a form of nonlinear dimensionality reduction.

The advantage of this approach is that feature selection becomes integrated into the model training process, rather than being a separate pre-processing step. However, it comes at the cost of increased complexity and reduced interpretability, as the contribution of individual features is harder to disentangle from the quantum circuit. Once a feature subset has been selected—whether via qPCA, QGA, Grover search, QICO, or implicit kernel methods—the features must be encoded into quantum states for input to the CDSS. The encoding strategy depends on the number of selected features and the quantum model to be used:

- *Amplitude encoding:* If the number of selected features $N_f = 2^n$ (after possible zero-padding), amplitude encoding into an n-qubit state is the most qubit-efficient option. This is suitable when feature selection has reduced the dimension to a few hundred or thousand, which can be encoded in $\log_2 N_f$ qubits.
- *Angle encoding:* For smaller feature sets ($n \sim 20-100$), angle encoding on individual qubits is simpler and more noise-resilient. This is appropriate if feature selection has identified a compact biomarker panel of moderate size.
- *Hierarchical encoding:* Features from different omics modalities can be encoded into separate quantum registers, with entangling gates between registers to model cross-modal interactions. This preserves modality identity and allows for modality-specific pre-processing.
- *Quantum random access memory (QRAM):* For very large feature sets that cannot be reduced further, QRAM provides a theoretical mechanism for loading classical data in superposition. However, practical QRAM remains a long-term goal.

The encoding choice also influences the design of subsequent quantum models. For amplitude-encoded data, kernel methods or variational circuits that respect the global structure of the state are appropriate. For angle-encoded data, local

variational circuits with parameterized rotations are more natural. The output of this stage is a set of quantum-ready feature representations: for each patient, a quantum state $|\psi_i\rangle$ that encodes the selected, compressed multi-omics profile. These states are then passed to the multi-omics fusion and representation learning stage (Section 3.2.3), where they may be combined or transformed to capture higher-order interactions. The quality of feature extraction and selection directly determines the information content available to all downstream components; a poorly chosen feature set cannot be compensated by even the most sophisticated quantum model.

Table 7: Quantum Methods for Feature Extraction, Selection and Encoding

Method	*Stage*	*Key Advantage*
Quantum PCA (qPCA)	Extraction	Exponential speedup in feature dimension; identifies principal modes of variation
Quantum Genetic Algorithm (QGA)	Selection	Superposition-based exploration of subset space
Bisection Grover's Search	Selection	Quadratic speedup for optimal subset identification
Quantum-Inspired Immune Clone Optimization (QICO)	Selection	Balances exploration and exploitation; demonstrated performance on biological data
Quantum Feature Maps / Kernels	Implicit selection	Integrates selection with model training; captures cross-omics interactions
Amplitude / Angle / Hierarchical Encoding	Encoding	Balances qubit efficiency with noise resilience and modality structure

In summary, feature extraction, selection, and encoding form the bridge between raw multi-omics data and quantum-ready representations in a CDSS. By leveraging quantum and quantum-inspired methods at each step, it is possible to reduce dimensionality, identify clinically relevant biomarkers, and prepare data in forms that maximally exploit the representational power of quantum models. The choice of methods must be guided by the specific characteristics of the data, the constraints of available hardware, and the clinical questions to be answered.

3.2.3 Multi-omics Fusion and Representation Learning

Following feature extraction, selection, and encoding, the next critical stage in a quantum-inspired clinical decision support system (CDSS) is the fusion of disparate omics modalities into a unified representation that captures their synergistic interactions. Modern radiotherapy outcome models increasingly rely on

integrating genomics, radiomics, proteomics, and transcriptomics to capture the full complexity of tumor biology and patient-specific treatment response [14]. However, classical fusion methods often struggle to model the higher-order interactions that span multiple omics layers, typically assuming linear relationships or limited non-linearities that may not reflect biological reality. Quantum and quantum-inspired approaches offer a fundamentally different paradigm: representing multi-omics data in exponentially large Hilbert spaces where correlations across modalities can be naturally encoded and learned. Multi-omics integration faces several fundamental challenges:

- *Heterogeneity:* Different omics modalities have distinct statistical properties—genomic data are discrete and sparse, transcriptomic data are continuous but often skewed, radiomic features are derived from images and exhibit spatial correlations. Fusing these disparate data types into a common representation without losing modality-specific structure is non-trivial.
- *Dimensionality:* The combined feature space can exceed 10^5 dimensions, far beyond what classical models can process without aggressive dimensionality reduction, which may discard biologically relevant signals.
- *Higher-order interactions:* Biological pathways involve complex interactions that may span multiple omics layers—for example, a genomic mutation may alter protein expression, which in turn affects cellular metabolism and ultimately manifests as a radiomic texture feature on imaging. Capturing these multi-way interactions requires models that can represent correlations beyond pairwise relationships.
- *Missing modalities:* In clinical practice, not all omics assays may be available for every patient. Fusion methods must be robust to missing data and able to leverage whatever modalities are present.

Quantum feature spaces address these challenges by providing a high-dimensional Hilbert space in which data from different modalities can be embedded and correlated through entanglement. Quantum Neural Networks (QNNs) provide a flexible framework for learning representations that fuse multiple input modalities. In a multi-omics QNN, each omics type is first encoded into a separate quantum register using an appropriate encoding scheme (for instance, angle encoding for selected features, amplitude encoding for compressed representations). These registers are then entangled through parameterized quantum circuits that learn to capture cross-modal correlations. Formally, let $|\psi_g(\mathbf{x}_g)\rangle$, $|\psi_p(\mathbf{x}_p)\rangle$, and $|\psi_r(\mathbf{x}_r)\rangle$ represent the quantum states encoding genomic, proteomic, and radiomic data, respectively. A fusion circuit $U_{\text{fuse}}(\boldsymbol{\theta})$ acts on the joint state:

$$|\Psi_{\text{fused}}\rangle = U_{\text{fuse}}(\boldsymbol{\theta})\left(|\psi_g\rangle \otimes |\psi_p\rangle \otimes |\psi_r\rangle\right).$$

The resulting entangled state $|\Psi_{\text{fused}}\rangle$ lives in the tensor product Hilbert space $\mathcal{H}_g \otimes \mathcal{H}_p \otimes \mathcal{H}_r$, whose dimension is the product of the individual dimensions. Even with modest numbers of qubits per modality, this joint space can be exponentially large, providing ample capacity for representing complex multi-way

interactions. The parameters $\boldsymbol{\theta}$ of the fusion circuit are trained end-to-end to minimize a loss function that reflects the clinical task (e.g., predicting treatment response, toxicity grade). During training, the circuit learns which cross-modal correlations are most predictive, effectively performing a form of representation learning that is guided by clinical outcomes. A key advantage of this approach is that it naturally handles missing modalities: if a particular omics assay is unavailable for a patient, the corresponding register can be initialized in a maximally mixed state or a fixed reference state, and the fusion circuit can still produce a meaningful representation using the available data. This robustness is essential for clinical deployment, where complete multi-omics profiles are the exception rather than the rule.

While QNNs treat omics data as flat feature vectors, some modalities have inherent structure that can be exploited. Radiomic features, for example, are derived from images and retain spatial relationships; genomic data may have known pathway structures. Quantum Convolutional Neural Networks (QCNNs) can leverage such structure through locality-preserving circuits. For radiomics, a QCNN can process 2D or 3D image patches directly, extracting features at multiple scales through hierarchical entanglement and pooling. These radiomic features can then be fused with genomic and proteomic data at a higher level of abstraction. The translational invariance of convolutional layers ensures that the model learns features that are robust to tumor position and orientation. For genomics, one can design QCNNs that operate on 1D sequences (e.g., along the chromosome) or on graph structures representing gene regulatory networks. By respecting the underlying biology in the circuit architecture, these models can learn representations that are more interpretable and data-efficient than fully connected alternatives. The fusion of QCNN-extracted features with other omics modalities can occur at multiple levels:

- *Early fusion:* Radiomic features extracted by a QCNN are concatenated with genomic and proteomic features before being passed to a joint QNN.
- *Mid-level fusion:* Intermediate representations from modality-specific QCNNs are combined in a fusion layer that learns cross-modal correlations.
- *Late fusion:* Predictions from modality-specific models are combined through a voting or averaging mechanism.

Early and mid-level fusion are generally more powerful as they allow cross-modal interactions to influence the learned representations, but they require more sophisticated circuit design and training. An alternative to discriminative QNNs is provided by Quantum Bayesian Networks (QBNs), which model the joint probability distribution over multi-omics features and clinical outcomes in a quantum framework. In a QBN, each random variable (e.g., expression level of a gene, presence of a mutation, radiomic texture feature) is associated with a quantum state, and conditional dependencies are represented through entangling operations. The quantum analogue of a classical Bayesian network

is a quantum circuit that prepares a state $|\Psi\rangle$ such that measurements on subsets of qubits yield marginal distributions consistent with the network structure. For a given graph structure G encoding conditional independencies, one can design a circuit that generates a quantum state that factorizes according to G:

$$|\Psi\rangle = \bigotimes_i U_{i|\mathrm{pa}(i)}|0\rangle,$$

where $U_{i|\mathrm{pa}(i)}$ are unitaries that entangle variable i with its parents $\mathrm{pa}(i)$ in the graph. Once the QBN is trained (i.e., the parameters of the conditional unitaries are learned from data), it can be used for inference: given evidence on some variables (e.g., observed genomic and radiomic features), one can condition on that evidence through post-selection or amplitude amplification and sample the posterior distribution over unobserved variables (e.g., treatment response). QBNs offer several advantages for multi-omics fusion:

- *Uncertainty quantification:* The quantum state naturally encodes a full probability distribution, providing not just point predictions but also measures of uncertainty.
- *Causal reasoning:* By incorporating causal structure (e.g., known biological pathways) into the network, QBNs can support counterfactual reasoning about interventions (e.g., "What would be the effect of inhibiting this pathway?")
- *Handling missing data:* Inference in a QBN can proceed even when some variables are unobserved, by marginalizing over the corresponding qubits.

However, training QBNs is challenging due to the need to learn both the graph structure and the conditional unitaries, and inference may require many circuit repetitions to achieve reliable estimates. Beyond explicit fusion architectures, quantum embedding circuits provide a mechanism for learning compact representations of multi-omics data that can be used with classical classifiers or similarity-based methods. A quantum embedding is a parameterized circuit $U_{\boldsymbol{\theta}}(\mathbf{x})$ that maps input data $\mathbf{x}$ (which may already be a fused multi-omics vector) to a quantum state $|\phi_{\boldsymbol{\theta}}(\mathbf{x})\rangle$. The inner product between two such states defines a quantum kernel:

$$k_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{x}') = |\langle\phi_{\boldsymbol{\theta}}(\mathbf{x})|\phi_{\boldsymbol{\theta}}(\mathbf{x}')\rangle|^2.$$

This kernel can be used in classical kernel methods (e.g., support vector machines, Gaussian processes) to perform prediction in the quantum feature space. The advantage is that the quantum circuit learns a data-dependent representation that is optimized for the task, while the final prediction is made classically, avoiding the need for quantum measurement during inference. In the context of multi-omics fusion, the embedding circuit can be designed to process each modality separately before entangling them, effectively learning a representation that captures cross-modal interactions. The resulting kernel can then be

used for tasks such as patient stratification, outcome prediction, or treatment response modeling. Quantum embedding circuits have been proposed for modeling patient-specific radiosensitivity, where the goal is to predict an individual's likelihood of benefiting from radiotherapy based on their multi-omics profile [14]. By learning embeddings that place similar patients (in terms of treatment response) close together in the quantum feature space, these methods can provide compact, interpretable representations that support clinical decision-making. Implementing multi-omics fusion on quantum hardware faces several practical challenges:

- *Qubit requirements:* Even with aggressive feature selection, encoding multiple omics modalities may require more qubits than are currently available. Hybrid approaches that use classical pre-fusion to reduce dimensionality before quantum encoding are often necessary.
- *Circuit depth:* Fusion circuits that entangle multiple registers can become deep, especially if they aim to capture all pairwise interactions. This depth increases sensitivity to noise and decoherence, limiting what can be implemented on NISQ devices.
- *Training complexity:* End-to-end training of fusion circuits requires gradient estimation via the parameter-shift rule, which may be sample-inefficient for deep circuits. Layer-wise training or pre-training of modality-specific components may be needed.
- *Interpretability:* While quantum representations are expressive, they are also opaque. Developing methods to extract clinically meaningful insights from learned quantum states is an ongoing challenge.
- *Validation:* Demonstrating that quantum-fused representations outperform classical alternatives requires careful benchmarking on clinically realistic datasets, with attention to both predictive accuracy and computational cost.

Despite these challenges, the potential of quantum methods to capture higher-order interactions across omics layers—interactions that may be invisible to classical models—makes this a vibrant area of research with significant implications for precision radiotherapy.

In summary, multi-omics fusion and representation learning represent the computational core of a quantum-inspired CDSS. By leveraging quantum principles to encode and correlate heterogeneous biological data, these methods aim to capture the full complexity of tumor biology and patient-specific treatment response. While still in early stages of development, they point toward a future where the integration of genomics, proteomics, radiomics, and other modalities is not a pre-processing afterthought but a fundamental part of the learning process, enabled by the unique representational power of quantum systems.

Table 8: Quantum Methods for Multi-omics Fusion and Representation Learning

Method	*Approach*	*Key Advantage*
Quantum Neural Networks (QNNs)	Entangled registers for each modality	End-to-end learning of cross-modal correlations; handles missing modalities
Quantum Convolutional Neural Networks (QCNNs)	Locality-preserving circuits for structured data	Exploits spatial/structure in radiomics and genomics
Quantum Bayesian Networks (QBNs)	Quantum circuits encoding probabilistic graphical models	Full uncertainty quantification; causal reasoning
Quantum Embedding Circuits / Kernels	Learned representations for classical kernel methods	Data-dependent similarity measures; compact representations
Hybrid pipelines	Classical pre-fusion + quantum fusion	Reduces qubit requirements; pragmatic for NISQ era

3.2.4 Predictive Modeling, Uncertainty Quantification, and Risk-Aware Decision Optimization

The ultimate objective of a quantum-inspired clinical decision support system (CDSS) for multi-omics precision radiotherapy is not merely to predict outcomes, but to guide clinical decisions under uncertainty. This requires an integrated framework that encompasses three interconnected capabilities: (i) predictive modeling that leverages quantum feature spaces to capture complex, higher-order interactions across omics layers; (ii) uncertainty quantification that distinguishes between aleatoric uncertainty (inherent randomness in outcomes) and epistemic uncertainty (lack of knowledge due to limited data); and (iii) risk-aware decision optimization that formulates treatment policies accounting for both predicted outcomes and their associated uncertainties. This subsection synthesizes these capabilities into a unified view of how quantum and quantum-inspired methods can support clinical decision-making. At the core of the CDSS lie predictive models that map a patient's multi-omics profile to clinically relevant outcomes—tumor control probability (TCP), normal tissue complication probability (NTCP), treatment response, toxicity grades, or survival endpoints. Quantum and quantum-inspired approaches offer several paradigms for this task, each with distinct advantages. *Quantum Support Vector Machines (QSVMs)* provide a kernel-based approach to classification and regression. By embedding multi-omics data into quantum feature spaces via parameterized circuits, QSVMs can capture non-linear decision boundaries that reflect complex interactions among genomic, proteomic, and radiomic features [53]. The quantum kernel $k(\mathbf{x}, \mathbf{x}') = |\langle\phi(\mathbf{x})|\phi(\mathbf{x}')\rangle|^2$ is estimated on quantum hardware and then fed to a classical SVM, combining quantum expressivity with

classical training efficiency. For multi-omics data, the feature map ϕ can be designed to emphasize cross-modal correlations, effectively learning a similarity measure that reflects biological pathway interactions.

Quantum Neural Networks (QNNs) offer an end-to-end variational approach in which a parameterized quantum circuit directly outputs predictions. For multi-omics inputs, QNNs can be structured with separate registers for each modality, entangled through learnable layers that capture higher-order interactions. The output is obtained through measurement of observables corresponding to clinical endpoints. QNNs are particularly suited to regression tasks (e.g., predicting continuous toxicity scores) and to settings where the relationship between inputs and outputs is highly non-linear [53]. *Deep Quantum Ordinal Regressors (DQORs)* address the common clinical scenario where outcomes are ordered categorical variables—toxicity grades (0–5), tumor stages (I–IV), or response categories (complete response, partial response, stable disease, progression). DQOR combines classical convolutional neural networks (for feature extraction from imaging data) with quantum measurement layers that naturally model the ordinal structure [53]. The quantum component outputs a probability distribution over the ordered classes, with the measurement operator designed to respect the ordering (e.g., through a cumulative link model). This yields predictions that are not only accurate but also well-calibrated, a critical property for clinical decision-making. *Quantum state-coupled gene regulatory network (qscGRN) models* represent a specialized approach for predictive analysis of complex gene regulatory networks [15]. These models encode the activity of genes as quantum states, with regulatory interactions represented through entangling operations. By learning the parameters of the quantum circuit from expression data, qscGRN models can capture higher-order interactions among genomic features that go beyond pairwise gene relationships—interactions that are central to radiation response but difficult to model classically. The learned quantum state encodes not only the marginal expression levels but also the correlations that reflect regulatory logic. When applied to predicting radiation response, qscGRN models have demonstrated the ability to reveal latent regulatory patterns that correlate with treatment outcomes, supporting improved patient stratification [15].

The unifying theme across these approaches is the use of quantum feature spaces to represent multi-omics data in a form that makes predictive patterns more accessible. Whether through kernels, variational circuits, or state-based models, quantum methods offer a principled way to capture the complex, non-linear, higher-order interactions that characterize biological systems. A prediction without an accompanying measure of confidence is of limited value in clinical decision-making. Clinicians must weigh the likelihood of benefit against the risk of harm, and this requires knowing not only the expected outcome but also the uncertainty around that expectation. Quantum methods offer novel approaches to uncertainty quantification that are naturally integrated with the predictive model. *Quantum Bayesian inference* provides a framework for updating beliefs

about patient outcomes as new data become available [14]. In this approach, prior clinical knowledge (e.g., population-level outcome distributions) is encoded as a quantum state $|\psi_{\text{prior}}\rangle$. Patient-specific multi-omics data are then used to condition this state through a quantum circuit that implements Bayes' rule:

$$|\psi_{\text{posterior}}\rangle \propto U_{\text{likelihood}}(\mathbf{x})|\psi_{\text{prior}}\rangle,$$

where $U_{\text{likelihood}}$ encodes the probability of observing the data given the outcome. Measurements on the posterior state yield predictive distributions that reflect both the prior and the evidence. Crucially, the quantum formalism naturally distinguishes between aleatoric uncertainty (reflected in the spread of measurement outcomes) and epistemic uncertainty (reflected in the purity of the quantum state). A mixed state indicates high epistemic uncertainty—the model is unsure because it has seen insufficient or conflicting data—while a pure state indicates confidence. *Quantum amplitude estimation (QAE) for Monte Carlo methods* accelerates the computation of probabilities and expectations that are central to uncertainty quantification [52]. In classical Monte Carlo, estimating the probability p of an event (e.g., grade ≥ 3 toxicity) to precision ϵ requires $O(1/\epsilon^2)$ samples. QAE achieves the same precision with $O(1/\epsilon)$ queries, a quadratic speedup. For a CDSS that needs to compute many such probabilities (e.g., for different dose levels, fractionation schemes, or supportive care interventions), this speedup could be clinically significant, enabling more thorough exploration of the decision space within the time constraints of a clinical workflow.

Quantum Monte Carlo (QMC) methods extend this idea to the estimation of integrals and expectations over high-dimensional spaces. In the context of treatment outcome modeling, QMC can accelerate the computation of TCP and NTCP by integrating over distributions of patient characteristics, dosimetric parameters, and biological factors [14]. The resulting estimates come with built-in uncertainty quantification, as the quantum algorithm can be designed to output both the mean and the variance of the estimated quantity. These uncertainty quantification techniques serve a critical function in the CDSS: they flag cases where the model's predictions are too uncertain to support a confident recommendation. In such cases, the system might recommend additional data collection (e.g., a confirmatory biopsy, an additional imaging study) or revert to a conservative default treatment strategy. This risk-aware behavior is essential for building clinician trust and ensuring patient safety. The final stage of the CDSS integrates predictive models and uncertainty estimates into a decision-making framework that selects optimal treatment policies. Unlike traditional approaches that first make a prediction and then separately apply a decision rule, quantum-enhanced decision systems incorporate uncertainty directly into the optimization process, formulating risk-aware policies rather than relying on post-hoc adjustments [14].

Quantum Reinforcement Learning (QRL) treats treatment planning as a sequential decision problem under uncertainty. The patient's state—encoded as

a quantum state incorporating multi-omics data, treatment history, and current clinical status—evolves over time in response to interventions. A quantum policy network, implemented as a parameterized quantum circuit, maps states to actions (e.g., dose per fraction, choice of modality, adaptation timing). The policy is optimized to maximize a cumulative reward that reflects both tumor control and normal tissue sparing, with penalties for high-uncertainty decisions. The quantum advantage in QRL arises from:

- *Superposition of actions:* The policy network can output a superposition over possible actions, enabling efficient exploration of the action space without sacrificing performance.
- *Compressed state representation:* Quantum states can represent high-dimensional patient information in a compact form, mitigating the curse of dimensionality that plagues classical RL in complex clinical domains.
- *Amplitude amplification for policy improvement:* Actions that lead to higher expected rewards can be preferentially amplified through Grover-like operations, accelerating convergence to optimal policies.

Quantum Deep Reinforcement Learning (QDRL) extends this framework by using classical neural networks to process raw observations (e.g., medical images) before passing compressed features to a quantum policy network [53]. This hybrid architecture leverages the representational power of deep learning for feature extraction while retaining the quantum advantage in policy representation and exploration. *Quantum annealing for combinatorial decision optimization* addresses settings where the decision space is discrete and the objective can be formulated as a quadratic unconstrained binary optimization (QUBO) problem. For example, selecting an optimal combination of dose levels, fractionation schedule, and concurrent therapies from a discrete set of options can be mapped to a QUBO and solved on a quantum annealer [14]. The annealer's ability to tunnel through energy barriers allows it to escape local optima and find globally better combinations than classical heuristics.

Quantum Approximate Optimization Algorithm (QAOA) provides an alternative circuit-based approach to combinatorial decision problems. In the CDSS context, QAOA can be used to optimize treatment policies by encoding the decision objective and constraints into a problem Hamiltonian and then using a parameterized quantum circuit to prepare a state that encodes high-quality solutions [53]. The variational parameters are optimized classically, and the final state is sampled to obtain candidate decisions. The key innovation across these decision optimization methods is the integration of uncertainty quantification directly into the optimization objective. Rather than optimizing the expected outcome, the system can optimize a risk-sensitive criterion such as:

$$\text{maximize} \quad \mathbb{E}[U(\text{outcome})] - \lambda \cdot \text{Var}[U(\text{outcome})],$$

where U is a utility function that encodes clinical preferences (e.g., the relative value of tumor control versus toxicity), and λ controls risk aversion. The

expectation and variance are estimated using quantum Monte Carlo or QAE, providing a principled way to balance efficacy and safety. For these methods to be clinically useful, they must be embedded in a workflow that respects the realities of clinical practice. The CDSS would typically operate in two modes:

- *Offline planning:* Before treatment begins, the system uses the patient's baseline multi-omics profile to generate a set of candidate treatment plans, each with associated outcome predictions and uncertainty estimates. The clinician reviews these options, possibly adjusting preferences or constraints, and selects a plan for delivery.
- *Online adaptation:* During treatment, as new data become available (daily imaging, fresh biomarker assays, emerging toxicities), the system updates its predictions and uncertainty estimates. If the uncertainty around the current plan becomes too high, or if an alternative plan is predicted to be significantly better, the system can flag the need for adaptation and suggest revised options.

Throughout both modes, the system must communicate not only its recommendations but also the confidence (or lack thereof) behind them. Visualizations of predictive distributions, uncertainty intervals, and trade-off curves can help clinicians make informed decisions that incorporate both the model's outputs and their own clinical judgment. The integration of predictive modeling, uncertainty quantification, and decision optimization into a unified CDSS faces several challenges:

- *Computational complexity:* Running quantum algorithms for prediction, uncertainty estimation, and optimization may exceed the capabilities of current NISQ hardware. Hybrid approaches that offload the most intensive tasks to classical computers while using quantum coprocessors for specific subroutines are the most promising near-term pathway.
- *Validation:* Demonstrating that risk-aware quantum policies improve patient outcomes compared to standard care requires prospective clinical trials, which are challenging to design and execute for rapidly evolving technologies.
- *Interpretability:* Clinicians must understand the basis for recommendations. Developing methods to explain quantum policy decisions—e.g., by highlighting which features drive the uncertainty or which trade-offs are being balanced—is essential for adoption.
- *Calibration:* Uncertainty estimates must be well-calibrated; an 80% confidence interval should contain the true outcome roughly 80% of the time. Achieving good calibration with quantum methods requires careful attention to both algorithmic design and the statistical properties of finite measurement samples.

- *Ethical considerations:* Risk-aware decision-making inherently involves trade-offs between different possible outcomes. The CDSS must be transparent about these trade-offs and allow clinicians to override recommendations based on patient preferences that may be difficult to encode in a utility function.

Despite these challenges, the vision of a CDSS that integrates predictive modeling, uncertainty quantification, and risk-aware decision optimization represents the ultimate goal of quantum-inspired precision radiotherapy. By leveraging quantum methods to capture complex biological interactions, quantify uncertainty, and optimize decisions under that uncertainty, such a system could significantly enhance the ability of clinicians to deliver personalized, effective, and safe radiation treatment.

Table 9: Quantum Methods for Predictive Modeling, Uncertainty Quantification, and Decision Optimization

Component	*Quantum Methods*	*Key Contribution*
Predictive Modeling	QSVM, QNN, DQOR, qscGRN	Capture higher-order multi-omics interactions; ordinal outcomes
Uncertainty Quantification	Quantum Bayesian inference, QAE, QMC	Distinguish aleatoric/epistemic uncertainty; quadratic speedup for Monte Carlo
Decision Optimization	QRL, QDRL, Quantum Annealing, QAOA	Risk-aware policies; superposition-based exploration; tunneling through local optima
Integration	Hybrid offline/online workflow	Continuous updating; uncertainty visualization; clinician-in-the-loop

In summary, the combination of quantum-inspired predictive modeling, rigorous uncertainty quantification, and risk-aware decision optimization forms the intellectual core of a next-generation CDSS for precision radiotherapy. While many of the methods described remain in early stages of development, they point toward a future where the full complexity of multi-omics data can be harnessed to support personalized, evidence-based, and transparent clinical decisions.

3.2.5 Clinical Interpretability, Recommendation Generation, Validation, and Deployment

The ultimate test of any clinical decision support system (CDSS) is not its predictive accuracy on retrospective data, but its ability to influence clinical decisions in real time, under the constraints of actual practice, in a manner that earns and maintains clinician trust. This requires an integrated framework that

encompasses four interconnected capabilities: (i) interpretability—making the reasoning behind recommendations transparent and clinically meaningful; (ii) recommendation generation—translating model outputs into actionable treatment guidance; (iii) validation—ensuring that recommendations are reliable, reproducible, and safe; and (iv) deployment—integrating the CDSS into clinical workflows in both online and offline modes. This subsection synthesizes quantum and quantum-inspired approaches to these challenges. Quantum models, for all their representational power, are notoriously opaque. The high-dimensional Hilbert spaces in which they operate, the complex entanglement structures they learn, and the probabilistic nature of their outputs create a "black box" that can be even more impenetrable than classical deep networks. In clinical contexts, where decisions have life-altering consequences, this opacity is unacceptable. Clinicians must understand not only what a model recommends, but why it recommends it, and under what circumstances the recommendation might be unreliable.

Quantum clinical decision support systems address this imperative by providing risk-aware, explainable recommendations rather than deterministic outputs [14]. By integrating quantum-enhanced predictive modeling with interpretability frameworks, these systems support transparent, patient-specific guidance for adaptive radiotherapy decision-making. The explanations they generate are not post-hoc rationalizations but are grounded in the quantum mechanical principles underlying the model's operation. Classical SHAP (SHapley Additive exPlanations) values provide a unified framework for feature attribution based on cooperative game theory. For a given prediction $f(\mathbf{x})$, the Shapley value ϕ_i quantifies the contribution of feature i by averaging its marginal contribution across all possible feature subsets. Extending this concept to quantum models requires redefining the "contribution" of a feature in terms of quantum mechanical quantities. Quantum-inspired SHAP achieves this by measuring the fidelity contribution of each encoded feature to the final prediction [53]. Given a quantum state $|\psi(\mathbf{x})\rangle$ encoding the full feature vector, and a state $|\psi(\mathbf{x}_{\setminus i})\rangle$ encoding all features except feature i (with feature i replaced by a baseline or marginalized), the fidelity $F_i = |\langle\psi(\mathbf{x})|\psi(\mathbf{x}_{\setminus i})\rangle|^2$ quantifies how much the quantum state changes when feature i is removed. The Shapley value for feature i is then computed by averaging such fidelity differences over all possible subsets of features, weighted appropriately. This approach has several advantages:

- It respects the quantum nature of the representation, measuring feature importance through state overlap rather than through perturbations of classical inputs.
- It naturally accounts for interactions between features, as fidelity captures the joint effect of multiple features on the quantum state.
- It can be estimated efficiently using swap tests or fidelity estimation protocols, avoiding the need for exponential classical computation.

The resulting attribution values provide clinicians with a clear picture of which features—which genes, proteins, radiomic textures—are driving the model's predictions, enabling them to assess whether the model's reasoning aligns with clinical knowledge and biological plausibility. Beyond feature attribution, clinicians often need to answer counterfactual questions: "If this patient had a different genetic profile, would they respond better to this treatment?" or "What would be the expected outcome if we had used a different fractionation schedule?" Classical models struggle with counterfactual reasoning because they learn correlations, not causal structures. Quantum causal models represent treatment-outcome relationships as quantum circuits, enabling principled counterfactual reasoning [14]. In this framework, causal relationships are encoded through controlled operations: a variable X that causally influences Y is represented by a unitary $U_{Y|X}$ that maps $|x\rangle \otimes |0\rangle_Y$ to $|x\rangle \otimes |\psi_{Y|x}\rangle$. The joint quantum state over all variables then factorizes according to the causal graph, with interventions represented by replacing the unitary corresponding to the intervened-upon variable with a fixed state preparation. Counterfactual queries are answered by:

1. Preparing the quantum state corresponding to the observed factual scenario.

2. Applying the intervention that changes the causal structure (e.g., forcing a variable to a different value).

3. Measuring the distribution over the outcome variable of interest.

For example, to answer "What would be the probability of grade ≥ 3 toxicity if this patient had received a lower dose per fraction?", the quantum causal model would intervene on the dose variable, setting it to the counterfactual value while keeping all other variables at their observed values, and then sample the resulting toxicity distribution. This capability is transformative for clinical decision support, as it allows clinicians to explore "what-if" scenarios and understand the causal drivers of outcomes, not just their statistical associations. The quantum formulation ensures that these counterfactual queries respect the constraints of probability theory and causal structure, providing mathematically sound answers that can inform treatment decisions. While Shapley values and causal models provide quantitative insights, clinicians often benefit from visual representations that summarize complex reasoning pathways. Probabilistic decision graphs (PDGs) extend classical decision trees and influence diagrams to the quantum setting, providing graphical explanations of how evidence propagates through the model to reach a recommendation [53]. In a PDG, nodes represent clinical variables or latent factors, and edges represent probabilistic dependencies, quantified by quantum states and operations. The graph is annotated with the quantum circuits that implement these dependencies, allowing clinicians to trace the flow of information from inputs (multi-omics data) to outputs (treatment recommendations). Critical decision points—where the model's uncertainty is high, or where different evidence paths lead to conflicting conclusions—are highlighted, prompting clinician review.

The graph can be interactively explored: clicking on a node reveals the marginal distribution of that variable (as a quantum state or as a classical probability histogram); clicking on an edge shows the conditional relationship; and "what-if" sliders allow clinicians to adjust variables and see how the recommendation changes in real time. This interactivity transforms the CDSS from a static predictor into a collaborative tool for exploring clinical scenarios. No matter how sophisticated the model, clinical decisions ultimately require human judgment. Human-in-the-loop (HITL) quantum frameworks integrate clinical expertise with quantum probabilistic inference, allowing iterative refinement of predictions and recommendations [14]. In these frameworks, the clinician is not merely a recipient of recommendations but an active participant in the reasoning process. A typical HITL interaction proceeds as follows:

1. The CDSS ingests the patient's multi-omics data and generates an initial set of predictions and recommendations, along with uncertainty estimates and explanations (via SHAP, causal models, and decision graphs).

2. The clinician reviews this output, focusing on regions of high uncertainty or on recommendations that conflict with their clinical intuition.

3. The clinician can provide feedback in several forms: (i) adjusting the importance weights of different outcomes (e.g., prioritizing toxicity avoidance over tumor control), (ii) overriding specific features that they believe are unreliable (e.g., due to poor image quality), or (iii) suggesting counterfactual scenarios to explore.

4. The quantum model incorporates this feedback through Bayesian updating: the clinician's input is treated as additional evidence that conditions the quantum state, refining the posterior distribution over outcomes and recommendations.

5. The updated recommendations are presented to the clinician, along with an explanation of how the feedback changed the model's reasoning.

This iterative process enhances both interpretability and reliability. The clinician gains insight into the model's reasoning, and the model benefits from the clinician's expertise, especially in cases where the data are sparse or ambiguous. By combining clinician feedback with uncertainty-aware quantum models, these frameworks support truly patient-specific decision support in precision radiotherapy [14]. Before any CDSS can be deployed clinically, it must undergo rigorous validation to ensure that its recommendations are safe, reliable, and generalizable. For quantum systems, validation must address both the standard challenges of medical AI and issues specific to quantum computing.

Technical validation: The quantum hardware and software components must be tested for reliability. This includes characterizing gate errors, decoherence times, and measurement fidelity; verifying that the circuit implementations match the

intended designs; and establishing that the stochastic outputs are statistically stable across repeated runs.

Predictive validation: The model's predictions must be evaluated on independent test sets, with particular attention to calibration (do predicted probabilities match observed frequencies?), discrimination (can the model distinguish between outcome classes?), and clinical utility (does using the model improve decisions compared to standard care?).

Causal validation: For causal models, the inferred causal structures must be tested against known biology and, where possible, against interventional data (e.g., from clinical trials). Sensitivity analysis can assess how robust the conclusions are to violations of the assumed causal graph.

Human factors validation: The interpretability and usability of the CDSS must be evaluated with clinicians. Does the explanation format support understanding? Do clinicians trust the recommendations? Does the HITL interaction improve decision quality without imposing excessive time burden?

Prospective validation: Ultimately, the CDSS must be tested in prospective studies, ideally randomized controlled trials, to demonstrate that its use leads to better patient outcomes. Given the rapid evolution of quantum hardware, such trials may need to be designed with built-in flexibility to accommodate technological upgrades.

The deployment of a quantum-inspired CDSS in clinical practice occurs in two complementary modes: online and offline, each with distinct requirements and opportunities [15].

Online deployment: During a treatment session, time is critical. The CDSS must integrate imaging, multi-omics, and clinical data streams in real time to support immediate decisions. For example, if daily cone-beam CT reveals anatomical changes that could affect dose distribution, the system might recommend online replanning to maintain target coverage while sparing organs at risk. Quantum-enhanced inference, accelerated by QAE for Monte Carlo dose calculation or by variational circuits for rapid outcome prediction, can provide guidance within the time constraints of a treatment fraction (typically 15–30 minutes). Online deployment requires:

- Ultra-fast data encoding and circuit execution, likely using pre-compiled circuits and optimized measurement protocols.
- Robustness to noise, as there is no time for extensive error mitigation.
- Clear, concise explanations that can be quickly assimilated by the clinician.

- Fail-safe mechanisms: if the quantum hardware is unavailable or the results are too uncertain, the system must gracefully degrade to classical backup models or default protocols.

Offline deployment: Between treatment fractions, there is more time for comprehensive analysis. Longer simulations—such as quantum digital twins that model tumor evolution over the entire treatment course—can update the patient model for subsequent sessions. These simulations might explore multiple alternative scenarios (e.g., different fractionation schedules, different adaptation strategies) and generate a set of candidate plans for clinician review before the next fraction. Offline deployment enables:

- Deeper exploration of the decision space using quantum annealing or QAOA to optimize multi-fraction plans.
- More thorough uncertainty quantification, using quantum Monte Carlo to compute predictive distributions over long-term outcomes.
- Iterative refinement through HITL interactions, with the clinician reviewing and adjusting plans well before the next treatment session.
- Validation checks, comparing model predictions with observed outcomes from previous fractions to detect model drift or unexpected patient responses.

For either deployment mode, the CDSS must integrate seamlessly with existing clinical infrastructure: electronic health records, treatment planning systems, picture archiving and communication systems (PACS), and laboratory information systems. This requires:

- Standardized APIs for data ingestion and result output.
- Secure, HIPAA-compliant handling of patient data.
- Audit trails that record all inputs, outputs, and clinician interactions for medicolegal purposes.
- User interfaces designed in collaboration with clinicians, optimized for clinical workflows rather than for research flexibility.

The vision of an interpretable, validated, and deployable quantum CDSS faces several challenges:

- *Hardware maturity:* Current NISQ devices lack the qubit counts, coherence times, and gate fidelities needed for many of the proposed methods. Progress in error correction and fault tolerance will be essential.
- *Algorithmic scalability:* Many quantum algorithms that work well on small, synthetic problems fail to scale to clinically realistic data sizes. Hybrid approaches that combine quantum and classical processing are the most promising near-term pathway.

- *Interpretability metrics:* Standardized metrics for evaluating the quality of quantum explanations are lacking. Developing such metrics and validating them with clinicians is an urgent research need.
- *Regulatory approval:* Quantum-based medical devices face a novel regulatory landscape. Early engagement with agencies like the FDA will be necessary to establish validation standards.
- *Workflow integration:* Even a perfect model will fail if it disrupts clinical workflows. Co-design with clinicians and iterative prototyping are essential.

Despite these challenges, the integration of interpretability, human-in-the-loop refinement, and dual-mode deployment positions quantum-inspired CDSS as a transformative technology for precision radiotherapy. By making quantum models transparent, engaging clinicians as active participants, and embedding the system in real-world workflows, we can begin to realize the promise of quantum computing for improving cancer care.

Table 10: Components of Quantum CDSS Interpretability, Validation, and Deployment

Component	*Quantum Methods*	*Clinical Function*
Feature Attribution	Quantum-inspired SHAP (fidelity-based)	Identifies which biomarkers drive predictions
Causal Reasoning	Quantum causal models (circuit-based)	Enables counterfactual "what-if" exploration
Visual Explanation	Probabilistic decision graphs	Provides interactive, traceable reasoning pathways
Human-in-the-Loop	Bayesian updating with clinician feedback	Iterative refinement of recommendations
Validation	Technical, predictive, causal, human factors	Ensures safety, reliability, and usability
Online Deployment	Real-time inference, QAE acceleration	Intra-fraction adaptation and guidance
Offline Deployment	Quantum digital twins, annealing, QAOA	Inter-fraction planning and scenario exploration

In summary, the final stage of a quantum-inspired CDSS for multi-omics precision radiotherapy is not merely about generating predictions, but about translating those predictions into trustworthy, actionable guidance that improves clinical decisions. By integrating interpretability tools that reveal the quantum reasoning process, human-in-the-loop frameworks that combine machine and clinician intelligence, and dual-mode deployment strategies that respect the realities of clinical practice, these systems can begin to fulfill the promise of personalized, data-driven radiation oncology.

3.2.6 Quantum Reinforcement Learning for Adaptive Radiotherapy

Adaptive radiotherapy (ART) is inherently a sequential decision-making problem under uncertainty. As a patient progresses through a course of fractionated treatment, their anatomy, tumor characteristics, and normal tissue responses evolve in response to delivered dose and biological processes. Treatment decisions—whether to adjust the dose per fraction, modify beam angles, or replan entirely—must be made incrementally, balancing immediate tumor control against cumulative normal tissue toxicity, while accounting for incomplete and noisy observations. Reinforcement learning (RL) provides a natural framework for optimizing such sequential decisions, and quantum reinforcement learning (QRL) extends this framework by representing states, actions, and policies as quantum objects, offering potential advantages in capturing uncertainty, exploring large action spaces, and compressing high-dimensional patient data.

Adaptive radiotherapy can be formalized as a Markov decision process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where:

- $\mathcal{S}$ is the state space, encompassing all relevant patient information at a given fraction: anatomical configuration (from daily imaging), cumulative dose distribution, biological markers, treatment history, and derived quantities such as tumor control probability (TCP) and normal tissue complication probability (NTCP).
- $\mathcal{A}$ is the action space, consisting of clinically permissible decisions: choice of dose per fraction, beam configuration, MLC aperture shapes, or a binary decision to replan.
- $\mathcal{P}(s_{t+1}|s_t, a_t)$ is the transition probability, capturing the stochastic evolution of the patient state in response to treatment.
- $\mathcal{R}(s_t, a_t, s_{t+1})$ is the reward function, quantifying the immediate clinical benefit—typically a weighted combination of TCP improvement and NTCP penalty.
- $\gamma \in (0, 1]$ is a discount factor that prioritizes near-term outcomes over distant ones, reflecting the time-sensitive nature of cancer treatment.

The goal is to learn a policy $\pi(a|s)$ that maximizes the expected cumulative discounted reward over the treatment course. Classical RL approaches to this problem face two fundamental challenges: (i) the state space is high-dimensional and includes significant uncertainty; (ii) the action space is large and combinatorial. Quantum reinforcement learning addresses these challenges by representing states and policies as quantum states and circuits, leveraging superposition, entanglement, and amplitude amplification. In QRL, the patient state at fraction t is represented as a quantum state ρ_t, a density matrix that encodes both the most likely values of clinical variables and the uncertainties associated with them [13]. This representation naturally captures:

- *Aleatoric uncertainty:* inherent randomness in biological response and imaging noise, reflected in the mixedness of ρ_t.
- *Epistemic uncertainty:* lack of knowledge due to limited observations, which can be reduced as more data are assimilated.
- *Temporal correlations:* the evolution from ρ_t to ρ_{t+1} under action a_t can be modeled as a quantum channel $\mathcal{E}_{a_t}$, a completely positive trace-preserving map that updates the density matrix based on the delivered dose and expected biological effect.

The initial state ρ_0 is constructed from pre-treatment data: imaging, multi-omics profiles, and population-level prior distributions. Techniques such as quantum principal component analysis (qPCA) may be used to compress high-dimensional data into a low-rank density matrix, while quantum feature maps (angle or amplitude encoding) embed specific patient features into the quantum state [58]. Quantum state tomography can be employed periodically to validate that the learned state faithfully represents the actual patient, although full tomography is impractical for large systems; instead, partial tomography or fidelity estimation with clinical observations is used. The action space $\mathcal{A}$ is discretized into a set of clinically feasible dose-per-fraction decisions, bounded by established safety constraints (e.g., maximum dose to organs at risk, minimum dose to target). These constraints are encoded into the quantum RL framework through several mechanisms:

- *Quadratic unconstrained binary optimization (QUBO):* For discrete action selection, the feasibility of an action can be encoded as a penalty term in a QUBO problem, which can be solved on a quantum annealer or via QAOA to identify actions that satisfy all constraints [13].
- *Penalty-augmented variational circuits:* In variational quantum policies, the cost function used for training includes penalty terms for actions that violate constraints, effectively shaping the policy to avoid infeasible actions.
- *Quantum Lagrangian methods:* Inspired by classical constrained optimization, these methods introduce Lagrange multipliers as additional qubits or parameters, enforcing constraints through a primal-dual optimization performed on quantum hardware.

The resulting safe action space ensures that the RL agent only considers actions that are clinically valid, reducing the risk of recommending harmful dose levels. To train QRL agents without endangering patients, a simulated environment—ARTE—is constructed that emulates the patient's response to treatment [13]. ARTE comprises three core components:

1. *Transition function:* A model that predicts the evolution of the patient state given the current state and action. This may be implemented using quantum stochastic process models or open quantum systems frameworks,

where the effect of radiation on tumor and normal tissues is represented as a quantum operation. Parameters of the transition function can be learned from historical data using techniques such as quantum Bayesian inference.

2. *Outcome estimator:* Models that compute TCP and NTCP from the state after each fraction. These estimators can be classical machine learning models trained on large datasets, or quantum models (e.g., QNNs) that capture non-linear dose-response relationships. Importantly, the outcome estimators provide not only point estimates but also uncertainty bounds, which are essential for risk-aware decision-making.

3. *Reward function:* A scalar reward that quantifies the clinical quality of the transition. A typical formulation is:

$$R_t = \alpha \cdot \Delta\text{TCP}_t - \beta \cdot \Delta\text{NTCP}_t - \lambda \cdot \text{Penalty(constraints)},$$

where ΔTCP_t and ΔNTCP_t are changes in predicted control and complication probabilities from fraction $t-1$ to t, and the penalty term enforces adherence to hard constraints. The weights α, β, λ are chosen to reflect clinical priorities.

To address data scarcity and improve the fidelity of outcome prediction, ARTE may incorporate a Wasserstein Generative Adversarial Network with gradient penalty (WGAN-GP) that learns to generate realistic patient trajectories, augmenting the training data and ensuring that the environment captures the full range of possible responses [13]. Two primary QRL approaches have been explored for adaptive radiotherapy: Quantum Reinforcement Learning (QRL) and Quantum Deep Reinforcement Learning (QDRL) [53, 14]. In QRL, both the policy and value functions are represented as parameterized quantum circuits. For a given state ρ_t, the policy circuit $U_\pi(\boldsymbol{\theta})$ produces a quantum state whose measurement yields a probability distribution over actions:

$$\pi(a|\rho_t) = \text{Tr}\left[M_a U_\pi(\boldsymbol{\theta}) \rho_t U_\pi^\dagger(\boldsymbol{\theta})\right],$$

where $\{M_a\}$ are positive operator-valued measures (POVMs) corresponding to each action. The parameters $\boldsymbol{\theta}$ are trained to maximize expected cumulative reward using gradient-based methods (e.g., parameter-shift rule) or evolutionary strategies. In QDRL, a classical deep Q-network (DQN) first processes raw observations (e.g., medical images) to produce a compact feature vector, which is then encoded into a quantum state and fed into a quantum policy network [13]. This hybrid architecture leverages the representational power of deep learning for feature extraction while exploiting quantum circuits for policy representation and exploration. The quantum component can encode multiple candidate actions in superposition, enabling efficient exploration of the action space. Model training and evaluation are conducted across a hierarchy of computational settings: (i) classical DRL baselines, (ii) QDRL on Qiskit-based quantum simulators, and (iii) hardware-executed QDRL on IBM quantum processors. This

allows systematic assessment of performance gains, noise sensitivity, and scalability [13].

At decision time, the current patient state ρ_t is fed into the trained policy network. In the hybrid QDRL architecture, a classical DQN first computes Q-values for each candidate action based on the current state. These Q-values are then embedded into a quantum controller circuit, where they are used to initialize a superposition over actions:

$$|\psi_{\text{actions}}\rangle = \sum_a \sqrt{p_a}|a\rangle,$$

with p_a proportional to the exponentiated Q-values (softmax). Quantum amplitude amplification is then applied to preferentially enhance the probability amplitude of the action with the highest Q-value (or the action that maximizes a risk-adjusted criterion). After a fixed number of amplification steps, the quantum state is measured, collapsing to a single action a_t which is then executed in ARTE (or ultimately, in the clinic) [13]. This quantum action selection mechanism offers two advantages: (i) it naturally implements a form of stochastic policy where actions are sampled according to their estimated value, and (ii) amplitude amplification can boost the probability of selecting near-optimal actions without requiring explicit maximization over a large discrete set. After action a_t is taken, ARTE simulates the transition to a new state ρ_{t+1} using the learned transition model. The outcome estimator computes updated TCP and NTCP values, and the reward R_t is calculated. This reward, along with the new state, is used to update the RL agent. Quantum Deep Reinforcement Learning has demonstrated robust performance in both high-fidelity radiotherapy simulators and early quantum hardware platforms (e.g., IBMQ), providing empirical support for adaptive treatment policies [53]. Comparative analysis between simulator-based and hardware-implemented QDRL reveals the influence of decoherence, gate errors, and circuit constraints on policy learning. These factors impact reward evaluation and convergence dynamics, underscoring the importance of hardware-aware optimization for reliable dose adaptation [53].

Unlike classical RL, where policy updates are typically deterministic (e.g., gradient ascent), QRL performs updates within a quantum probabilistic framework. The learning agent maintains a quantum state that encodes its current knowledge of the value function and policy. After each interaction, this state is updated using quantum Bayesian inference, which combines the prior quantum state with the likelihood of the observed reward and transition [13]. This can be implemented via density matrix evolution:

$$\rho_{\text{new}} = \frac{\mathcal{E}_{\text{obs}} \circ \mathcal{E}_{\text{transition}}(\rho_{\text{old}})}{\text{Tr}[\mathcal{E}_{\text{obs}} \circ \mathcal{E}_{\text{transition}}(\rho_{\text{old}})]},$$

where $\mathcal{E}_{\text{transition}}$ is the quantum channel corresponding to the state transition, and $\mathcal{E}_{\text{obs}}$ is a channel that conditions on the observed reward. Alternatively,

quantum neural gradient descent can be employed, where gradients of the expected reward with respect to policy parameters are estimated using quantum circuits and then used to update the parameters classically. The key distinction from classical methods is that the gradient estimates themselves may be obtained through quantum amplitude estimation, providing quadratic speedups in sample complexity. This iterative learning procedure proceeds across treatment fractions until convergence criteria are met, either by reaching terminal clinical states (e.g., treatment completion or failure) or by satisfying predefined episode length constraints imposed by the clinical protocol [13]. The ultimate output of QRL is a policy that can be deployed to guide dose fractionation decisions in real time. These policies are optimized to jointly maximize TCP and minimize NTCP, a multi-objective optimization problem. Hybrid quantum–deep reinforcement learning frameworks address this by incorporating both objectives into the reward function and using quantum-enhanced state representations to capture the trade-offs [14]. Validation of the learned policies involves several steps:

- *Comparison with clinician prescriptions:* For a held-out set of patient trajectories, the actions recommended by the QRL policy are compared with those taken by human clinicians. Metrics include mean absolute error in dose per fraction, frequency of agreement, and cumulative reward.
- *Longitudinal self-assessment:* The policy's performance is evaluated over complete simulated treatment courses, tracking cumulative TCP, NTCP, and reward. Improvements over baseline policies (e.g., fixed fractionation, classical RL) are quantified.
- *Hardware-in-the-loop testing:* Policies trained on simulators are executed on quantum hardware (e.g., IBMQ) to assess the impact of noise and decoherence. Performance degradation relative to simulators is measured, and error mitigation techniques are evaluated.

Results from early studies suggest that QDRL can modestly enhance decision quality compared to conventional clinical planning and classical reinforcement learning approaches, indicating its potential to better capture uncertainty and complex treatment–response relationships [13]. These improvements, while modest, are encouraging given the current NISQ hardware limitations and point toward greater gains as quantum technology matures. The QRL module for adaptive dose selection is designed to operate within the larger quantum-inspired clinical decision support system described in Section 3.2. It receives as input the patient's current quantum state ρ_t, which has been constructed by the multi-omics fusion and representation learning components (Sections 3.2.2–3.2.3). It outputs a recommended action (dose per fraction, replanning decision) that is then presented to the clinician through the interpretability and recommendation generation interface (Section 3.2.7). The clinician may accept, modify, or override this recommendation, and the resulting decision (whether quantum-recommended or clinician-modified) is fed back to update the patient state

and the QRL agent's knowledge, closing the human-in-the-loop. This tight integration ensures that the quantum RL component benefits from the rich, uncertainty-aware state representations provided by upstream quantum models, while its outputs are contextualized and validated by clinical expertise before execution.

Table 11: Quantum Reinforcement Learning Components for Adaptive Radiotherapy

Component	*Quantum Method*	*Clinical Function*
State Representation	Density matrices, qPCA, quantum feature maps	Encodes patient state with uncertainty
Action Space	QUBO, penalty-augmented circuits, Lagrangian methods	Enforces clinical constraints
Environment (ARTE)	Quantum channels, QNN outcome estimators, WGAN-GP	Simulates patient response
Policy	Variational quantum circuits, hybrid DQN+quantum	Maps states to action probabilities
Action Selection	Amplitude amplification	Enhances probability of optimal actions
Reward Computation	Quantum expectation estimation	Quantifies clinical benefit
Policy Update	Quantum Bayesian inference, density matrix evolution	Uncertainty-aware learning
Validation	Simulator vs. hardware comparison	Assesses real-world performance

In summary, quantum reinforcement learning provides a principled and powerful framework for adaptive radiotherapy, addressing the challenges of sequential decision-making under uncertainty. By representing patient states as quantum density matrices, encoding clinical constraints into the action space, simulating patient response through quantum environments, and learning policies via hybrid quantum–classical algorithms, QRL offers a pathway to truly personalized, dynamically optimized radiation treatment. While still in early stages, the empirical results to date justify continued investment in this approach, with the expectation that advances in quantum hardware will translate into tangible clinical benefits.

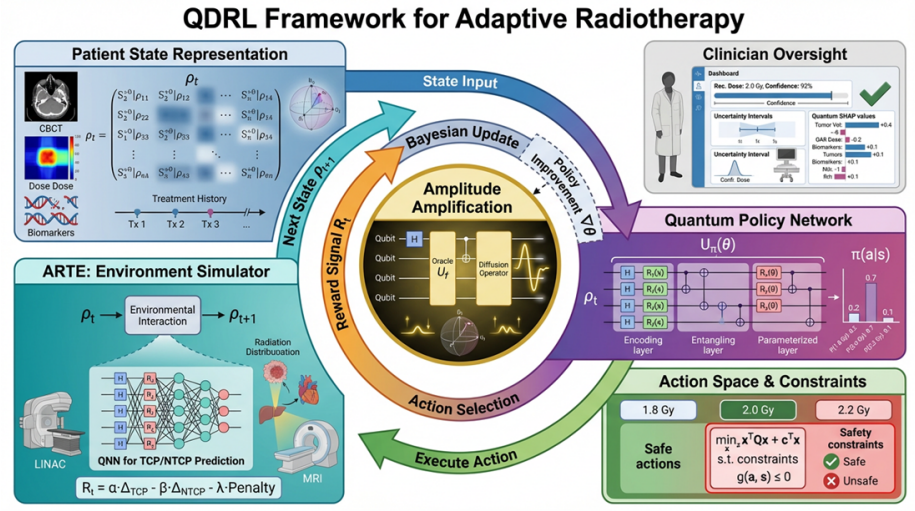

Figure 4: Quantum deep reinforcement learning (QDRL) framework for adaptive radiotherapy. The circular learning loop includes: (1) patient state representation ρ_t (density matrix encoding multi-modal data with uncertainty); (2) quantum policy network $U_\pi(\theta)$ (variational circuit outputting superposition over actions); (3) safe action space (discretised dose-per-fraction decisions with clinical constraints encoded via QUBO); (4) Artificial Radiotherapy Environment (ARTE) comprising transition function (quantum channel), outcome estimator (QNN for TCP/NTCP), and reward function $R_t = \alpha\Delta\text{TCP} - \beta\Delta\text{NTCP} - \lambda\text{Penalty}$; (5) amplitude amplification module (Grover-like) enhancing optimal action probability; (6) human-in-the-loop interface (clinician dashboard with quantum SHAP explanations and feedback controls). Arrows show the learning iteration: state $\rightarrow$ policy $\rightarrow$ action $\rightarrow$ environment $\rightarrow$ reward $\rightarrow$ Bayesian update $\rightarrow$ improved policy.

3.3 Quantum Digital Twin for System-Level Oncology Integration

The concept of a quantum digital twin (QDT) represents the apex of quantum-enabled personalized oncology: a unified, continuously updating virtual replica of the patient that integrates diagnostic models, decision support systems, and adaptive control mechanisms into a single computational framework. Unlike conventional machine learning models that perform isolated tasks—classification, regression, policy learning—a digital twin aims to capture the full complexity of the patient's disease state and its evolution under treatment, enabling real-time simulation, optimization, and intervention guidance throughout the radiotherapy course. The central question, as highlighted in the literature, is: how can diagnostic models, decision systems, and adaptive controllers be integrated into a single, continuously updating clinical system? The QDT paradigm addresses this by constructing a dynamic, patient-specific representation that co-evolves with the actual clinical state, providing a computational substrate for system-level oncology modeling [14, 53].

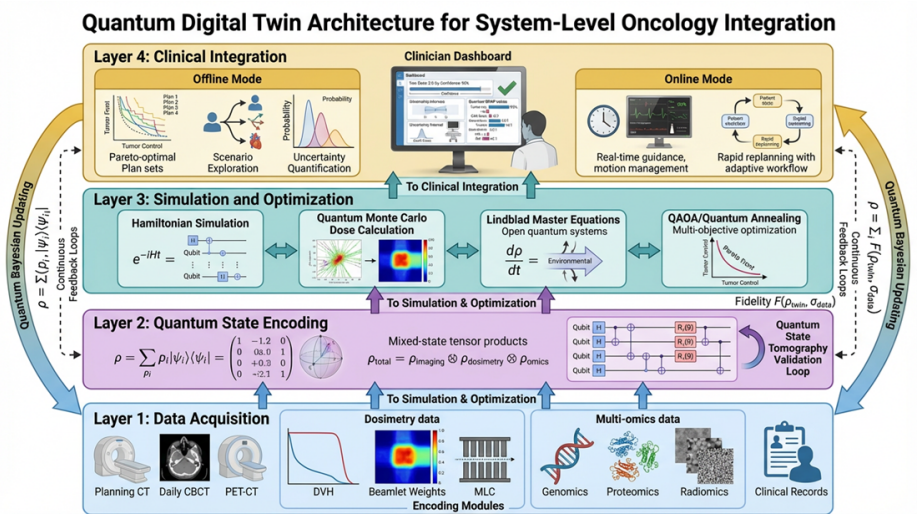

Figure 5: Quantum digital twin architecture. Four layers: (1) Data acquisition (imaging, dosimetry, multi-omics, clinical records); (2) Quantum state encoding (density matrices $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, mixed-state representations $\rho_{\text{total}} = \rho_{\text{imaging}} \otimes \rho_{\text{dosimetry}} \otimes \rho_{\text{omics}}$, controlled encoding circuits for cross-modal correlations); validation via quantum state tomography and fidelity $F(\rho_{\text{twin}}, \sigma_{\text{data}})$; (3) Simulation and optimisation (Hamiltonian simulation e^{-iHt}, quantum Monte Carlo, open quantum systems via Lindblad master equations, QAOA/quantum annealing for multi-objective optimisation); (4) Clinical integration (offline mode for scenario exploration, online mode for real-time guidance). Continuous feedback loops (quantum Bayesian updating) maintain twin alignment with patient state throughout treatment.

3.3.1 Digital Twin Construction

The foundation of any digital twin is a patient-specific virtual replica that models tumor dynamics, anatomical changes, and individual treatment responses [14]. In the quantum context, this replica is constructed by integrating multimodal data—imaging, dosimetry, genomics, proteomics, radiomics—into a unified quantum representation that captures the complex, multiscale interactions inherent in the patient's disease and treatment course. Variational quantum circuits (VQCs) serve as the primary building blocks for encoding the relationships between different data modalities. A VQC can be designed to map input features (e.g., radiomic texture descriptors, gene expression levels, dose-volume metrics) into an entangled quantum state that reflects the correlations among these features. Quantum neural networks (QNNs) extend this by introducing trainable parameters that allow the twin to learn patient-specific patterns from historical data. Hybrid quantum graph models are particularly suited to capturing the network structure of biological systems—for example, gene regulatory networks or tumor-atlas connectivity—by representing nodes as quantum states and edges as entangling operations. The construction process begins with baseline data acquired before treatment: planning CT, diagnostic MRI or PET, biopsy-derived genomic profiles, and population-level priors. These data are encoded into an initial quantum state $|\Psi_0\rangle$ (or density matrix ρ_0) using techniques described in Section 3.4.2. This initial state serves as the seed from which the twin will evolve.

3.3.2 Quantum State Encoding and System Representation

A foundational requirement for any digital twin is the ability to faithfully represent the current state of the patient—including anatomical configuration, dosimetric history, and biological response—in a format that supports both simulation and continuous updating. Quantum machine learning provides a suite of encoding techniques that map high-dimensional clinical data into quantum states, creating a computational substrate for the digital twin [53]. *Density matrix encoding* offers a particularly powerful representation because it naturally captures both the most probable patient state and the uncertainty surrounding that state. In the density matrix formalism, a mixed quantum state

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$$

encodes a probability distribution over possible clinical trajectories, where each pure state $|\psi_i\rangle$ corresponds to a specific anatomical or biological configuration (e.g., a particular tumor shape, a specific gene expression pattern) and the weights p_i reflect the likelihood of that configuration given the available data. This probabilistic representation is essential for adaptive radiotherapy, where incomplete observations and inter-fraction variability introduce genuine uncertainty that must be propagated through subsequent decisions. *Mixed-state quantum representations* extend this idea by allowing the digital twin to simultaneously encode multiple modalities within a single composite quantum state.

For example, imaging data (planning CT, CBCT, MRI) can be encoded in one register, dosimetric data (dose-volume histograms, beamlet weights) in another, and multi-omics data (genomics, proteomics, radiomics) in a third. The joint state is formed via tensor products:

$$\rho_{\text{total}} = \rho_{\text{imaging}} \otimes \rho_{\text{dosimetry}} \otimes \rho_{\text{omics}},$$

and cross-modal correlations are introduced through controlled encoding circuits that entangle the registers. This enables the twin to capture relationships such as how a genetic mutation influences the radiographic appearance of the tumour or how prior dose deposition affects normal tissue complication probabilities. The ability to represent these cross-modal correlations in a unified Hilbert space is a distinct advantage of quantum over classical representations.

Quantum state tomography provides the essential validation mechanism for the digital twin. By performing measurements on the twin's quantum state and comparing the resulting statistics with newly acquired clinical data (e.g., a daily CBCT or a fresh biomarker assay), one can assess the fidelity of the twin and detect when it has diverged from the actual patient. The fidelity between the twin state ρ_{twin} and a state σ_{data} reconstructed from new observations,

$$F(\rho_{\text{twin}}, \sigma_{\text{data}}) = \left(\text{Tr}\sqrt{\sqrt{\rho_{\text{twin}}}\sigma_{\text{data}}\sqrt{\rho_{\text{twin}}}} \right)^2,$$

quantifies their agreement. Tomographic reconstruction techniques, adapted to the constraints of NISQ hardware, allow partial or full recovery of the density matrix, enabling clinicians to gauge the reliability of the twin's predictions before acting upon them. In the context of online adaptation, rapid tomography protocols can flag the need for re-initialisation or Bayesian updating of the twin's state. Together, these encoding and validation techniques establish the quantum digital twin as a dynamic, uncertainty-aware representation that co-evolves with the patient throughout the radiotherapy course. The ability to encode complex, multimodal data into a single quantum state—and to update that state as new information arrives—provides the computational foundation for system-level integration.

3.3.3 Quantum-Enhanced Simulation and Prediction

With the patient's current state represented as a quantum density matrix ρ_t, the digital twin must be able to simulate its evolution under proposed treatments and predict future outcomes. Quantum computing offers several powerful paradigms for this task, each suited to different aspects of the simulation problem. *Hamiltonian simulation* provides a framework for modeling the continuous-time evolution of quantum states under a specified Hamiltonian H. In the context of tumor dynamics, one can construct a Hamiltonian that captures the effects of radiation on cell populations, including cell kill, repair, repopulation,

and reoxygenation. The time evolution operator $U(t) = e^{-iHt}$ is then applied to the current state to obtain the state at a future time:

$$\rho_{t+\Delta t} = U(\Delta t)\rho_t U^\dagger(\Delta t).$$

For biologically relevant Hamiltonians, which may be sparse and structured, Hamiltonian simulation can be performed efficiently on quantum computers, providing exponential speedups over classical methods for large systems [53]. *Quantum Monte Carlo (QMC)* methods leverage quantum amplitude estimation to accelerate the stochastic simulations that are ubiquitous in radiation oncology. Classical Monte Carlo dose calculation, which tracks millions of particle histories through the patient anatomy, converges as $O(1/\sqrt{M})$ with the number of histories M. QMC achieves a quadratic speedup, reaching the same precision with $O(1/M)$ histories, by encoding the probability distribution over particle trajectories into a quantum state and using amplitude estimation to compute expectation values (e.g., dose deposited in a voxel) [52]. In the digital twin context, QMC can be used to rapidly compute dose distributions for candidate treatment plans, enabling real-time comparison of alternatives. *Open quantum systems modeling* provides a framework for representing the dissipative dynamics of tumor response, including effects such as repair, repopulation, and reoxygenation. The evolution of an open quantum system is described by a Lindblad master equation:

$$\frac{d\rho}{dt} = -i[H, \rho] + \sum_k \left(L_k \rho L_k^\dagger - \frac{1}{2}\{L_k^\dagger L_k, \rho\} \right),$$

where the Lindblad operators L_k model the effects of the environment (e.g., oxygen concentration, immune response). By solving this equation on a quantum computer, the digital twin can predict not only the mean behavior of the tumor but also the fluctuations and uncertainties that are critical for risk-aware decision-making. These quantum-enhanced simulation capabilities allow the digital twin to explore "what-if" scenarios rapidly, generating predictions of tumor control probability (TCP), normal tissue complication probability (NTCP), and other clinically relevant endpoints for multiple candidate treatment strategies.

3.3.4 System-Level Optimization and Decision Support

The predictive power of the digital twin is harnessed for decision support through quantum optimization algorithms that search for optimal treatment parameters. Unlike the specialized optimization tasks described in Section 3.1.6 (which focus on individual planning steps), system-level optimization within the digital twin framework considers the entire treatment course as an integrated whole, accounting for interactions between fractions, cumulative dose effects, and the evolving patient state. *Quantum Approximate Optimization Algorithm (QAOA)* can be applied to combinatorial problems that arise in treatment planning, such as beam angle selection or MLC sequencing. For beam angle optimization, the

problem is formulated as selecting a subset of angles from a discrete set to maximize target coverage while sparing organs at risk. This can be mapped to a quadratic unconstrained binary optimization (QUBO) problem and solved using QAOA [12]. The hybrid quantum-classical nature of QAOA allows it to run on near-term hardware, with the quantum circuit exploring the solution space and a classical optimizer refining the parameters. *Quantum annealing* provides an alternative approach to QUBO problems, particularly suited to problems with thousands of variables. In the digital twin context, quantum annealing could be used to optimize fractionation schedules across the entire treatment course, balancing acute and late effects. The ability of quantum annealers to tunnel through energy barriers enables them to escape local optima that trap classical simulated annealing, potentially finding superior global solutions [14].

Quantum-enhanced multi-objective optimization addresses the inherent trade-offs in radiotherapy: maximizing TCP while minimizing NTCP, often with multiple organs at risk each having their own dose constraints. By encoding the Pareto front as a quantum state, one can use amplitude amplification to sample from different regions of the front, presenting the clinician with a diverse set of Pareto-optimal plans. This approach, combined with the uncertainty quantification capabilities of the digital twin, allows for risk-aware decision-making where plans are evaluated not only on their expected outcomes but also on the variance around those expectations. A quantum computing approach to beam angle optimization, for example, formulates the problem as a combinatorial binary decision and addresses it using a hybrid quantum–classical optimization framework. Discrete beam selection is mapped to a quantum-amenable representation, while dose optimization is handled classically. The approach demonstrates improved target conformity and organ-at-risk sparing compared to classical heuristics, illustrating how quantum-enhanced optimization can support system-level treatment planning decisions within digital-twin–enabled radiotherapy workflows [12].

3.3.5 Continuous Feedback, Model Updating and Co-Evolution

A defining characteristic of a digital twin is its ability to continuously update itself as new data become available, ensuring that it remains an accurate representation of the patient throughout the treatment course. Asghar et al. define digital twin construction as the creation of a continuously updated, patient-specific computational model that integrates multimodal clinical, imaging, and molecular data into a unified representation capable of simulating disease dynamics and treatment responses for personalized oncology workflows [59]. In the quantum context, this updating is achieved through several mechanisms. *Quantum Bayesian updating* provides a principled framework for incorporating new observations into the twin's quantum state. Given a prior state ρ_{t-1} and a new observation (e.g., a daily CBCT image, a fresh biomarker measurement), the posterior state ρ_t is obtained by applying a quantum channel that conditions

on the observed data:

$$\rho_t = \frac{\mathcal{E}_{\text{obs}}(\rho_{t-1})}{\text{Tr}[\mathcal{E}_{\text{obs}}(\rho_{t-1})]},$$

where $\mathcal{E}_{\text{obs}}$ is a completely positive trace-non-increasing map corresponding to the likelihood of the observation. This approach naturally handles noisy or incomplete observations and maintains the uncertainty representation throughout the treatment course. *Density matrix evolution* under the influence of treatment and time advances the twin in lockstep with the patient. Between observations, the twin evolves according to the simulated dynamics described in Section 3.4.3, using Hamiltonian simulation or Lindblad master equations. This predicted evolution provides a baseline against which new observations can be compared, enabling early detection of deviations from expected response. *Quantum neural gradient descent* provides a mechanism for fine-tuning the twin's parameters to maintain alignment with observed outcomes. If systematic discrepancies emerge between the twin's predictions and actual patient data, the parameters of the underlying models (e.g., radiobiological parameters in the Hamiltonian, weights in the encoding circuits) can be adjusted using gradient-based optimization. The gradients themselves may be estimated using quantum techniques such as the parameter-shift rule, ensuring that the twin remains a faithful representation even as the patient's condition evolves in unexpected ways. This continuous feedback loop ensures that the digital twin co-evolves with the patient, becoming more accurate and personalized as treatment progresses.

3.3.6 Clinical Validation and Performance Assessment

Before a quantum digital twin can be deployed clinically, it must undergo rigorous validation to ensure that its predictions are reliable and that its recommendations are safe. Hybrid quantum–classical workflows and quantum simulators serve as intermediate platforms for implementing patient-specific modeling and adaptive scenario testing, enabling preliminary evaluation of treatment strategies and paving the way toward fully quantum-powered oncology ecosystems [53]. *Quantum kernel-based validation* uses fidelity measures to compare the twin's predicted state distribution with observed outcomes. For a cohort of patients, one can compute the fidelity between the twin's state $\rho^{(i)}_{\text{twin}}$ for patient i and a state $\sigma^{(i)}_{\text{outcome}}$ constructed from their actual clinical outcomes. High average fidelity indicates that the twin accurately captures the patient population; low fidelity for specific patients flags cases where the twin may need refinement. *Quantum–classical Monte Carlo comparison* benchmarks the twin's predictions against gold-standard classical simulations. For dose calculation, for example, the twin's QMC-based dose estimates can be compared with classical Monte Carlo results for the same anatomy and beam configuration. Agreement within clinically acceptable tolerances validates the quantum approach; discrepancies may indicate errors in the Hamiltonian or encoding.

Fidelity and distance measures on quantum states provide a rich set of tools

for assessing twin accuracy. Beyond simple fidelity, one can compute trace distance, quantum relative entropy, or Wasserstein distance between predicted and observed distributions, each offering a different perspective on twin performance. These measures can be tracked over time to detect drift or degradation in twin fidelity. Hybrid quantum recurrent models approximate temporal tumor evolution and inter-fraction setup variability as stationary quantum processes. This approach enables continuous, uncertainty-aware prediction of patient-specific responses, supporting high-fidelity simulation and adaptive decision-making in digital twin frameworks [14].

3.3.7 Integration with the Clinical Workflow

The quantum digital twin is not a standalone system but an integral component of the broader clinical workflow, interacting with data acquisition systems, treatment planning software, and the clinical team. Its deployment mirrors the two modes described in Section 3.2.8. The first is the *offline mode*, whereby between treatment fractions, the twin operates in a comprehensive simulation mode, exploring multiple "what-if" scenarios for future fractions. It may run hundreds or thousands of simulated trajectories, each corresponding to a different treatment policy, and present the clinician with a summary of expected outcomes and their uncertainties. This offline analysis informs the initial plan for the next fraction and can be used to update population-level models. The second is the *online mode* whereby during a treatment session, the twin operates in a rapid-response mode, providing real-time guidance for adaptation. If daily imaging reveals unexpected anatomical changes, the twin can quickly simulate the dosimetric consequences of proceeding with the original plan versus adapting, and recommend a course of action within the time constraints of the treatment slot. Quantum speedups in simulation and optimization are critical here, as classical methods would be too slow. Throughout both modes, the twin communicates with the clinician through the interpretability and visualization interfaces described in Section 3.2.7, ensuring that its recommendations are transparent and actionable. The vision of a fully realized quantum digital twin for radiation oncology faces several challenges:

- *Hardware requirements:* The qubit counts, coherence times, and gate fidelities needed for full-scale patient simulation are far beyond current NISQ capabilities. Progress in error correction and fault tolerance will be essential.
- *Model complexity:* Capturing the full complexity of tumor biology, including spatial heterogeneity, immune interactions, and microenvironmental effects, requires models that may be too complex even for quantum computers. Hierarchical modeling, where detailed submodels are used for critical processes and coarse-grained models elsewhere, may be necessary.
- *Validation burden:* Validating a digital twin requires extensive clinical data, ideally from prospective trials. The rapid evolution of quantum

hardware complicates this validation, as a twin validated on one generation of hardware may not perform identically on the next.

- *Integration challenges:* Integrating a quantum digital twin into existing clinical information systems requires standardized data formats, secure APIs, and workflows that respect clinical time constraints. Co-design with clinicians and IT professionals is essential.

Despite these challenges, the quantum digital twin represents the ultimate expression of personalized, predictive, and adaptive radiotherapy. By unifying diagnostic models, simulation engines, and decision support systems into a single, continuously updating representation of the patient, it offers the promise of truly individualized cancer care—care that anticipates rather than reacts, optimizes rather than satisfies, and learns rather than static.

Table 12: Components of a Quantum Digital Twin for Radiation Oncology

Component	*Quantum Methods*	*Clinical Function*
Twin Construction	VQCs, QNNs, hybrid graph models	Integrates multimodal data into unified representation
State Encoding	Density matrices, mixed states, tomography	Captures patient state with uncertainty
Simulation	Hamiltonian simulation, QMC, open systems	Predicts tumor evolution and treatment response
Optimization	QAOA, quantum annealing, multi-objective	Identifies optimal treatment parameters
Updating	Bayesian inference, density matrix evolution	Maintains alignment with actual patient
Validation	Kernel methods, fidelity measures, MC comparison	Ensures reliability and safety
Integration	Offline/online deployment modes	Supports both planning and real-time adaptation

In summary, the quantum digital twin paradigm offers a comprehensive framework for system-level oncology integration, bringing together the full range of quantum methods—encoding, simulation, optimization, and updating—into a unified, patient-specific model. While significant challenges remain, the potential rewards—truly personalized, dynamically adaptive, optimally delivered radiotherapy—justify continued investment in this vision.

4 Challenges and Clinical Barriers

Despite the theoretical promise of quantum machine learning (QML) in radiation oncology, the path to clinical deployment is obstructed by a series of formidable challenges. These barriers span hardware limitations, algorithmic

instability, data representation difficulties, and the stringent requirements of medical validation and regulation. A clear understanding of these obstacles is essential to temper expectations and to guide research toward realistically achievable goals.

4.1 Hardware Constraints in the NISQ Era

Current quantum processors operate in the noisy intermediate-scale quantum (NISQ) regime [8], characterized by limited qubit counts (typically 50–1000), short coherence times, and significant gate errors. These imperfections restrict the depth of circuits that can be executed reliably, thereby limiting the complexity of implementable QML models. For radiation oncology applications, which often involve high-dimensional data (e.g., 512×512 medical images) or large combinatorial spaces (e.g., beam angle selection), the available qubit resources and error rates pose a severe bottleneck [15]. Moreover, the need for repeated measurements to estimate expectation values with sufficient precision further erodes any potential quantum advantage, especially when compared to increasingly powerful classical hardware (GPUs, TPUs) that continue to scale.

4.2 Algorithmic Limitations and Trainability Issues

Variational quantum algorithms, the workhorses of near-term QML, are prone to the *barren plateau* phenomenon: for sufficiently expressive circuits, the gradient of the cost function vanishes exponentially with the number of qubits, rendering training impractical [26]. This effect is exacerbated by noise and by the use of global cost functions. In the context of oncology, where problem sizes are large and loss landscapes are complex, barren plateaus may prevent any meaningful learning unless the circuit architecture is carefully constrained or pretrained. Additionally, finite sampling noise (shot noise) introduces stochasticity into gradient estimates, requiring many circuit repetitions and reducing the effective speedup [60]. The interplay of noise, vanishing gradients, and sampling overhead means that many QML proposals that look good on paper fail to deliver tangible improvements in practice [52].

4.3 Data Encoding and Input/Output Bottlenecks

A fundamental challenge for QML in any medical application is the efficient loading of classical data into quantum states. Amplitude encoding, which can represent 2^n classical values using n qubits, is exponentially efficient in theory but requires coherent arithmetic operations that are themselves difficult to implement on NISQ hardware. Angle encoding is simpler but uses qubits inefficiently. Moreover, the final step of extracting classical information through measurement collapses the quantum state and yields only probabilistic outputs, necessitating repeated runs. For high-resolution medical images or high-dimensional radiomic feature vectors, these encoding and readout costs can

dominate the total runtime, potentially offsetting any theoretical quantum advantage [52]. Furthermore, most clinical datasets are small (hundreds to thousands of patients), which may not justify the overhead of quantum encoding when classical models already perform well.

4.4 Lack of Standardized Clinical Quantum Datasets

Machine learning, whether classical or quantum, thrives on large, high-quality, and well-annotated datasets. In radiation oncology, publicly available repositories (e.g., The Cancer Imaging Archive) have accelerated classical AI research. However, no equivalent datasets exist that are specifically formatted or preprocessed for QML experiments. Researchers often resort to downsampling images or reducing feature dimensions to fit quantum hardware, which may discard clinically relevant information and limit the validity of conclusions. The absence of standardized benchmarks makes it difficult to compare different QML approaches or to assess their true clinical potential [15].

4.5 Regulatory and Interpretability Hurdles

Any computational tool intended to influence patient care must undergo rigorous regulatory approval (e.g., FDA clearance in the United States). Quantum-based models introduce novel challenges: their stochastic nature, hardware-dependent behavior, and lack of established interpretability methods make it difficult to demonstrate safety, efficacy, and reproducibility. Regulators expect clear documentation of algorithm behavior, validation on independent cohorts, and robustness to variations in input data. Current QML models, especially those implemented on real quantum hardware, exhibit variability across runs and even across different days on the same machine, complicating the validation process [52]. Moreover, clinicians are rightfully hesitant to trust "black-box" recommendations, and quantum models are often even more opaque than classical deep networks. The development of quantum-specific explainability tools (e.g., quantum SHAP, fidelity-based attribution) is still in its infancy and must mature before clinical adoption becomes feasible.

4.6 The Gap Between Theoretical Speedups and Practical Gains

Many quantum algorithms offer provable asymptotic speedups only under idealized conditions—sparse matrices, well-conditioned systems, or access to fault-tolerant hardware. In realistic oncology problems, these conditions rarely hold. For instance, the HHL algorithm requires a well-conditioned matrix and efficient block-encoding, which is non-trivial for the dense, noisy matrices arising in deformable image registration. QAOA lacks rigorous performance guarantees for problem sizes of clinical interest. Even Grover's quadratic speedup is eroded if the oracle call requires a significant overhead. Consequently, translating theoretical advantages into measurable clinical benefits requires careful end-to-end

benchmarking against state-of-the-art classical solvers on relevant problem instances—a step that is still largely missing in the literature [52].

4.7 Toward Realistic Evaluation and Incremental Progress

Overcoming these barriers will demand a concerted, interdisciplinary effort. Near-term research should focus on:

- Identifying small, well-defined subproblems where quantum circuits can be rigorously compared to classical baselines (e.g., optimizing a few beam angles, classifying a reduced set of radiomic features).
- Developing hardware-efficient ansätze that mitigate barren plateaus and exploit problem symmetries.
- Creating publicly available quantum-friendly oncology benchmarks that preserve clinical realism while being tractable on current simulators and hardware.
- Integrating QML components into existing clinical workflows in a hybrid fashion, so that failure of the quantum part does not compromise the entire pipeline.
- Engaging with regulatory bodies early to understand requirements for validation and to shape the development of interpretability methods.

Only through such pragmatic, stepwise progress can the field move from theoretical promise to clinically impactful reality. The challenges are substantial, but they are not insurmountable; they define the research agenda for the coming decade.

5 Discussion

The convergence of quantum computing and radiation oncology presents both unprecedented opportunities and substantial challenges. This review has systematically mapped quantum algorithms to the clinical workflow, revealing several areas where quantum methods could, in principle, outperform classical approaches: combinatorial optimisation (beam angle selection, MLC sequencing), linear algebra (deformable registration, inverse planning), Monte Carlo simulation (dose calculation, uncertainty quantification), and high-dimensional pattern recognition (radiomics, multi-omics integration). Despite these theoretical advantages, the path to clinical deployment is fraught with obstacles. The barren plateau phenomenon renders many variational circuits untrainable as qubit counts increase, necessitating careful architecture design and problem-specific initialisations. Finite sampling noise forces repeated measurements that erode quantum speedups, especially when high precision is required. Quantum decoherence and gate errors restrict circuit depth on current hardware, limiting

the complexity of implementable algorithms. Moreover, the input/output bottleneck—encoding classical medical images into quantum states and extracting classical results—can dominate runtime, offsetting any computational gain. These limitations underscore the importance of hybrid quantum-classical approaches that delegate only the most quantum-amenable subtasks to the quantum processor while keeping the majority of the workflow on classical hardware. The NISQ era will likely see the first clinical pilots using such hybrids, focusing on well-defined, low-scale problems where a clear advantage can be demonstrated (e.g., small-scale optimisation for stereotactic radiosurgery or feature selection from a moderate number of radiomic variables).

Another critical consideration is data availability and quality. Quantum models are notoriously data-hungry in terms of the number of circuit evaluations required, yet oncology datasets are often small and heterogeneous. Transfer learning, where classical networks pretrained on large natural image corpora are fine-tuned with quantum layers, offers a pragmatic path forward. Similarly, quantum generative models can augment limited data, but their fidelity must be rigorously validated against clinical ground truth. Interpretability remains a cornerstone of clinical acceptance. While quantum circuits can be more expressive, they are also more opaque than classical models. Emerging work on quantum-aware explainability (e.g., quantum SHAP, fidelity-based attribution) must be integrated into any QML tool intended for patient care. Clinicians will not act on recommendations they cannot understand, especially when they involve complex trade-offs between tumour control and normal tissue toxicity. From a regulatory perspective, quantum-inferred decisions will face the same scrutiny as any AI-based medical device. The "black-box" nature of many quantum models, combined with the sensitivity of healthcare data, demands that any clinical QML system undergo rigorous prospective validation, preferably in multi-centre trials. Moreover, the rapid evolution of quantum hardware poses a challenge for regulatory stability: a model validated on one generation of quantum processors may not perform identically on the next. Looking forward, we envision a roadmap with three horizons:

1. *Near-term (1–5 years):* Hybrid quantum-classical proof-of-concept studies on well-circumscribed problems (beam angle selection, small-scale radiomics) using NISQ devices and simulators. Development of open-source benchmarks and clinical datasets tailored for QML.

2. *Medium-term (5–10 years):* Integration of quantum coprocessors into treatment planning systems for specific subroutines (for instance, Monte Carlo acceleration via QAE). Prospective observational studies comparing quantum-enhanced predictions with classical standards.

3. *Long-term (¿10 years):* Fault-tolerant quantum computers enabling full-scale simulation of patient-specific radiation transport, real-time adaptive replanning, and system-level digital twins that integrate multi-omics, imaging, and dosimetry into a unified predictive model.

This roadmap must be pursued in close collaboration between quantum physicists, computer scientists, medical physicists, and radiation oncologists. Only through such interdisciplinary efforts can we ensure that quantum innovation translates into tangible improvements in patient outcomes.

6 Conclusions

Quantum machine learning offers a compelling new frontier for radiation oncology, with the potential to overcome computational bottlenecks that limit current treatment planning, image analysis, and outcome prediction. This review has provided a comprehensive taxonomy of quantum algorithms relevant to the field, ranging from fault-tolerant methods (Grover, HHL, QAE) to NISQ-era techniques (QAOA, QCNN, QSVM) and emerging hybrid architectures (QDRL, quantum digital twins, VQE, QGANs, quanvolutional NNs, QGNNs, and tensor networks). We have mapped these algorithms to specific stages of the radiation oncology workflow—consultation, simulation, planning, delivery, adaptation, and follow-up—highlighting opportunities for optimisation, simulation, and pattern recognition. At the same time, we have critically examined the formidable challenges that lie ahead: hardware constraints (limited qubits, decoherence, gate errors), algorithmic barriers (barren plateaus, finite sampling noise), data encoding bottlenecks, the lack of standardised clinical quantum datasets, and the stringent requirements of clinical validation and interpretability. The path to clinical adoption will be incremental, with near-term successes likely confined to well-defined subproblems where quantum advantages can be unequivocally demonstrated (e.g., small-scale beam angle optimisation, feature selection from moderate-sized radiomic sets). Hybrid quantum-classical frameworks, coupled with transfer learning and explainability tools, represent the most viable strategy for the coming decade. Ultimately, the integration of quantum computing into radiation oncology is not merely a technological upgrade—it is a paradigm shift that could redefine what is computationally possible in personalised cancer care. Realising this vision will require sustained investment, rigorous evaluation, and a community committed to bridging the gap between quantum theory and clinical practice. We hope this review serves as both a foundation and a catalyst for that journey.

Acknowledgments. We acknowledge that discussions with Dr. Bhaskar Vishwanathan and Dr. Deepak Thaper at Amrita Hospital, Faridabad have helped in the writing of this paper.

References

[1] Edward C. Halperin, Carlos A. Perez, and Luther W. Brady. *Perez and Brady's Principles and Practice of Radiation Oncology*. Wolters Kluwer, Philadelphia, 8 edition, 2023.

[2] Faiz M. Khan and John P. Gibbons. *Khan's The Physics of Radiation Therapy*. Wolters Kluwer, Philadelphia, 6 edition, 2021.

[3] Indrin J. Chetty, Bruce Curran, Johanna E. Cygler, John J. DeMarco, Gary Ezzell, Bruce A. Faddegon, Iwan Kawrakow, Paul J. Keall, Helen Liu, C.-M. Charlie Ma, D. W. O. Rogers, Jan Seuntjens, Daryoush Sheikh-Bagheri, and Jeffrey V. Siebers. Report of the aapm task group no. 105: Issues associated with clinical implementation of monte carlo-based photon and electron external beam treatment planning. *Medical Physics*, 34(12):4818–4853, 2007.

[4] Martin J. Murphy, Kristy K. Brock, Marc L. Kessler, and Todd R. McNutt. Image registration in radiation oncology: A review of methods and applications. *Medical Physics*, 47(5):e123–e136, 2020.

[5] Maximilian Brettner, Peter Ziegenhein, and Uwe Oelfke. Multi-criteria optimization in radiation therapy: A review of methods and applications. *Physics in Medicine & Biology*, 66(4):04TR01, 2021.

[6] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[7] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 10th anniversary edition edition, 2010.

[8] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[9] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv*, 2014.

[10] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019.

[11] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, 2017.

[12] Nimita Shinde, Ya-Nan Zhu, Haozheng Shen, and Hao Gao. A quantum computing approach to beam angle optimization. *arXiv preprint*, 2025.

[13] Dipesh Niraula, Jamalina Jamaluddin, Martha M. Matuszak, Randall K. Ten Haken, and Issam El Naqa. Quantum deep reinforcement learning for clinical decision support in oncology: application to adaptive radiotherapy. *Scientific Reports*, 11(1):23545, 2021.

[14] Dipesh Niraula, Sunan Cui, Julia Pakela, Lise Wei, Yi Luo, Randall K Ten Haken, and Issam El Naqa. Current status and future developments in predicting outcomes in radiation oncology. *British Journal of Radiology*, 95(1139):20220239, 2022.

[15] Nasna Nassir, Mohammad Amiruddin Hashmi, Kavya Gopan Raji, Bassam Jamalalail, Andrew Maksymowsky, Stephen W. Scherer, Alawi Alsheikh-Ali, and Mohammed Uddin. Quantum computing and the implementation of precision medicine. *npj Genomic Medicine*, 10(1), 2025.

[16] David L. Craft, Theodore S. Hong, Helen A. Shih, and Thomas R. Bortfeld. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 81(5):e83–e91, 2011.

[17] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.

[18] Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81(2):865–942, 2009.

[19] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv*, 2018.

[20] Iris Cong, Soonwon Choi, and Mikhail D. Lukin. Quantum convolutional neural networks. *Nature Physics*, 15:1273–1278, 2019.

[21] Philippe Lambin, Emmanuel Rios-Velazquez, Ralph Leijenaar, Sara Carvalho, Ruud G.P.M. van Stiphout, Patrick Granton, Catharina M.L. Zegers, Robert Gillies, Ronald Boellard, André Dekker, and Hugo J.W.L. Aerts. Radiomics: Extracting more information from medical images using advanced feature analysis. *European Journal of Cancer*, 48(4):441–446, 2012.

[22] Hugo J. W. L. Aerts, Emmanuel Rios Velazquez, Ralph T. H. Leijenaar, Chintan Parmar, Patrick Grossmann, Sara Carvalho, Johan Bussink, René Monshouwer, Benjamin Haibe-Kains, Derek Rietveld, Frank Hoebers, Michelle M. Rietbergen, C. René Leemans, André Dekker, John Quackenbush, Robert J. Gillies, and Philippe Lambin. Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature Communications*, 5:4006, 2014.

[23] Issam El Naqa. Detection and prediction of radiotherapy errors. In Issam El Naqa, Ruijiang Li, and Michael Murphy, editors, *Machine Learning in Radiation Oncology*, pages 155–176. Springer, Cham, 2015.

[24] Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, 2015.

[25] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.

[26] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9:4812, 2018.

[27] Marco Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature communications*, 12(1):1791, 2021.

[28] Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. In *SIGACT News*, volume 48, pages 41–76, 2017.

[29] Seymour H. Levitt, James A. Purdy, Carlos A. Perez, and Srinivasan Vijayakumar. *Technical Basis of Radiation Therapy: Practical Clinical Applications*. Springer, Berlin, 5 edition, 2012.

[30] Leonard L. Gunderson and Joel E. Tepper. *Clinical Radiation Oncology*. Elsevier, Philadelphia, 5 edition, 2020.

[31] Carsten Nieder, Nicolaus H. Andratschke, and Anca Grosu. Re-irradiation: New frontiers. *Radiation Oncology*, 16(1):216, 2021.

[32] Charles M. Washington and Dennis T. Leaver. *Principles and Practice of Radiation Therapy*. Elsevier, St. Louis, 4 edition, 2016.

[33] Juliette Thariat, Sophie Marae, and Jérôme Doyen. Immobilization in radiation therapy: A review. *Cancer/Radiothérapie*, 25(6):582–588, 2021.

[34] Daniela Thorwarth. Functional imaging for radiotherapy treatment planning: Current status and future directions—a review. *The British Journal of Radiology*, 93(1107):20190020, 2020.

[35] Paul J. Keall, Sastry S. Vedam, and Rohini George. A review of real-time 4d image-guided radiotherapy for lung cancer. *Medical Physics*, 49(3):1456–1472, 2022.

[36] Benedick A. Fraass. The evolution of treatment planning systems. *Medical Physics*, 48(10):6541–6554, 2021.

[37] ICRU. Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (imrt). ICRU Report 91, International Commission on Radiation Units and Measurements, 2016. Journal of the ICRU, Vol. 16, No. 1.

[38] Lawrence B. Marks, Ellen D. Yorke, Andrew Jackson, Randall K. Ten Haken, Louis S. Constine, Avraham Eisbruch, Søren M. Bentzen, Jiho Nam, and Joseph O. Deasy. Use of normal tissue complication probability models in the clinic. *International Journal of Radiation Oncology*Biology*Physics*, 76(3):S10–S19, 2010.

[39] Carlos E. Cardenas, Jinzhong Yang, Brian M. Anderson, Laurence E. Court, and Kristy K. Brock. Auto-segmentation in radiation oncology: A review. *Medical Physics*, 47(5):e178–e191, 2020.

[40] Steve Webb. *Intensity-Modulated Radiation Therapy*. Institute of Physics Publishing, Bristol, 2003.

[41] Chen-Shou Chui and Thomas LoSasso. *Inverse planning in radiation therapy*. Medical Physics Publishing, Madison, 2001.

[42] Stanley H. Benedict, Kamil M. Yenice, David Followill, James M. Galvin, William Hinson, Brian Kavanagh, Paul Keall, Michael Lovelock, Sanford Meeks, Lech Papiez, Thomas Purdie, Ramani Sadagopan, Michael C. Schell, Bill Salter, David J. Schlesinger, Almon Shiu, Timothy Solberg, Danny Y. Song, Valery Stieber, Robert Timmerman, Wolfgang A. Tomé, Dirk Verellen, Lu Wang, and Fang-Fang Yin. Stereotactic body radiation therapy: The report of aapm task group 101. *Medical Physics*, 37(8):4078–4101, 2010.

[43] Sara Broggi, Gian Mauro Cattaneo, Claudio Fiorino, Filippo Alongi, and Riccardo Calandrino. Deformable image registration for adaptive radiotherapy: A review. *Physica Medica*, 78:56–67, 2020.

[44] Eric E. Klein, Joseph Hanley, John Bayouth, Fang-Fang Yin, William Simon, Sean Dresser, Christopher Serago, Francisco Aguirre, Lijun Ma, Bijan Arjomandy, Chihray Liu, Carlos Sandin, and Sean Holmes. Quality assurance of medical accelerators. *Medical Physics*, 36(9):4197–4212, 2009. AAPM Task Group 142 report.

[45] Daniel A. Low, Jean M. Moran, James F. Dempsey, Lei Dong, and Mark Oldham. Quality assurance of intensity-modulated radiation therapy. *Medical Physics*, 37(5):2226–2250, 2010. AAPM Task Group 119 report.

[46] David A. Jaffray. *Image-guided radiation therapy: From concept to practice*. Lippincott Williams & Wilkins, Philadelphia, 2 edition, 2012.

[47] S. van der Merwe, F. Van den Heuvel, D. Verellen, W. Wells, and P. Lambin. Image guidance in radiation therapy: A review of the literature on clinical implementation. *Clinical Oncology*, 29(4):226–235, 2017.

[48] Jan-Jakob Sonke, Marianne Aznar, Coen Rasch, Jose Belderbos, Marc Kessler, Philippe Lambin, Dirk Verellen, and Dirk De Ruysscher. Adaptive radiotherapy: A review of the literature and recommendations from the estro-acrop task group. *Radiotherapy and Oncology*, 141:4–13, 2019.

[49] Olga L. Green, Lauren E. Henke, Geoffrey D. Hugo, Jeffrey R. Olsen, Maria A. Thomas, Sasa Mutic, and Rojano Kashani. Magnetic resonance image-guided radiotherapy: A review of the technology and clinical applications. *Journal of Medical Imaging and Radiation Oncology*, 65(5):519–529, 2021.

[50] Dennis Winkel, Gijsbert H. Bol, Petra S. Kroon, Bram van Asselen, Sara S. Hackett, Anita M. Werensteijn-Honingh, Martijn P. W. Intven, Wietse S. C. Eppinga, Rob H. N. Tijssen, Linda G. W. Kerkmeyer, Hans C. J. de Boer, Helena M. Verkooijen, and Bas W. Raaymakers. Adaptive radiotherapy: The elekta unity mr-linac concept. *Clinical and Translational Radiation Oncology*, 18:54–59, 2019.

[51] Jia Wu, Khin Khin Tha, Lei Xing, and Ruijiang Li. Radiomics and radiogenomics for precision radiotherapy. *Journal of Radiation Research*, 59(suppl_1):i25–i31, 2018.

[52] Siddhi Ramesh, Teague Tomesh, Samantha J. Riesenfeld, Frederic T. Chong, and Alexander T. Pearson. Quantum computing for oncology. *Nature Cancer*, 5:811–816, 2024.

[53] Milad Rahimi and Farkhondeh Asadi. Oncological applications of quantum machine learning. *Technology in Cancer Research & Treatment*, 22:15330338231215214, 2023.

[54] Nadine Matondo-Mvula and Khaled Elleithy. Breast cancer detection with quanvolutional neural networks. *Entropy*, 26(8):630, 2024.

[55] Xianzhi Huang, Fangyi Xu, Wenchao Zhu, Lin Yao, Jiahuan He, Junhao Su, Wending Zhao, and Hongjie Hu. An integrated strategy based on radiomics and quantum machine learning: diagnosis and clinical interpretation of pulmonary ground-glass nodules. *BMC Medical Imaging*, 25(1):279, 2025.

[56] Tony Felefly. *Quantum-Classical Machine Learning for Brain Tumor Imaging Analysis*. PhD thesis, Université de Strasbourg; Université Saint-Joseph (Beyrouth), Strasbourg, France, 2024. PhD Thesis, defended December 13, 2024.

[57] M Priyadharshini, B Deevena Raju, A Faritha Banu, P Jagdish Kumar, V Murugesh, and Oleg Rybin. A quantum machine learning framework for predicting drug sensitivity in multiple myeloma using proteomic data. *Scientific Reports*, 15(1):26553, 2025.

[58] Chih-Wei Chang, Sri Sai Akkineni, Mingzhe Hu, Keyur Shah, Yuan Gao, Pretesh Patel, Ashesh B Jani, Greeshma Agasthya, Jun Zhou, and Xiaofeng Yang. A digital twin framework for adaptive treatment planning in radiotherapy. *arXiv preprint arXiv:2506.14701*, 2025.

[59] Uzma Saddia Asghar and Caroline Chung. Application of digital twins for personalized oncology. *Nature Reviews Cancer*, 25:823–825, 2025.

[60] Theresa Sharu Jose and Osvaldo Simeone. Error mitigation-aided optimization of parameterized quantum circuits: Convergence analysis. *arXiv preprint arXiv:2209.11514*, 2022.