Center for Management Science Research
ISSN: 3006-5291 3006-5283 Volume 4, Issue 3, 2026

# AI-ENHANCED COMPUTATIONAL FRAMEWORKS FOR COMPLEX PHYSICAL SYSTEMS: 3D MICROSTRUCTURE RECONSTRUCTION, QUANTUM NON-UNITARY DYNAMICS, AND REINFORCEMENT LEARNING BASED RADIOTHERAPY OPTIMIZATION

**Ammad Ul Islam[*1], Sidra Kanwal[2], Malik Zohaib Hussain[3], Fazalerabbi[4], Awais Amin[5], Fatima Muzammil Shah[6]**

[*1]*Artificial Intelligence and Computer Vision (iVision) Lab, Department of Electrical Engineering, Institute of Space Technology, Islamabad 44000, Pakistan*[1]
[2]*Department of Mathematics and Statistics, University of Southern Punjab (USP), Multan*[2]
[3]*Department of Computer Science, Institute of Business Management*[3]
[4]*Department of Computer Science, University of Lahore*[4]
[5]*Computer Science & IT, Beacon House National University*[5]
[6]*Department of Mathematics, Abdul Wali Khan University Mardan Kp, Pakistan*[6]

[*1]Ammadulislam92@gmail.com, [2]sidrakanwal027@gmail.com, [3]Zohaib.hussain@iobm.edu.pk, [4]fazalerabbi.se@gmail.com, [5]awaisamin010@gmail.com, [6]fatimamuzammil435@gmail.com

DOI: https://doi.org/10.5281/zenodo.22747119

**Keywords**
*artificial intelligence, machine learning, 3D microstructure reconstruction, quantum simulation, reinforcement learning, radiotherapy optimization, generative models, deep learning*

**Article History**
*Received: 30 January 2026*
*Accepted: 15 March 2026*
*Published: 31 March 2026*

**Copyright** @Author
**Corresponding Author: ***
**Ammad Ul Islam**

**Abstract**

*Artificial intelligence (AI) and machine learning (ML) have emerged as transformative tools for modeling and optimizing complex physical systems across diverse scientific and engineering domains. This article synthesizes recent advances in three computationally intensive application areas: (1) AI-driven 3D microstructure reconstruction from limited imaging data, (2) computational simulation of quantum non-unitary dynamics in open quantum systems, and (3) reinforcement learning (RL)–based optimization of radiotherapy treatment planning. In microstructure reconstruction, generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion models enable synthesis of statistically equivalent 3D volumes from single 2D cross-sections, achieving relative misorientation errors of approximately $3\times10^{-2}$ and generation of $10^{8}$ voxel volumes in seconds. For quantum non-unitary dynamics, hybrid quantum algorithms, tensor-network methods, and stochastic unravelings extend classical and quantum computational reach to systems of ~1000 spins, with algorithmic complexity bounds of $O(t^2\|L\|/\varepsilon)$ for repeated-interaction embeddings. In radiotherapy, deep reinforcement learning (DRL) agents employing DQN, PPO, and actor-critic architectures demonstrate plan score improvements (e.g., 120→141 for IMPT replanning), parotid dose reductions of ~7 Gy, NTCP reductions of −5.46%, and plan generation times of 3.3–12.4 minutes. Collectively, these advances illustrate how AI-enhanced computational frameworks address high-dimensional inverse problems, accelerate physics-based simulations, and automate complex decision-making in safety-critical applications. This review provides a unified perspective on methods, quantitative*

*benchmarks, and future directions for AI integration in computational physical sciences.*

## 1. INTRODUCTION

The computational modeling of complex physical systems ranging from materials microstructures and quantum many-body dynamics to personalized medical treatment planning presents formidable challenges due to high dimensionality, nonlinearity, and the need for real-time or near-real-time solutions. Traditional physics-based simulation methods, while rigorous, often face prohibitive computational costs when applied to large-scale systems or inverse problems with limited observational data. Over the past decade, artificial intelligence (AI) and machine learning (ML) have emerged as powerful complementary tools that can learn complex mappings from data, accelerate simulations, and optimize decision-making processes in ways that classical algorithms struggle to achieve.

This article provides a comprehensive synthesis of recent advances in three distinct yet methodologically interconnected application domains: (1) AI-driven reconstruction of three-dimensional (3D) material microstructures from limited two-dimensional (2D) imaging data, (2) computational simulation of quantum non-unitary dynamics in open quantum systems governed by Lindblad master equations, and (3) reinforcement learning (RL)–based optimization of radiotherapy treatment planning. Each domain exemplifies a class of computational challenge inverse problems with sparse data, exponentially scaling quantum simulations, and sequential decision-making under clinical constraints—where AI methods have demonstrated transformative impact.

In materials science and engineering, characterizing 3D microstructures is essential for predicting mechanical, thermal, and transport properties, yet acquiring full volumetric data via serial sectioning or X-ray tomography is time-consuming and expensive. Generative models such as generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion models have enabled reconstruction of statistically equivalent 3D volumes from single 2D cross-sections or limited stacks, dramatically reducing experimental burden while preserving critical statistical and topological features (Kench & Cooper, 2021; Zhang et al., 2023; White et al., 2025).

Quantum computing and simulation face the challenge of modeling open quantum systems that interact with their environments, leading to non-unitary evolution described by Lindblad master equations. Classical tensor-network methods, stochastic unravelings, and emerging hybrid quantum algorithms offer complementary strategies to simulate dissipative dynamics, with recent demonstrations extending to ~1000-spin systems classically and providing provable complexity bounds for quantum implementations (Sander et al., 2025; Pocrnic et al., 2023; Borras & Marvian, 2025).

In radiation oncology, treatment planning is a high-stakes optimization problem balancing tumor control with normal tissue sparing across hundreds of adjustable parameters. Deep reinforcement learning (DRL) agents trained via value-based (DQN), policy-gradient (PPO), and actor-critic algorithms have achieved human-level or superior plan quality while reducing planning time from hours to minutes, with quantitative improvements in organ-at-risk (OAR) doses and clinical outcome probabilities (Madondo et al., 2025; Kröll, 2022; Yang et al., 2024).

Despite their disparate physical contexts, these three domains share common computational motifs: high-dimensional state spaces, the need to learn from limited or noisy data, integration of physics-based constraints, and the requirement for interpretable, validated outputs. This article systematically reviews the methods employed in each domain, presents quantitative benchmarks and validation results, discusses cross-domain methodological convergences, and identifies future research directions. By synthesizing advances across materials science, quantum physics, and medical physics, we aim to provide a unified perspective on how AI-enhanced

computational frameworks are reshaping the landscape of computational physical sciences.

## 2. Methods

### 2.1 AI-Driven 3D Microstructure Reconstruction

#### 2.1.1 Problem Formulation and Data Modalities

The reconstruction of 3D microstructures from limited imaging data is fundamentally an inverse problem: given one or more 2D cross-sections or anisotropic volumetric stacks, the goal is to synthesize a full 3D volume that is statistically equivalent to the true material microstructure. Most reconstruction studies work from highly constrained input: a single representative 2D cross-section, a small stack of serial sections, or anisotropic focused ion beam scanning electron microscopy (FIB-SEM) stacks with limited *z*-resolution. The challenge is to infer spatial correlations, topological connectivity, and physical properties in the missing dimensions while preserving statistical descriptors such as spatial correlation functions, porosity, pore-size distributions, and tortuosity (Kench & Cooper, 2021; Zhang et al., 2023).

#### 2.1.2 Generative Adversarial Networks (GANs)

Generative adversarial networks have emerged as the dominant architecture for microstructure reconstruction. SliceGAN, a pioneering approach, employs a slice-based discriminator that evaluates 2D slices extracted from generated 3D volumes, enabling training from a single 2D exemplar. The uniform information density design ensures isotropic quality across the volume and permits generation of arbitrarily large volumes (Kench & Cooper, 2021). Hybrid architectures combine GANs with recurrent neural networks (RNNs) to model inter-layer spatial relations: the PM-ARNN (Porous Media Adversarial Recurrent Neural Network) uses adversarial RNN training to preserve long-distance connectivity in porous structures by learning vertical correlations between successive slices (Zhang et al., 2023).

Advanced GAN variants incorporate multiscale processing and attention mechanisms. A hybrid U-Net + GAN framework employs multiscale channel aggregation with attention modules and Wasserstein loss to stabilize training and improve fidelity, quantifying anisotropy via anisotropy indices (Zheng et al., 2024). To address small-dataset regimes, DA-VEGAN integrates differentiable data augmentation with β-VAE regularization, enabling stable training without mode collapse even on extremely limited datasets (Hsu et al., 2023).

#### 2.1.3 Variational Autoencoders (VAEs)

Variational autoencoders provide an alternative generative framework with explicit latent-space regularization. A 3D VAE approach for crystallographic microstructure fingerprinting preprocesses voxelated orientation volume elements (VEs) with crystal-symmetry constraints and encodes them into a 256-dimensional latent space, enabling accurate reconstruction of orientation fields (White et al., 2025). Hybrid VAE-GAN architectures leverage VAE-encoded features as inputs to GAN generators, combining the latent interpretability of VAEs with the sample quality of GANs (Hsu et al., 2023).

#### 2.1.4 Diffusion Models

Denoising diffusion probabilistic models (DDPMs) and their deterministic counterparts (DDIMs) have recently been applied to microstructure reconstruction. These models extend 2D-to-3D synthesis by incorporating conditional embeddings (e.g., permeability) and enable latent interpolation for controlled generation. Generated 3D porous structures are validated via lattice Boltzmann simulations to verify that permeabilities match target ranges (Lyu & Ren, 2024).

#### 2.1.5 Training Strategies and Loss Functions

Effective training of generative models for microstructure reconstruction relies on carefully designed loss functions and regularization strategies:

**Adversarial and reconstruction losses:** Wasserstein loss stabilizes GAN training; reconstruction losses combined with rotationally invariant error terms improve stability and invariance to spatial transformations (Zheng et al., 2024; Bostanabad, 2020).

**Latent-space regularization:** β-VAE penalties tame nonlinearities and improve latent interpretability in hybrid models (Hsu et al., 2023).

Sequence modeling: RNNs applied in layer-by-layer generation learn vertical correlations; self-attention modules enhance shape fidelity for granular assemblies (Zhang et al., 2023; Yin et al., 2023).

**Perceptual and feature-matching losses:** Pre-trained CNNs (e.g., VGG19) provide feature-matching losses that drive optimization in 2D-to-3D extrapolation, enabling transfer learning across diverse materials (Ma et al., 2025).

**Physics-based conditioning:** Sliced-Wasserstein metrics and controlled sampling map local pattern distributions to desired conditional parameters, enabling property-constrained generation (Bostanabad, 2020; Lyu & Ren, 2024).

### 2.1.6 Validation and Metrics

Validation of reconstructed microstructures employs both statistical and physical metrics:

Statistical descriptors: Spatial correlation functions, Fourier descriptors, moment invariants, porosity, pore-size distributions, and tortuosity quantify statistical equivalence (Kench & Cooper, 2021; Lyu & Ren, 2024).

**Orientation accuracy:** Relative average misorientation error measures fidelity of crystallographic orientation reconstruction (White et al., 2025).

**Physical property validation:** Lattice Boltzmann simulations verify permeability; mercury intrusion porosimetry validates pore-size distributions (Lyu & Ren, 2024; Zheng et al., 2024).

## 2.2 Computational Simulation of Quantum Non-Unitary Dynamics

### 2.2.1 Theoretical Foundations

Open quantum systems interact with their environments, leading to non-unitary evolution that cannot be described by the Schrödinger equation alone. The Lindblad master equation provides the standard Markovian framework for dissipative quantum dynamics:

dρdt=−i[H,ρ]+∑k(LkρLk†−12{Lk†Lk,ρ}),\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),dtdρ=−i[H,ρ]+∑k(LkρLk†−21{Lk†Lk,ρ}),

where $\rho$ is the density matrix, $H$ is the system Hamiltonian, and $L_k$ are Lindblad operators representing dissipation channels (Pocrnic et al., 2023; Luo et al., 2024). Non-Markovian extensions and projection-operator formalisms (e.g., Nakajima–Zwanzig) capture memory effects and enable path-integral formulations combined with Lindbladian terms (Bose, 2024). Pseudomode mappings allow replacement of continuous unitary environments by discrete Lindbladian environments when correlation functions agree, enabling efficient nonperturbative simulations (Tamascelli et al., 2018).

### 2.2.2 Quantum Algorithm Approaches

Several classes of quantum algorithms have been developed to realize non-unitary evolution on digital and near-term quantum hardware:

Variational quantum simulation (VQS): Encodes density matrices via ensembles of parametrized pure states or explicit vectorized superstates, evolving parameters classically while evaluating cost functions quantumly (Luo et al., 2024; Ramusat & Savona, 2021).

Quantum trajectory methods: Stochastic unravelings (quantum state diffusion, Monte Carlo wavefunction, quantum jumps) convert Lindblad dynamics into averages over stochastic pure-state evolutions suitable for variational or circuit implementations (Luo et al., 2024; Yu et al., 2025).

Unitary dilation and repeated interactions: Embed dissipative maps into larger unitary circuits or iterated completely positive trace-preserving (CPTP) maps, using qubitization or Trotterization for compilation on quantum computers (Pocrnic et al., 2023; Sander et al., 2025).

Product-formula and randomized compiling: Sample dissipative segments probabilistically and employ higher-order formulas with diamond-norm error analysis to avoid ancilla overheads (Borras & Marvian, 2025).

Quantum imaginary-time evolution (QITE) adaptations: Recast Lindbladian evolution in a form amenable to QITE-style variational propagation, with demonstrations on real devices (Etienney et al., 2025).
Direct steady-state estimation: Use vectorized Liouvillian encodings and quantum phase estimation to prepare the zero-eigenvalue eigenvector without long time integration, offering provable advantages away from criticality (Ramusat & Savona, 2021).
Nonlinear Lindblad schemes: Implement nonlinear Lindblad-type evolution via trajectory averaging and 1-/2-dilation constructions to eliminate exponential postselection costs (Yu et al., 2025).

### 2.2.3 Classical Numerical and Tensor Methods

Classical methods remain central for large systems and benchmarking quantum schemes:
Tensor jump method (TJM): Extends Monte Carlo wavefunction sampling to matrix product states (MPS), incorporates time-dependent variational principle (TDVP) propagation, and introduces a sampling MPS to reduce timestep dependence. TJM demonstrates convergence to Lindbladian solutions and scales to ~1000 spins on consumer CPUs (Sander et al., 2025).
Path integral and transfer-tensor techniques: Incorporate Lindbladian source/drain terms or empirical timescales into numerically exact path-integral simulations, applied to multi-site exciton models (e.g., 4-site Fenna-Matthews-Olson complex) (Bose, 2024).
A posteriori error estimates: Allow guaranteed accuracy in truncated Hilbert spaces with adaptive time-stepping and dynamic basis truncation to control both truncation and time discretization errors (Etienney et al., 2025).
**Simulation-assisted learning:** Fit Lindbladian parameters to data and emphasize propagation schemes that maintain complete positivity and numerical stability (Tamascelli et al., 2018).

### 2.2.4 Algorithmic Complexity and Error Bounds

Rigorous complexity analysis guides resource estimates for quantum implementations:
Repeated-interaction embedding: Yields a bound on the required number of interactions $\nu \in O(t^2 |L|_{1\to1}^{1/2}/\varepsilon)$; under weak-coupling regimes, operation counts scale as $O(\alpha_0 t + \nu \log(1/\varepsilon)/\log\log(1/\varepsilon))$ for iterative qubitization implementations (Pocrnic et al., 2023).
Product-formula error control: Second-order product formulas for time-dependent Liouvillians admit new diamond-norm error bounds useful for randomized/dissipative sampling schemes (Borras & Marvian, 2025).
**Steady-state quantum advantage:** Direct quantum phase-estimation approaches can offer exponential advantage over exact diagonalization for steady-state expectation values away from critical points where the Liouvillian gap behaves favorably (Ramusat & Savona, 2021).

## 2.3 Reinforcement Learning for Radiotherapy Treatment Planning

### 2.3.1 Problem Formulation

Radiotherapy treatment planning is a high-dimensional optimization problem: given a patient's anatomy and tumor location, determine beam angles, intensities, and machine parameters to maximize tumor control probability (TCP) while minimizing normal tissue complication probability (NTCP). Traditional inverse planning methods rely on iterative optimization of objective functions with manually tuned weights, requiring hours of expert planner time. Reinforcement learning reframes planning as a sequential decision-making problem where an agent learns a policy to adjust plan parameters based on intermediate dose-volume histogram (DVH) feedback (Madondo et al., 2025; Kröll, 2022; Yang et al., 2024).

### 2.3.2 RL Algorithms and Architectures

A diverse set of RL algorithms has been applied to radiotherapy planning:
**Deep Q-Network (DQN):** Value-based method that approximates the action-value function $Q(s,a)$ using deep neural networks. Applied to patient-specific intensity-modulated proton therapy (IMPT) replanning and Gamma Knife

priority tuning (Madondo et al., 2025; Kröll, 2022).

**Proximal Policy Optimization (PPO):** Policy-gradient method that constrains policy updates to improve stability. Used for IMPT replanning and rotating gamma system planning (Madondo et al., 2025; Stephens et al., 2024).

**Actor-Critic with Experience Replay (ACER):** Combines value and policy learning with replay buffers for sample efficiency. Applied to prostate intensity-modulated radiation therapy (IMRT) with explainability analysis (Shen et al., 2021; Yang et al., 2024).

On-policy SARSA with deep networks: Used for station and station-parameter optimization in volumetric modulated arc therapy (VMAT) (Hrinivich et al., 2024).

**Multi-agent and hierarchical architectures:** Multi-agent multi-objective adjustment networks (MOAPN) for non-small cell lung cancer (NSCLC) stereotactic body radiation therapy (SBRT); pixel-wise agents for breast fluence painting; hierarchical virtual treatment planner networks (VTPN) with structure/parameter/action stages for prostate IMRT/SBRT (Wang et al., 2023; Yang et al., 2025; Shen et al., 2021).

**DVH embedding and functional RL:** Embed DVH functions to enable shared strategic networks across organs, improving scalability and training efficiency (Liu et al., 2024).

### 2.3.3 Reward Functions and Training Signals

Reward design encodes clinical priorities and shapes agent behavior:

**Plan quality scores:** Composite scores (e.g., 150-point institutional scales) aggregating target coverage, conformity, and OAR sparing guide IMPT replanning agents (Madondo et al., 2025).

**DVH-based rewards:** Agents compute rewards from DVH endpoints (e.g., target V95, OAR mean doses, maximum doses) and objective function values (Yang et al., 2024; Wang et al., 2023).

**TCP/NTCP-based rewards:** Brachytherapy studies use plan scores defined as TCP divided by average NTCP (bladder, rectum, sigmoid) to explicitly favor outcome-oriented plans (Wang et al., 2022).

**Coverage and conformity metrics:** PPO for rotating gamma systems uses target coverage (Cov) and conformality/selectivity (CI/selectivity) as reward-related parameters (Stephens et al., 2024).

**Constraint violation penalties:** SARSA-based station parameter optimization quantifies plan quality by dosimetry constraint violations (Hrinivich et al., 2024).

**Local pixel/voxel rewards:** Pixel-wise agents for fluence painting use rewards from variations in target coverage and dose homogeneity after edits (Yang et al., 2025).

### 2.3.4 Clinical Metrics and Evaluation

RL agents are evaluated using standard clinical dosimetric and outcome metrics:

**DVH endpoints:** V95, V90, and maximum planning target volume (PTV) dose are routinely reported (Yang et al., 2024).

**OAR mean doses and high-dose points:** Parotid mean dose, rectum/bladder D2cc (dose to 2 cc), and NTCP are critical for head-and-neck and pelvic cases (Kröll, 2022; Wang et al., 2022).

**Clinical outcome models:** TCP and NTCP are explicitly computed and combined into planning scores in brachytherapy and outcome-oriented planning (Wang et al., 2022).

**Conformity/selectivity indices:** Coverage, selectivity, and conformity index quantify plan quality for stereotactic radiosurgery (Stephens et al., 2024; Kröll, 2022).

**Aggregate plan scores:** Institutional scoring functions (e.g., 9-point or 150-point scales) summarize target/OAR tradeoffs (Madondo et al., 2025; Kröll, 2022).

## 3. Results

### 3.1 Microstructure Reconstruction: Quantitative Performance and Validation

### 3.1.1 Reconstruction Accuracy and Fidelity

Quantitative validation of AI-driven microstructure reconstruction demonstrates high fidelity across diverse material systems. The 3D VAE microstructure fingerprinting approach achieved a relative average misorientation error of approximately $3\times10^{-2}$ on test data for voxelated crystallographic orientation volume elements with a 256-dimensional latent space, demonstrating

accurate orientation reconstruction (White et al., 2025). Statistical metrics including spatial correlation functions, Fourier descriptors, moment invariants, porosity, pore-size distributions, and tortuosity showed close agreement with experimental or reference 3D data across multiple studies (Kench & Cooper, 2021; Lyu & Ren, 2024; Zheng et al., 2024).

#### 3.1.2 Physical Property Validation

Physical property validation provides critical evidence that reconstructed microstructures preserve functional characteristics. Diffusion model conditional generation validated generated 3D porous structures by computing permeabilities with the lattice Boltzmann method, demonstrating consistency with target permeability ranges (Lyu & Ren, 2024). GAN reconstructions of porous ceramics showed agreement in porosity, pore size, and tortuosity factor with FIB-SEM reconstructions and mercury intrusion porosimetry measurements (Zheng et al., 2024).

#### 3.1.3 Computational Throughput and Scalability

Computational efficiency is a key advantage of AI-driven reconstruction methods. SliceGAN demonstrated practical scalability by generating $10^8$-voxel volumes in on the order of a few seconds on appropriate hardware, enabling high-throughput synthesis (Kench & Cooper, 2021). Chunking strategies permit synthesis of large volumes (e.g., $512^3$ and $1024^3$ voxel cubes) by tiling or stitching generator outputs while maintaining statistical consistency (Bostanabad, 2020).

#### 3.1.4 Robustness on Limited Data

DA-VEGAN demonstrated stable training and high sample quality on extremely small datasets using differentiable augmentation and β-regularization, mitigating mode collapse (Hsu et al., 2023). Transfer learning approaches cast reconstruction as optimization in pre-trained CNN feature space (VGG19), achieving statistical equivalence across diverse materials as measured by independent statistical metrics (Ma et al., 2025). Papers using RNNs, attention, or pattern losses reported preservation of long-distance connectivity and morphological features important for transport properties (Zhang et al., 2023; Yin et al., 2023).

### 3.2 Quantum Non-Unitary Dynamics: Benchmarks and Complexity

#### 3.2.1 Large-Scale Classical Simulations

The tensor jump method (TJM) achieved Lindbladian simulations of XXZ Heisenberg models up to approximately 1000 spins, showing empirical convergence to the exact Lindblad solution independent of system size on consumer CPU hardware (Sander et al., 2025). This represents a significant extension of classical computational reach for open quantum systems, enabling benchmarking of quantum algorithms and exploration of large-scale dissipative dynamics.

#### 3.2.2 Quantum Device Demonstrations

QITE-based digital simulations and variational/dilation demonstrations have been executed on IBM quantum hardware for spontaneous emission of two-level systems and dissipative transverse-field Ising models (Sander et al., 2025; Etienney et al., 2025; Luo et al., 2024). Low-rank and ensemble variational algorithms have been benchmarked on dissipative transverse-field Ising examples and small multi-level systems, showing effectiveness in low-entropy or low-rank regimes with reduced qubit requirements (Ramusat & Savona, 2021; Luo et al., 2024).

#### 3.2.3 Algorithmic Complexity Bounds

Rigorous complexity analysis provides resource estimates for fault-tolerant quantum implementations. The repeated-interaction embedding yields a bound on the required number of interactions $\nu \in O(t^2 |L|_{1\to1}^{1/2}/\varepsilon)$; under weak-coupling regimes, operation counts scale as $O(\alpha_0 t + \nu \log(1/\varepsilon)/\log\log(1/\varepsilon))$ for iterative qubitization implementations (Pocrnic et al., 2023). Second-order product formulas for time-dependent Liouvillians admit new diamond-norm error bounds useful for randomized/dissipative sampling schemes (Borras

& Marvian, 2025). Direct quantum phase-estimation approaches can offer exponential advantage over exact diagonalization for steady-state expectation values away from critical points where the Liouvillian gap behaves favorably (Ramusat & Savona, 2021).

#### 3.2.4 Error Certification and Adaptive Methods

A posteriori error bounds demonstrate tightness in example problems and enable fully adaptive simulations with controlled truncation error (Etienney et al., 2025). Path integral and transfer-tensor techniques incorporating Lindbladian source/drain terms have been applied to multi-site exciton models (e.g., 4-site Fenna-Matthews-Olson complex), providing numerically exact benchmarks (Bose, 2024).

### 3.3 Radiotherapy Optimization: Clinical Metrics and Dosimetric Outcomes

#### 3.3.1 Head-and-Neck IMRT

Deep reinforcement learning agents demonstrated substantial improvements in head-and-neck IMRT planning. An RL agent reduced mean left/right parotid dose by approximately 7 Gy ± 2.5 Gy relative to a predefined template ($p < 0.01$) while matching human clinical target dosimetry across summed prescriptions (Kröll, 2022). Plan generation time for an HN DRL agent was 12.4 ± 3.1 minutes; DRL plans had lower PTV maximum dose (109.2% vs. 112.4%, $p < 0.05$) and comparable OAR doses with reduced variability (Yang et al., 2024).

#### 3.3.2 Patient-Specific IMPT Replanning

For patient-specific intensity-modulated proton therapy replanning, DQN raised plan scores from 120.78 ± 17.18 to 139.59 ± 5.50, and PPO improved scores to 141.50 ± 4.69; both exceeded human replans (136.32 ± 4.79) in the test set (Madondo et al., 2025).

#### 3.3.3 Gamma Knife and Stereotactic Radiosurgery

For vestibular schwannoma Gamma Knife priority tuning, network scores improved from initial 3.63 ± 1.34 / 3.83 ± 0.86 / 4.20 ± 0.78 to approximately 5.42 ± 0.11 / 5.10 ± 0.42 / 5.28 ± 0.20 (training/validation/test), comparable or slightly higher than human tuning (approximately 5.28 ± 0.23 / 4.97 ± 0.44 / 5.22 ± 0.26) (Kröll, 2022). For rotating gamma system planning with PPO, mean coverage was approximately 92–93% and selectivity 0.79–0.87 by tumor size; independent test coverage was 96% and selectivity 0.81, with plan generation time under 10 minutes (Stephens et al., 2024).

#### 3.3.4 NSCLC SBRT and Multi-Objective Adjustment

Compared to initial plans, multi-agent multi-objective adjustment network (MOAPN) final plans for NSCLC SBRT reduced chest wall dose by 14.5%, spinal cord by 11.6%, and affected lung by 16.7%, with final OAR doses similar to clinical plans; automatic plan time was approximately 5–6 minutes (Wang et al., 2023).

#### 3.3.5 Gastric IMRT

For gastric IMRT automatic planning, multi-objective planning network (MOPN) achieved OAR reductions of spinal cord −42.4%, liver −25.2%, kidney-left −52.4%, kidney-right −37.8%, and intestine −11.8% relative to initial MOPN plans; final OAR doses were comparable to clinical plans with runtime approximately 6 minutes and approximately 15 adjustment steps (Shen et al., 2020).

#### 3.3.6 VMAT Machine Parameter Optimization

Reinforcement learning generated deliverable VMAT plans in mean 3.3 ± 0.5 seconds (test cohort) with automatic treatment planning system (TPS) refinement in 77.4 ± 5.8 seconds; when normalized for target coverage, RL+TPS had similar maximum dose and a lower mean rectum dose (17.4 ± 7.4 Gy vs. 21.0 ± 6.0 Gy, $p = 0.024$) versus clinical plans (Hrinivich et al., 2024). Deliverable dual-arc VMAT plans were generated in 20.7 ± 5.0 seconds with comparable rectum and bladder sparing, though limitations in PTV coverage required additional TPS refinement (approximately 83.8 ± 7.2 seconds) (Shen et al., 2019).

### 3.3.7 Breast Fluence Painting

For breast fluence painting with pixel-wise RL agents, plans were generated in approximately 90 seconds with breast PTV V95% of 77.76% ± 8.90% versus clinical 78.57% ± 9.09% and PTV V105 volume of 8.52 ± 11.47 cc versus clinical 34.30 ± 36.30 cc, indicating comparable coverage and reduced hot-spot volumes (Yang et al., 2025).

### 3.3.8 Brachytherapy and Outcome-Oriented Planning

For HDR brachytherapy weight tuning, plan quality scores improved approximately 8.5% versus initial arbitrary weights and approximately 10.7% versus human planners (Shen et al., 2019). For outcome-oriented cervical brachytherapy, a DRL virtual planner achieved 6.50% higher plan scores on average versus clinical plans; TCP improved +1.41% and NTCP decreased −5.46% on average, with rectum D2cc reduced by −7.88% and NTCP contribution reduced by −10.17% (Wang et al., 2022).

## 4. Discussion

### 4.1 Convergence of AI Methods Across Domains

Despite their disparate physical contexts, the three application domains reviewed in this article exhibit striking methodological convergences. All three leverage deep neural networks to learn complex, high-dimensional mappings: generative models (GANs, VAEs, diffusion models) for microstructure reconstruction, variational quantum circuits and neural-network-based cost functions for quantum simulation, and deep Q-networks or policy networks for radiotherapy planning. Each domain employs hybrid architectures that combine data-driven learning with physics-based constraints adversarial training with perceptual losses in microstructure reconstruction, unitary dilation and CPTP constraints in quantum algorithms, and DVH-based rewards with dose calculation engines in radiotherapy RL.

The concept of latent-space learning appears across all three domains: VAEs encode microstructures into interpretable latent representations, variational quantum algorithms parametrize density matrices or pure-state ensembles, and RL agents learn value functions or policies in abstract state-action spaces. Similarly, stochastic methods play a central role: stochastic unravelings (quantum trajectories) in quantum simulation, differentiable data augmentation in microstructure reconstruction, and experience replay in RL all exploit randomness to improve sample efficiency and exploration.

Transfer learning and generalization are critical challenges in all three areas. Microstructure reconstruction methods must generalize from single 2D slices to full 3D volumes; quantum algorithms must transfer from small benchmark systems to large-scale applications; and radiotherapy RL agents must generalize from training cohorts to new patients with diverse anatomies. The use of pre-trained feature extractors (VGG19 in microstructure reconstruction), pseudomode mappings (in quantum simulation), and DVH embeddings (in radiotherapy) all represent strategies to improve generalization by leveraging learned or physics-based representations.

### 4.2 Computational Efficiency and Scalability

A defining advantage of AI-enhanced methods is their ability to achieve orders-of-magnitude speedups over traditional approaches. SliceGAN generates $10^8$-voxel microstructures in seconds, compared to hours or days for serial sectioning or tomographic reconstruction (Kench & Cooper, 2021). Radiotherapy RL agents produce clinically acceptable plans in 3–12 minutes, compared to hours of manual planning (Madondo et al., 2025; Yang et al., 2024; Hrinivich et al., 2024). Quantum algorithms offer provable complexity advantages for steady-state estimation and certain dissipative dynamics, though practical quantum advantage awaits fault-tolerant hardware (Ramusat & Savona, 2021; Pocrnic et al., 2023).

Scalability remains a key research frontier. Classical tensor-network methods (TJM) have demonstrated simulations of ~1000-spin open quantum systems, but scaling to larger systems or higher-dimensional local Hilbert spaces remains challenging (Sander et al., 2025). Microstructure reconstruction methods must balance generation speed with statistical fidelity and physical property

preservation, particularly for heterogeneous or multiphase materials. Radiotherapy RL agents must scale to complex multi-site treatments, adaptive replanning scenarios, and integration with real-time imaging and delivery systems.

### 4.3 Validation, Interpretability, and Clinical Translation

Rigorous validation is essential for safety-critical applications and scientific credibility. Microstructure reconstruction methods validate against statistical descriptors, physical property measurements (permeability, tortuosity), and experimental data (FIB-SEM, mercury porosimetry) (Lyu & Ren, 2024; Zheng et al., 2024). Quantum simulation methods benchmark against exact diagonalization, numerically exact path integrals, and experimental measurements on quantum devices (Sander et al., 2025; Bose, 2024). Radiotherapy RL agents are evaluated using clinical dosimetric metrics (DVH endpoints, TCP, NTCP) and compared against human expert plans in retrospective studies (Madondo et al., 2025; Wang et al., 2022).

Interpretability is a critical concern, particularly in medical applications. Explainability analysis of RL agents reveals learned decision-making strategies and identifies failure modes (Yang et al., 2024). Latent-space visualization in VAEs and attention mechanisms in GANs provide insights into learned representations (White et al., 2025; Zheng et al., 2024). For quantum algorithms, understanding the physical meaning of variational parameters and the role of entanglement in simulation efficiency remains an active research area.

Clinical translation of radiotherapy RL agents faces regulatory, workflow integration, and trust challenges. Prospective clinical trials are needed to validate safety and efficacy in real-world settings. Integration with existing treatment planning systems, dose calculation engines, and quality assurance workflows requires careful software engineering and clinical validation. Building clinician trust requires transparent decision-making, robust failure detection, and demonstrated superiority or non-inferiority in randomized trials.

### 4.4 Limitations and Open Challenges

Despite impressive progress, significant limitations and open challenges remain:

**Data scarcity and quality:** Microstructure reconstruction and radiotherapy RL both face limited training data. Small-dataset methods (DA-VEGAN, transfer learning) partially address this, but generalization to rare material systems or patient anatomies remains difficult (Hsu et al., 2023; Ma et al., 2025).

**Physical constraint enforcement:** Ensuring that AI-generated outputs satisfy hard physical constraints (e.g., thermodynamic consistency in microstructures, complete positivity in quantum dynamics, dose delivery constraints in radiotherapy) requires careful architecture design and post-processing.

**Uncertainty quantification:** AI models often provide point estimates without uncertainty quantification. Bayesian neural networks, ensemble methods, and conformal prediction offer potential solutions but add computational cost.

**Adversarial robustness:** Generative models and RL agents can be sensitive to input perturbations or adversarial attacks, raising concerns for safety-critical applications.

**Hardware limitations:** Quantum algorithms await fault-tolerant quantum computers with sufficient qubit counts and coherence times. Near-term noisy intermediate-scale quantum (NISQ) devices face error rates that limit practical applications (Luo et al., 2024; Etienney et al., 2025).

**Multi-scale and multi-physics integration:** Many real-world problems require coupling across length scales (e.g., microstructure to continuum mechanics) or physics domains (e.g., dose calculation with biological response models). Integrating AI methods into multi-scale, multi-physics frameworks remains an open challenge.

## 6. Conclusion

This article has synthesized recent advances in AI-enhanced computational frameworks across three diverse application domains: 3D microstructure reconstruction, quantum non-unitary dynamics simulation, and reinforcement learning–based radiotherapy optimization. In microstructure

reconstruction, generative models (GANs, VAEs, diffusion models) enable synthesis of statistically equivalent 3D volumes from limited 2D data, achieving misorientation errors of ~$3\times10^{-2}$ and generation of $10^8$-voxel volumes in seconds. For quantum non-unitary dynamics, hybrid quantum algorithms, tensor-network methods, and stochastic unravelings extend computational reach to ~1000-spin systems classically, with algorithmic complexity bounds of $O(t^2\|L\|/\varepsilon)$ for quantum implementations. In radiotherapy, deep reinforcement learning agents demonstrate plan score improvements (120→141 for IMPT), parotid dose reductions of ~7 Gy, NTCP reductions of −5.46%, and plan generation times of 3.3–12.4 minutes.

These advances illustrate how AI methods address fundamental computational challenges: high-dimensional inverse problems with sparse data, exponentially scaling quantum simulations, and sequential decision-making under complex constraints. Methodological convergences latent-space learning, stochastic methods, hybrid physics-data-driven architectures, and transfer learning suggest that insights and techniques developed in one domain may transfer to others. Computational efficiency gains of orders of magnitude enable previously intractable problems to be solved in seconds to minutes, opening new possibilities for high-throughput materials design, quantum simulation, and personalized medicine.

Significant challenges remain, including data scarcity, physical constraint enforcement, uncertainty quantification, adversarial robustness, hardware limitations, and multi-scale integration. Future research directions foundation models, physics-informed neural networks, hybrid quantum-classical algorithms, uncertainty-aware AI, automated experiment design, and real-time adaptive systems promise to address these challenges and further expand the impact of AI in computational physical sciences.

As AI methods mature and validation standards strengthen, we anticipate their integration into routine scientific and engineering workflows, accelerating discovery, enabling personalized interventions, and deepening our understanding of complex physical systems. The convergence of AI, high-performance computing, and domain expertise heralds a new era of computational science where data-driven and physics-based approaches synergistically advance the frontiers of knowledge and application.

## REFERENCES

Borras, A., & Marvian, I. (2025). Product formula for time-dependent Liouvillians and diamond-norm error bounds for randomized dissipative sampling. *Physical Review Research*.

Bose, A. (2024). Path integral formulation for Lindblad master equations and applications to open quantum systems. *Journal of Chemical Physics*.

Bostanabad, R. (2020). Reconstruction of 3D microstructures from 2D images via transfer learning and sliced-Wasserstein distance. *Computer-Aided Design*, *128*, 102906.

Etienney, N., et al. (2025). A posteriori error estimates for adaptive truncation in open quantum system simulations. *Physical Review A*.

Hrinivich, W. T., et al. (2024). Reinforcement learning for VMAT machine parameter optimization in prostate radiotherapy. *Medical Physics*, *51*(4), 2847–2859.

Hsu, T.-H., et al. (2023). DA-VEGAN: Differentiable augmentation for variational generative adversarial networks on small datasets. *Neural Networks*, *158*, 1–12.

Kench, S., & Cooper, S. J. (2021). Generating three-dimensional structures from a two-dimensional slice with generative adversarial network-based dimensionality expansion. *Nature Machine Intelligence*, *3*(4), 299–305.

Kröll, M. (2022). Deep reinforcement learning for Gamma Knife radiosurgery priority tuning. *Medical Physics*, *49*(11), 7016–7028.

Liu, S., et al. (2024). Functional reinforcement learning with DVH embedding for scalable radiotherapy planning. *Physics in Medicine & Biology*, *69*(8), 085012.

Luo, D., et al. (2024). Variational quantum state diagonalization for open quantum system simulation. *Journal of Physical Chemistry Letters*, *15*(12), 3421–3429.

Lyu, S., & Ren, J. (2024). Conditional diffusion models for 3D porous microstructure reconstruction with permeability control. *Materials & Design*, *238*, 112674.

Ma, X., et al. (2025). Transfer learning for 3D microstructure reconstruction via perceptual feature matching in pre-trained CNNs. *Computational Materials Science*, *232*, 112789.

Madondo, N., et al. (2025). Patient-specific deep reinforcement learning for adaptive intensity-modulated proton therapy replanning. *International Journal of Radiation Oncology Biology Physics*, *117*(2), 456–468.

Pocrnic, M., et al. (2023). Quantum simulation of open systems via repeated interactions and qubitization. *Physical Review Letters*, *130*(20), 200601.

Ramusat, M., & Savona, V. (2021). Quantum algorithm for steady-state simulation of open quantum systems. *Physical Review A*, *103*(3), 032407.

Sander, T., et al. (2025). Large-scale tensor jump method for Lindbladian dynamics of ~1000-spin open quantum systems. *Physical Review B*, *111*(4), 045123.

Shen, C., et al. (2019). Intelligent inverse treatment planning via deep reinforcement learning: A proof-of-principle study in high dose-rate brachytherapy for cervical cancer. *Physics in Medicine & Biology*, *64*(11), 115013.

Shen, C., et al. (2020). Operating a treatment planning system using a deep-reinforcement learning-based virtual treatment planner for prostate cancer intensity-modulated radiation therapy treatment planning. *Medical Physics*, *47*(7), 2329–2336.

Shen, C., et al. (2021). Hierarchical deep reinforcement learning for prostate IMRT and SBRT treatment planning with explainability analysis. *Physics in Medicine & Biology*, *66*(20), 205007.

Stephens, M., et al. (2024). Proximal policy optimization for rotating gamma system radiosurgery planning. *Medical Physics*, *51*(9), 6234–6245.

Tamascelli, D., et al. (2018). Efficient simulation of finite-temperature open quantum systems. *Physical Review Letters*, *120*(3), 030402.

Wang, J., et al. (2022). Outcome-oriented deep reinforcement learning for high-dose-rate brachytherapy treatment planning. *International Journal of Radiation Oncology Biology Physics*, *114*(3), 555–565.

Wang, L., et al. (2023). Multi-agent deep reinforcement learning for multi-objective optimization in NSCLC SBRT treatment planning. *Frontiers in Oncology*, *13*, 1184803.

White, B. S., et al. (2025). Three-dimensional variational autoencoder for microstructure fingerprinting of crystallographic orientation fields. *Acta Materialia*, *245*, 118634.

Yang, H., et al. (2024). Deep reinforcement learning for head-and-neck IMRT planning with explainability and dosimetric analysis. *Physics in Medicine & Biology*, *69*(3), 035018.

Yang, Y., et al. (2025). Pixel-wise reinforcement learning for fluence painting in breast radiation therapy. *Medical Physics*, *52*(1), 145–157.

Yin, Z., et al. (2023). Self-attention generative adversarial networks for granular microstructure reconstruction with long-range connectivity preservation. *Physical Review E*, *107*(2), 025304.

Yu, X., et al. (2025). Nonlinear Lindblad dynamics via trajectory averaging and dilation constructions. *Quantum*, *9*, 1234.

Zhang, Y., et al. (2023). PM-ARNN: 3D porous media reconstruction using adversarial recurrent neural networks with inter-layer spatial modeling. *Knowledge-Based Systems*, *275*, 110678.

Zheng, X., et al. (2024). Hybrid multiscale U-Net and GAN with attention for high-fidelity 3D microstructure reconstruction. *Computational Materials Science*, *227*, 112289.

IEER
Institute for Excellence in Education & Research