ADVANCING FLASH RADIOTHERAPY WITH QUANTUM COMPUTING: OPPORTUNITIES AND CHALLENGES

SUMMARY: FLASH-RT advances hinge on ultra-high dose rate machines, validated dosimetry, and biology; quantum computing aids the supporting calculations.


By **JAMES C. L. CHOW**[1,2,3] <james.chow@uhn.ca>, **HARRY E. RUDA**[3] <harry.ruda@utoronto.ca>

[1]Department of Medical Physics, Princess Margaret Cancer Centre, University Health Network, Toronto, ON, M5G 1X6, Canada
[2]Department of Radiation Oncology, University of Toronto, Toronto, ON, M5T 1P5, Canada
[3] Department of Materials Science and Engineering, University of Toronto, Toronto, ON, M5S 3E4, Canada

Radiotherapy is used in nearly half of all cancer treatments, and its precision has improved steadily to better spare healthy tissues [1]. A newer approach, FLASH radiotherapy (FLASH-RT), pushes this further by delivering radiation at ultra-high dose rates (UHDR) with the potential to reduce side effects while maintaining tumour control [2]. Today, translation is limited chiefly by non-computational barriers: (i) the lack of clinically deployable UHDR treatment machines, (ii) incomplete understanding of the biological mechanism underlying the FLASH effect [3], and (iii) the difficulty of measuring dose accurately at UHDR with conventional detectors [4]. Computation plays a supporting role: classical Monte Carlo (MC) already underpins beamline and dosimeter design and benchmarking, and prospectively, quantum computing (QC) could accelerate specific calculations (e.g., MC-adjacent studies, optimisation, and data analysis) within broader experimental and clinical programmes. QC does not replace the need for validated UHDR dosimetry, FLASH-capable hardware, or biological insights, but it may help move those efforts faster and at larger scale [5]. In this review, we explore how QC might one day help us meet demanding computational needs of FLASH-RT, and what challenges still lie ahead.

WHAT IS FLASH-RT?

FLASH-RT is a relatively new approach to deliver radiation treatment for cancer. Unlike conventional radiotherapy, which typically delivers radiation over several minutes at standard dose rates, FLASH-RT administers the entire treatment dose in a fraction of a second using dose rates of 40 Gy per second or higher. This ultra-rapid delivery has sparked interest because of a phenomenon known as the “FLASH effect” [6]. In preclinical studies, mainly in small animal models, researchers observed that FLASH-RT could damage cancer cells, similar to conventional treatments, but with reduced harm to nearby healthy tissues. This suggests a potential breakthrough: delivering effective treatment with fewer adverse side effects [7]. Table 1 shows the differences between FLASH-RT and conventional radiotherapy (CONV-RT) in various aspects such as treatment time, dose rate and normal cell sparing [3].

TABLE 1

Comparison of FLASH-RT and conventional radiotherapy [3].

Aspect	FLASH-RT	CONV-RT
Treatment Time	Ultra-fast (milliseconds)	Typically seconds to minutes
Dose Rate	Extremely high (>40 Gy/s)	Moderate to high (0.001-0.4 Gy/s)
Normal Cell Sparing	Enhanced due to UHDR	Limited, increased risk to normal cells
Oxygen Effect	Reduced due to ultra-short exposure	Present, potential impact on tumour response
Radiobiological Effect	Increased therapeutic index	Standard radiobiological principles
Fractionation	Single or few fractions possible	Multiple fractions common
Patient Comfort	Reduced overall treatment time	Longer treatment sessions
Machine Wear and Tear	Potentially reduced	Standard wear and tear
Integration with Imaging	Compatibility with advanced imaging	Standard imaging requirements
Organ Motion during Treatment	Reduced impact due to faster delivery if the tumour position is known immediately prior to treatment	Continuous monitoring and adaptation
Patient Throughput	Potentially increased	Treatment duration may impact throughout
Clinical Trial Status	Investigational, ongoing research	Established, widely practiced
Cost and Accessibility	Potential for higher costs	Generally more accessible

Status snapshot: Clinical deployment awaits (a) machines that can reliably deliver clinically useful UHDR beams, (b) mechanistic clarity of the FLASH effect, and (c) validated dosimetry at UHDR where standard detectors face response limitations. These non-computational gaps currently dominate the translation agenda [4, 8].

The exact mechanisms behind the FLASH effect are still being investigated, but several hypotheses exist. One idea is that the extremely short duration of radiation exposure leads to different biological responses in healthy and cancerous tissues, possibly involving oxygen depletion or unique stress pathways in normal cells. While these theories remain under study, the biological promise has been compelling enough to encourage early-stage clinical trials [3].

Yet, bringing FLASH-RT into routine clinical practice is far from straightforward. The technology needed to produce such high dose rates safely and reliably is still under development. This intersection of innovation and complexity has led researchers to explore new computational tools, including QC, to support the future of FLASH-RT.

COMPUTATION'S ROLE

While FLASH-RT holds exciting potential for cancer treatment, its implementation introduces serious computational challenges. Delivering a high dose in just milliseconds leaves very little room for error. To ensure safety and effectiveness, clinicians must plan the treatment with great precision, model how the radiation interacts with tissue, and monitor delivery in real time. Each of these steps depends on intensive calculations rooted in complex medical physics and radiobiology [9].

One of the gold-standard methods for simulating radiation transport and dose distribution is the Monte Carlo method, which uses random sampling to model particle interactions with high accuracy [10]. However, this approach is computationally expensive, it can take hours or even days on classical computers to complete a full simulation. In the context of FLASH-RT, where speed is critical, these calculations remain important but are supporting tasks relative to today's hardware, dosimetry, and biology challenges.

Moreover, real-time treatment adaptation, adjusting the dose or beam in response to patient movement or changing conditions, is an emerging goal in radiotherapy. For FLASH-RT, real-time adaptation becomes even more important, and even more difficult, due to the sheer speed of dose delivery. This pushes current computing systems to their limits. Figure 1 shows key applications of Monte Carlo simulations in FLASH-RT, including dose calculation, beam optimization, biological modeling, and treatment personalization. Monte Carlo is indispensable for UHDR beamline and dosimeter design and for dose transport benchmarking. However, standard clinical Monte Carlo dose engines generally compute time-integrated dose and usually do not explicitly represent the pulsed or microsecond structure of UHDR beams; specialized time-resolved modeling or paired measurements are required when interrogating FLASH-specific dose-rate effects.

An important challenge is performing these calculations efficiently, an area where novel approaches, including QC, may help. This is where new computing paradigms, like QC, are being considered [11]. By rethinking how we process information at the most fundamental level, researchers hope to unlock faster, more efficient ways to meet the computational demands of FLASH-RT.

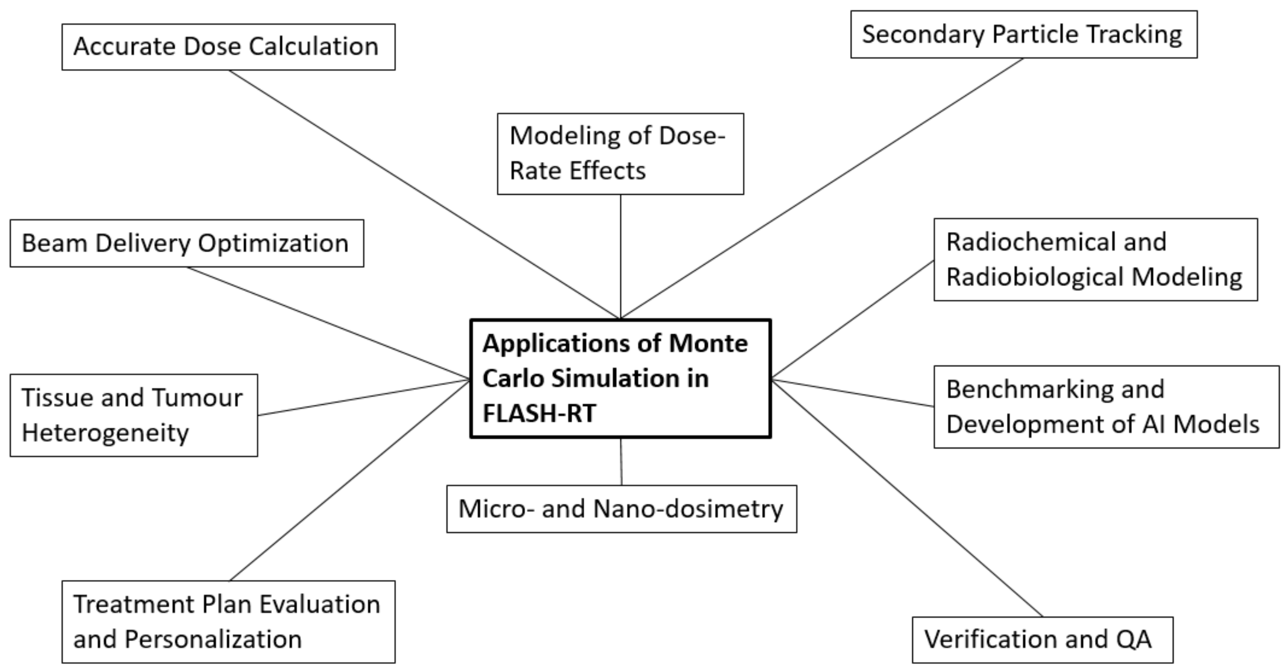

Figure 1. Applications of Monte Carlo simulations in FLASH-RT R&D. Monte Carlo methods support critical tasks such as accurate dose calculation, modeling dose-rate effects, optimizing beam delivery, and enhancing biological understanding, enabling the development and refinement of FLASH-RT techniques. Note: conventional Monte Carlo dose engines are typically time-integrated; micro-temporal beam structure relevant to FLASH is not represented unless custom time-tagging or extensions are used.

WHAT IS QC?

QC is a new way of thinking about computation, one that harnesses the principles of quantum mechanics to solve certain types of problems more efficiently than classical computers. At the heart of a quantum computer are qubits, the quantum counterpart to classical bits. While a classical bit can be either 0 or 1, a qubit can be in a superposition of both states at once. This allows quantum computers to process information in parallel and explore many possibilities simultaneously. Unlike classical parallelism, which distributes tasks across multiple processors, quantum parallelism leverages superposition and entanglement to evaluate many computational paths within a single quantum state, offering a fundamentally different scaling advantage for certain problems.

Another key concept is entanglement, a quantum phenomenon where the state of one qubit is linked to another, no matter how far apart they are. This connection enables qubits to work together in ways that classical bits cannot, potentially giving quantum computers an edge in solving highly complex problems, like optimizing large systems or simulating physical interactions at the atomic level.

Table 2 shows a comparison between classical and quantum computing, highlighting their fundamental differences in data representation, processing, and application. While classical computers use bits and perform deterministic operations, quantum computers use qubits and leverage quantum phenomena like superposition and entanglement to solve certain complex problems more efficiently.

TABLE 2

Comparison between classical and quantum computing. The table highlights key differences in how information is represented and processed, as well as their respective strengths, limitations, and areas of application.

Aspect	Classical Computing	Quantum Computing
Basic Unit of Information	Bit (0 or 1)	Qubit (0, 1, or both simultaneously via superposition)
Data Representation	Binary states	Quantum states (superposition and entanglement)
Processing	Sequential or parallel via multiple cores	Quantum parallelism – can explore many states at once
Computational Power	Scales linearly or polynomially	Can scale exponentially for certain problems
Key Operations	Logic gates (AND, OR, NOT)	Quantum gates (Hadamard, CNOT, etc.) manipulate probabilities
Strengths	General-purpose, reliable, good for most everyday tasks	Ideal for complex optimization, simulation, and factoring problems
Limitations	Slower for problems with vast probabilities (e.g. combinatorics, many-body simulations)	Prone to noise, requires error correction, still early-stage hardware
Memory	Deterministic and easily readable	Probabilistic outcomes that require measurement and repetition
Examples of Applications	Web browsing, spreadsheets, image processing	Drug discovery, cryptography, optimization, quantum physics simulation
Maturity	Fully developed and widely used	Emerging, in research and development phase

QC is still in its early days. Most current devices have a limited number of qubits and are prone to noise and error. But progress is moving quickly, and researchers around the world are exploring how even today's small-scale quantum devices might outperform classical computers in specific tasks.

Canada is actively contributing to the global advancement of quantum science, with leading companies such as D-Wave and Xanadu [12], as well as growing access to quantum hardware through IBM's quantum computing infrastructure in Canada [13]. Research institutions such as the Centre for Quantum Information and Quantum Control and Quantum Software Consortium at the University of Toronto are also playing a key role in exploring practical applications of quantum technologies. This expanding ecosystem positions Canada well to pursue real-world innovations in QC [14], including those in health and medical research. As FLASH-RT stretches the capabilities of classical computing, QC presents a fundamentally new approach that may help meet its unique computational demands.

How Could QC Support FLASH-RT?

Applying QC to FLASH-RT is still a developing idea, but there are several promising directions where it can make a meaningful impact. One key area is dose calculation. In radiotherapy, precise knowledge of how radiation deposits energy in tissue is critical for treatment planning. The most accurate way to simulate this is through Monte Carlo methods, which model millions of particle interactions. However, these simulations are extremely time-consuming on classical computers, especially at the speed and precision needed for FLASH-RT. Quantum algorithms could one day accelerate these simulations by using quantum parallelism to explore many possible particle paths simultaneously [15]. In FLASH- RT, QC should be viewed as a complementary tool. It does not replace the need for validated UHDR dosimetry, FLASH-capable treatment hardware, or mechanistic biology. Its most credible near-term role is to speed up or improve specific calculations within larger experimental and clinical programs.

Another important challenge is treatment plan optimization. Radiation doses must be shaped and delivered to fit the patient's anatomy while avoiding sensitive structures. This involves solving large, complex optimization problems, which are often computationally intensive. QC can offer advantages here through techniques like quantum annealing or variational algorithms, which may find solutions faster or more effectively than classical methods [16].

Moreover, as researchers try to understand the biological effects of FLASH-RT, there is growing interest in quantum machine learning to help analyze large datasets and uncover patterns in biological response. Personalized FLASH-RT will require accurate models of patient-specific anatomy and tissue heterogeneity, which influence scattering and dose deposition. Quantum algorithms, particularly in optimization and machine learning, could help integrate these complex datasets into individualized treatment plans, complementing classical approaches. These tools could contribute to better models for predicting patient outcomes and tailoring treatments [17]. Though practical implementation is still a way off, these examples highlight how QC can help address some of the most pressing computational challenges in FLASH-RT. The goal is not to replace classical computers, but to augment them, offloading the most difficult parts of the problem to quantum processors, and using hybrid approaches that combine the strengths of both.

Figure 2 shows an example of a quantum deep reinforcement learning (qDRL) algorithm designed to support optimal decision-making in knowledge-based adaptive radiotherapy [11]. In this approach, a quantum AI agent interacts with a simulated radiotherapy environment (ARTE) to determine the most effective radiation dose for a given patient state. The agent uses a deep Q-network to evaluate a range of dose options, selects the dose with the highest predicted value (Q-value), and applies quantum amplification to enhance this decision on a quantum state. A quantum measurement is then performed to finalize the dose selection. This decision, along with the patient's current state, is input into ARTE, which models how the patient responds to treatment. ARTE consists of three components: a transition function that simulates state progression, an outcome estimator that predicts treatment effects, and a reward function that scores each decision. The predicted outcomes, including tumour control and risk of side effects, are used to update the agent's learning. This cycle continues until an optimal treatment strategy is learnt, and the process is repeated for new patients. The patient state is defined using five key biological features, including radiomics, radiation sensitivity, cytokine levels, and genetic markers.

Challenges and the Road Ahead

Despite its promise, QC is not yet ready to transform FLASH-RT for radiation treatment, or most other applications, overnight. Today's quantum hardware is still in a phase known as "Noisy Intermediate-Scale Quantum" (NISQ) computing [18]. This means current quantum computers have a limited number of qubits, and those qubits are prone to errors caused by noise and imperfect control. These constraints make it difficult to run large, reliable quantum programs. There are also important practical challenges in applying QC to healthcare problems. Translating a medical physics or radiobiology problem into a form that a quantum computer can process is not straightforward. It requires close collaboration between experts in quantum algorithms, medical physics, oncology, and computer science. Even for researchers, the learning curve can be steep [19].

In addition, clinical validation and safety are paramount in radiotherapy. Before any quantum-enhanced tool could be used in patient care, it would need to go through extensive testing, regulatory approval, and integration with existing clinical workflows—a process that can take years. That said, progress is happening rapidly. As quantum hardware improves and algorithms become more sophisticated, researchers are already exploring hybrid quantum-classical models that use quantum processors for specific tasks within a larger workflow [3]. This "division of labor" may be one of the first practical steps toward bringing quantum computing into radiotherapy research and, eventually, into the clinic. Canada is well-positioned to contribute to this progress. With strong research institutions, access to QC platforms like IBM's, and a growing community of interdisciplinary researchers, the foundations are being laid for innovation at the intersection of quantum science and cancer care.

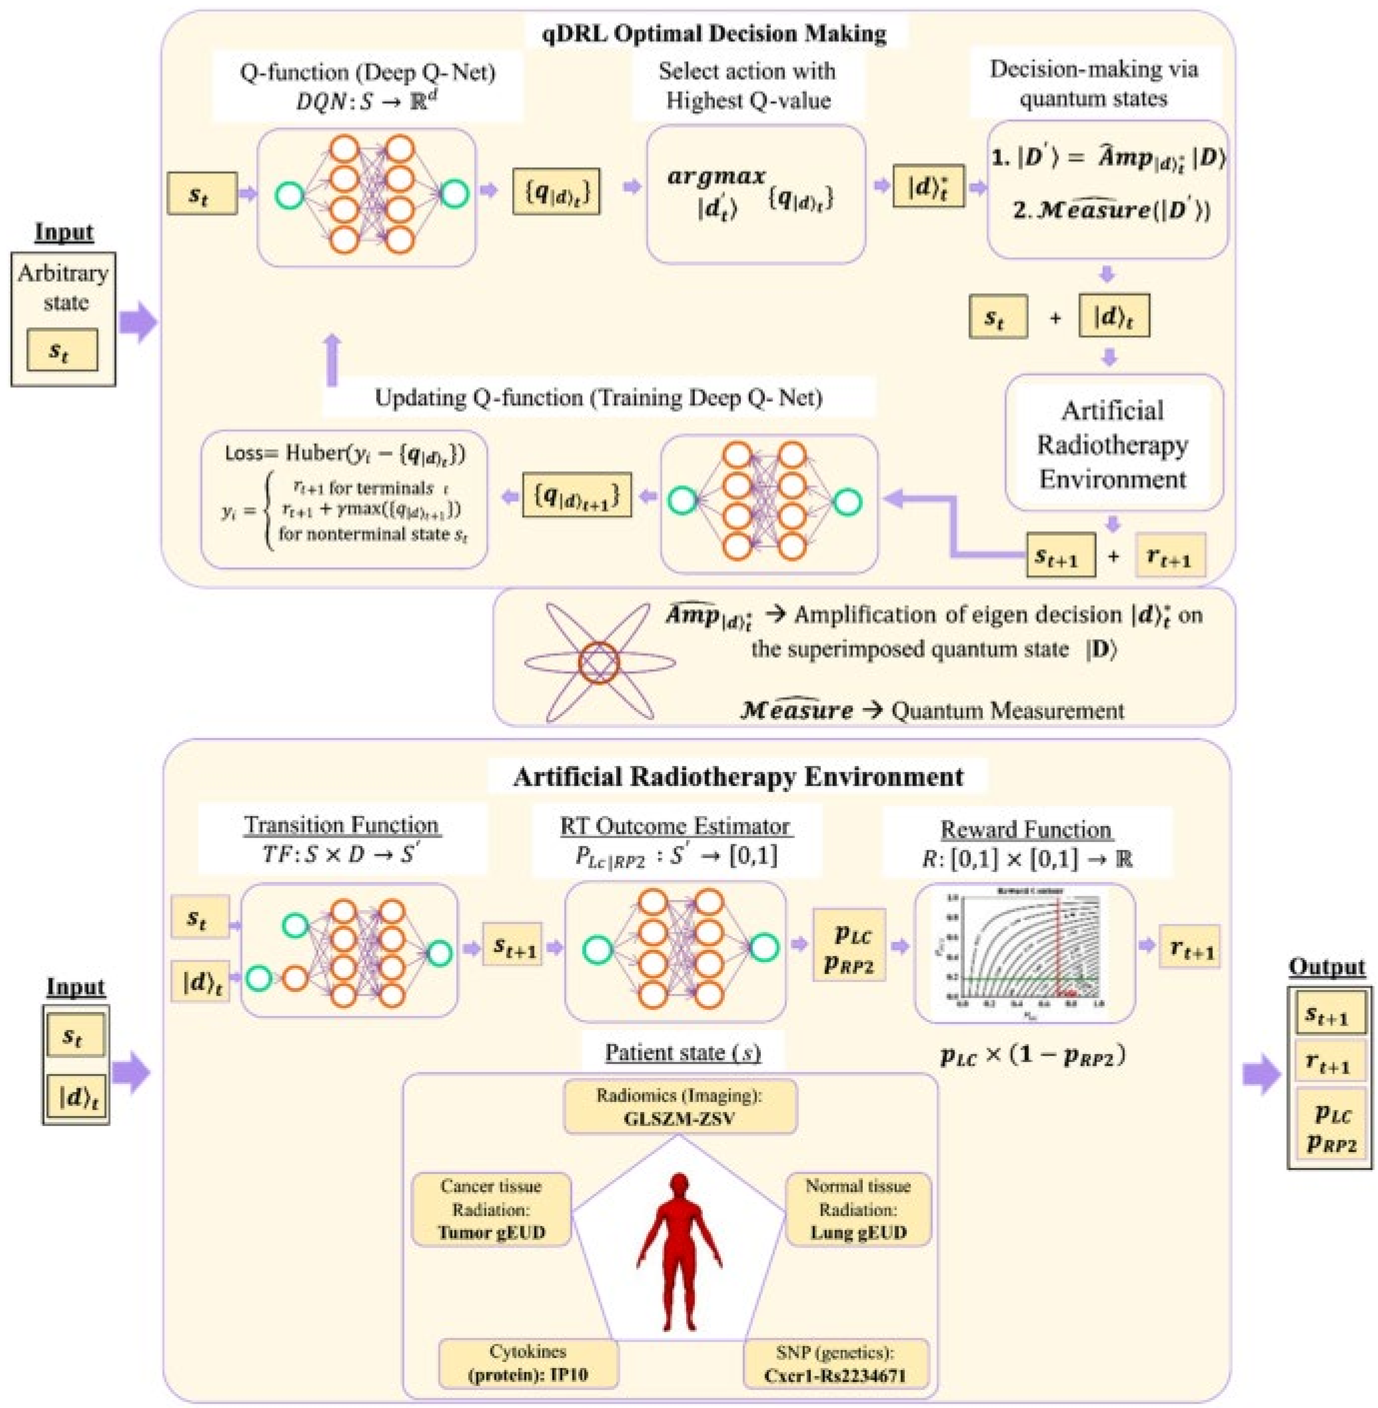

Figure 2. Schematic of a quantum deep reinforcement learning (qDRL) algorithm for adaptive radiotherapy decision-making. The quantum agent interacts with a simulated treatment environment to optimize radiation dosing based on patient-specific biological features [11].

CONCLUSION

FLASH-RT is a promising frontier in cancer treatment, with the potential to deliver highly effective radiation therapy while reducing side effects for patients. To make FLASH-RT practical and widely available, the foremost needs are engineering FLASH-capable clinical systems, establishing accurate and traceable UHDR dosimetry, and resolving key biological questions. Computation including Monte Carlo and, prospectively, quantum approaches will support these efforts by accelerating design, optimization, and analysis.

QC offers a compelling avenue for addressing the unique challenges of FLASH-RT, from speeding up dose calculations to optimizing treatment plans and improving our understanding of biological responses. While much work remains to be done, both in refining quantum technologies and adapting them to medical problems, the potential is there, and the first steps are already being taken. By fostering collaboration between physicists, computer scientists, engineers, and medical professionals, we can begin to explore how this emerging technology might one day support safer, faster, and more personalized cancer treatments. As Canada continues to build its strength in quantum science and technology, initiatives at the intersection of QC and healthcare, like those targeting FLASH-RT, represent an exciting opportunity to lead in both fields.

ACKNOWLEDGEMENTS

This work was supported by the Government of Canada's New Frontiers in Research Fund—Exploration (Grant number: NFRFE-2022-00707), through the three federal research funding agencies (CIHR, NSERC, and SSHRC).

REFERENCES

1. J. Thariat, J. M. Hannoun-Levi, A. Sun Myint, T. Vuong, J. P. Gérard, Past, present, and future of radiotherapy for the benefit of patients, Nat. Rev. Clin. Oncol. **10**, 52-60 (2013).
2. J. C. L. Chow, H. E. Ruda, Flash radiotherapy: Innovative cancer treatment, Encyclopedia **3**, 808-823 (2023).
3. J. C. L. Chow, H. E. Ruda, Mechanisms of action in FLASH radiotherapy: A comprehensive review of physicochemical and biological processes on cancerous and normal cells, Cells **13**, 835 (2024).
4. S. Siddique, H. E. Ruda, J. C. L. Chow, FLASH radiotherapy and the use of radiation dosimeters, Cancers **15**, 3883 (2023).
5. J. C. L. Chow, Quantum computing in medicine, Med. Sci. **12**, 67 (2024).
6. E. Hageman, P. P. Che, M. Dahele, B. J. Slotman, P. Sminia, Radiobiological aspects of FLASH radiotherapy, Biomolecules **12**, 1376 (2022).
7. E. Schüler, M. Acharya, P. Montay-Gruel, B. W. Loo Jr., M. C. Vozenin, P. G. Maxim, Ultra-high dose rate electron beams and the FLASH effect: from preclinical evidence to a new radiotherapy paradigm, Med. Phys. **49**, 2082-2095 (2022).
8. M. C. Vozenin, J. Bourhis, M. Durante, Towards clinical translation of FLASH radiotherapy, Nat. Rev. Clin. Oncol. **19**, 791-803 (2022).

9. J. C. L. Chow, Quantum computing and machine learning in medical decision-making: A comprehensive review, Algorithms **18**, 156 (2025).
10. D. W. O. Rogers, Fifty years of Monte Carlo simulations for medical physics, Phys. Med. Biol. **51**, R287 (2006).
11. D. Niraula, J. Jamaluddin, M. M. Matuszak, R. K. Haken, I. E. Naqa, Quantum deep reinforcement learning for clinical decision support in oncology: application to adaptive radiotherapy, Sci. Rep. **11**, 23545 (2021).
12. F. Hu, B. N. Wang, N. Wang, C. Wang, Quantum machine learning with D-wave quantum computer, Quantum Eng. **1**, e12 (2019).
13. M. AbuGhanem, IBM quantum computers: Evolution, performance, and future directions, J. Supercomput., **81** 687 (2025).
14. B. Sussman, P. Corkum, A. Blais, D. Cory, A. Damascelli, Quantum Canada, Quantum Sci. Technol. **4**, 020503 (2019).
15. G. Mazzola, Quantum computing for chemistry and physics applications from a Monte Carlo perspective, J. Chem. Phys. **160**, (2024).
16. J. M. Pakela, H. H. Tseng, M. M. Matuszak, R. K. Ten Haken, D. L. McShan, I. El Naqa, Quantum-inspired algorithm for radiotherapy planning optimization, Med. Phys. **47**, 5-18 (2020).
17. D. Maheshwari, B. Garcia-Zapirain, D. Sierra-Sosa, Quantum machine learning applications in the biomedical domain: A systematic review, IEEE Access **10**, 80463-80484 (2022).
18. J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).
19. R. Ur Rasool, H. F. Ahmad, W. Rafique, A. Qayyum, J. Qadir and Z. Anwar, Quantum computing for healthcare: A review, Future Internet **15**, 94 (2023).