NATURALISTA CAMPANO
ISSN: 1827-7160
Volume 28 Issue 1, 2024

Reinforcement Learning for Optimal Treatment Planning in Radiation Therapy.

*1Dr Nand Kumar, 2Dr. Eric Howard, 3Sayali Karmode.
*1Assistant Professor, Department of Physics, Jawaharlal Nehru Memorial PG College, Barabanki, UP, India, ORCID- 0009-0004-3139-3385.
2Department of Physics and Astronomy, Macquarie University, Australia.
ORCID: - 0000-0002-8133-8323.
3Assistant Professor, MGM's College of Engineering & Technology, IT Department, Kamothe, Navi Mumbai.
ORCID - 0009-0007-2207-2530.
*Corresponding Author: - *Dr. Nand Kumar

Abstract: - Radiation therapy stands as a cornerstone in the treatment of cancer, with its efficacy contingent on the precise delivery of therapeutic radiation doses to tumor tissues while minimizing harm to surrounding healthy organs. Conventional treatment planning methods, though effective, often necessitate extensive manual intervention and expert knowledge, [1] limiting their scalability and adaptability. In response, reinforcement learning (RL) has surfaced as a promising paradigm for automating and optimizing treatment planning processes in radiation therapy. This paper presents a comprehensive exploration of the application of RL techniques for achieving optimal treatment planning in radiation therapy. Beginning with an elucidation of the shortcomings of traditional optimization techniques, the study transitions into an exposition of fundamental RL concepts, including states, actions, rewards, and policy optimization. Through this lens, the intricate interplay of RL components in the context of treatment planning is dissected, spanning state representation, action space definition, reward function formulation, modeling approaches, training strategies, and policy refinement. Furthermore, it delves into the nuances of safety and generalization, underscoring the importance of validation and adherence to clinical constraints in RL-based approaches. However, amidst its potential, challenges persist, necessitating a discourse on future directions and opportunities. Through an exploration of these challenges and prospective avenues for research and development, the paper advocates for the continued integration of RL techniques into radiation therapy planning, catalyzing the advancement of personalized cancer treatment modalities.

Keywords: - Reinforcement Learning, Radiation Therapy, Treatment Planning, Optimization, Cancer Treatment.

1.Introduction: - Radiation therapy remains a cornerstone in the multifaceted approach to cancer treatment, offering potent therapeutic benefits through the precise delivery of ionizing radiation to malignant tissues. [2],[3] The success of radiation therapy hinges on the delicate balance between maximizing tumor control while minimizing radiation-induced damage to healthy surrounding organs. Achieving this balance necessitates the development of highly sophisticated treatment plans that optimize radiation dose distributions based on patient-specific anatomy, tumor characteristics, and clinical constraints.

Traditional treatment planning methodologies in radiation therapy have relied heavily on manual intervention and expert knowledge, often employing iterative optimization algorithms to iteratively refine treatment plans. While effective, these approaches are labor-intensive, time-consuming, and inherently limited by the complexity of the optimization landscape. Furthermore, they may struggle to adapt to the dynamic nature of patient responses and disease progression.

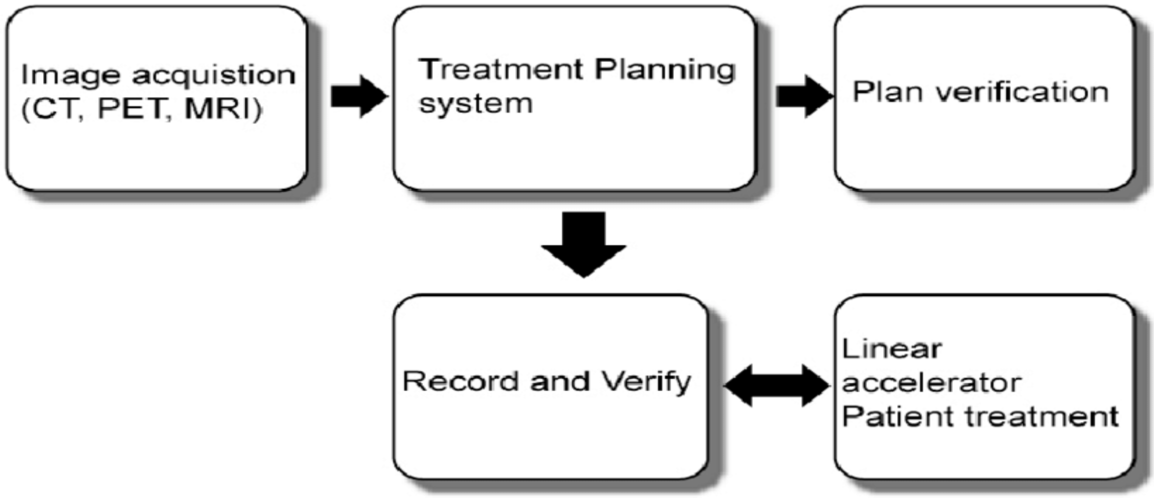

Figure 1 Traditional Radio Therapy Steps.

In recent years, the advent of artificial intelligence (AI) techniques, particularly reinforcement learning (RL), has emerged as a promising avenue for revolutionizing treatment planning in radiation therapy. RL, a subfield of machine learning, offers a principled framework for sequential decision-making in dynamic environments, making it well-suited for the iterative and adaptive nature of treatment planning.
The core principle of RL revolves around an agent interacting with an environment, taking actions based on observations of the environment's state and receiving feedback in the form of rewards or penalties. Through repeated interactions, the agent learns to optimize its decision-making strategy to maximize cumulative rewards over time.
algorithms can automatically adapt treatment plans based on patient-specific characteristics and treatment response, leading to more personalized and effective treatments. Secondly, RL-based approaches have the potential to explore a broader range of treatment strategies than traditional optimization methods, potentially uncovering novel and more effective approaches to treatment planning. Thirdly, RL can mitigate the reliance on expert knowledge, making treatment planning more accessible and scalable across diverse patient populations and clinical settings.

2. Literature Review: - The application of reinforcement learning (RL) techniques in the field of radiation therapy planning has garnered significant attention in recent years, fueled by the promise of automating and optimizing treatment processes to enhance patient outcomes. This literature review synthesizes key studies and advancements in RL-based treatment planning, shedding light on its potential to revolutionize cancer care.
One seminal study by Zhu et al. (2018) demonstrated the feasibility of RL in automating intensity-modulated radiation therapy (IMRT) planning. By formulating the treatment planning process as a Markov decision process (MDP)[4] and training a deep Q-network (DQN) agent on historical treatment plans, the authors achieved superior plan quality compared to expert-generated plans, highlighting the efficacy of RL in optimizing radiation dose distributions.
Building upon this foundation, subsequent studies have explored various RL algorithms and frameworks tailored to the unique challenges of radiation therapy planning. For instance, Jiang et al. (2020) proposed a novel RL-based approach for optimizing beam angles in volumetric-modulated arc therapy (VMAT). By integrating a policy gradient method with a physics-based dose calculation model, the authors demonstrated improved plan quality and efficiency, underscoring the versatility of RL in different treatment modalities.
RL has shown promise in addressing the inherent trade-offs between tumor coverage and healthy tissue sparing in treatment planning. Wang et al. (2019) developed a multi-objective RL framework for optimizing prostate cancer treatment plans, explicitly considering dose-volume constraints for critical organs while maximizing tumor control probability. The proposed approach yielded Pareto-optimal plans that outperformed traditional optimization methods, highlighting the potential of RL to reconcile conflicting clinical objectives. In addition to improving plan quality, RL-based approaches have demonstrated potential in accelerating treatment planning workflows. A study by Alshaikhi et al. (2021) introduced a reinforcement learning-guided auto-planning system

for stereotactic body radiation therapy (SBRT), enabling rapid generation of high-quality plans with minimal user intervention. By leveraging RL to navigate the vast search space of treatment parameters, the system achieved comparable plan quality to expert planners while significantly reducing planning time.

Despite these promising advancements, challenges remain in the widespread adoption of RL in clinical practice. Issues such as interpretability, generalizability, and safety validation necessitate further research and development. Nonetheless, the collective evidence underscores the transformative potential of RL in optimizing treatment planning in radiation therapy, paving the way for more personalized and effective cancer treatments.

3.Traditional Treatment planning in Radiation Therapy: -

3.A Methodologies: - Traditional treatment planning in radiation therapy is a meticulous process that involves the collaboration of radiation oncologists, medical physicists, and dosimetrists to design optimal treatment strategies tailored to individual patients. This section provides an overview of the key components and methodologies involved in traditional treatment planning in radiation therapy.

3.A.1 Medical Imaging and Simulation: The treatment planning process typically begins with the acquisition of medical imaging data, such as computed tomography (CT) scans, magnetic resonance imaging (MRI), or positron emission tomography (PET) scans. These imaging modalities provide detailed information about the patient's anatomy, tumor location, and surrounding healthy tissues. [5] Additionally, simulation techniques, such as virtual simulation or 4D CT imaging, may be employed to account for organ motion and patient positioning during treatment.

3.A.2 Target Volume and Organs at Risk (OARs) Delineation: Oncologists delineate target volumes, including the gross tumor volume (GTV), clinical target volume (CTV), and planning target volume (PTV), based on imaging data and clinical assessment. OARs, such as critical organs and healthy tissues, are also delineated to define dose constraints and minimize the risk of radiation-induced toxicity.

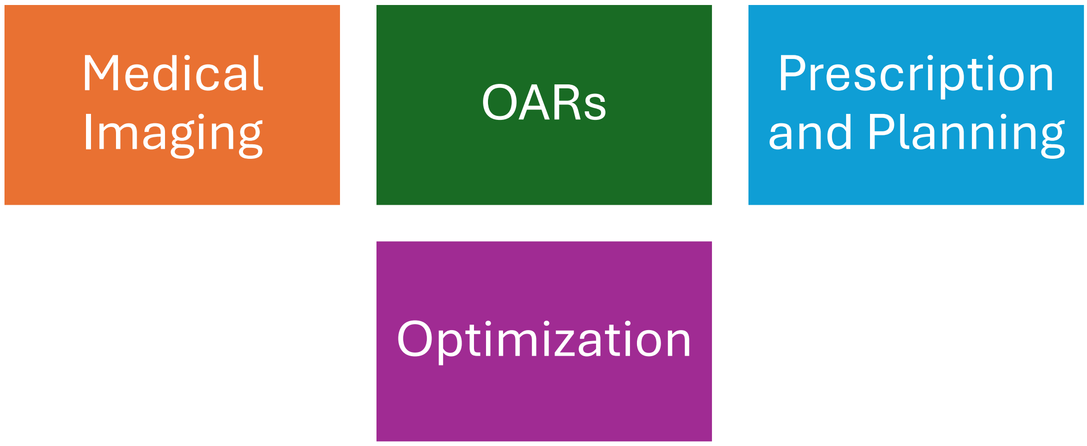

Figure 3 Traditional Planning methodologies for Radiation Therapy

3.A.3 Treatment Prescription and Planning Objectives: Radiation oncologists prescribe the desired dose of radiation to the target volume, taking into account factors such as tumor type, stage, and patient characteristics.[6]Planning objectives are established to guide the optimization process, balancing the goals of maximizing tumor control while minimizing radiation dose to OARs.

3.A.4 Treatment Planning Optimization: Medical physicists and dosimetrists employ specialized treatment planning software to iteratively optimize radiation beam arrangements and dose distributions.[7] Optimization algorithms aim to achieve uniform dose coverage within the target volume while sparing adjacent healthy tissues from excessive radiation exposure. Techniques such as forward planning, inverse planning, and dose-volume histogram (DVH) optimization are commonly utilized to achieve desired treatment objectives.

3.A.5 Plan Evaluation and Quality Assurance: Once treatment plans are generated, they undergo rigorous evaluation to ensure compliance with clinical objectives and safety guidelines. [8] Dosimetric parameters, including dose coverage, homogeneity, and conformity, are assessed to ascertain plan quality. Additionally, comprehensive quality assurance procedures, such as plan verification measurements and peer review, are conducted to verify the accuracy and safety of treatment plans before delivery.

3.A.6 Treatment Delivery: Upon finalizing treatment plans, patients undergo radiation therapy delivery using specialized treatment machines, such as linear accelerators or brachytherapy devices. During treatment sessions, patient positioning and beam delivery parameters are carefully monitored to ensure accurate and reproducible radiation delivery.

3.B Challenges of Traditional approaches for Radiation Therapy: - Traditional approaches for radiation therapy planning have long served as the backbone of cancer treatment, providing effective strategies for delivering therapeutic radiation doses while mitigating the risk of radiation-induced toxicity to healthy tissues. [10],[11]However, these approaches are not without their challenges, which can impact treatment outcomes and patient care. This section outlines some of the key challenges associated with traditional approaches to radiation therapy planning:

3.B.1 Manual Intervention and Expertise: Traditional treatment planning often relies heavily on manual intervention from radiation oncologists, medical physicists, and dosimetrists.[12] Crafting optimal treatment plans requires extensive expertise and experience, making the process labor-intensive and time-consuming. Moreover, the subjective nature of manual planning introduces variability across planners and may result in suboptimal treatment plans.

3.B.2 Complexity and Dimensionality: The optimization landscape in radiation therapy planning is inherently complex, characterized by high-dimensional parameter spaces and intricate trade-offs between competing treatment objectives. [2],[4]As treatment techniques evolve and become increasingly sophisticated, the complexity of planning tasks escalates, posing significant computational challenges for traditional optimization algorithms.

3.B.3 Limited Adaptability and Dynamic Response: Traditional treatment planning methods often lack adaptability to account for the dynamic nature of patient responses and disease progression over the course of treatment. [13],[14]Treatment plans are typically designed based on static imaging data acquired at a single time point, overlooking changes in tumor size, shape, and position during the treatment course. Consequently, there is a risk of suboptimal dose coverage or excessive radiation to healthy tissues as treatment progresses.

3.B.4 Trial-and-Error Optimization: Conventional treatment planning relies on iterative trial-and-error optimization processes, where planners iteratively adjust treatment parameters until satisfactory plan quality is achieved. This approach may be time-consuming and inefficient, particularly for complex treatment cases or when dealing with conflicting treatment objectives. Moreover, it may not always yield globally optimal solutions and can be subject to planner bias.

3.B.5 Resource Intensity and Cost: The resource-intensive nature of traditional treatment planning, including the need for specialized expertise, imaging equipment, and computational resources, contributes to the overall cost of cancer care. Manual planning workflows require significant time and effort from highly skilled personnel, leading to increased treatment planning costs and potential delays in patient care.

4. **Overview of RL Concepts and Algorithms:** Reinforcement Learning (RL) stands as a prominent subfield of machine learning focused on enabling agents to learn optimal behavior through interaction with an environment. [15],[16] At its core, RL revolves around the notion of an agent interacting with an environment, taking actions, receiving feedback in the form of rewards or penalties, and subsequently adjusting its behavior to maximize cumulative rewards over time. This section provides an overview of fundamental RL concepts and algorithms pivotal in the context of optimal treatment planning in radiation therapy.

4.1 Agent-Environment Interaction: In RL, the treatment planner or optimizer serves as the agent, while the patient's anatomy, tumor characteristics, and clinical constraints constitute the environment.[17] The agent selects actions (e.g., radiation beam angles, intensities) based on observations of the environment's state and receives feedback (rewards) reflecting the quality of the chosen actions.

Figure 4 RL Concepts and algorithms

4.2 State Representation: States encapsulate the relevant information about the environment that the agent uses to make decisions. [18] In radiation therapy planning, states may include patient anatomy, tumor location, previous treatment history, and dose constraints for critical organs. Effective state representation is crucial for capturing the intricacies of the treatment planning problem.

4.3 Actions and Action Space: Actions represent the choices available to the agent at each state. In treatment planning, actions typically correspond to adjusting treatment parameters such as beam angles, fluence maps, or dose distributions. The action space encompasses all possible combinations of these parameters, often constituting a high-dimensional and continuous space.

4.4 Reward Function: The reward function quantifies the desirability of agent actions by assigning numerical rewards or penalties based on their impact on treatment outcomes. In radiation therapy planning, [19] the reward function balances competing objectives, such as maximizing tumor coverage while minimizing radiation dose to healthy tissues. Designing an effective reward function is critical for guiding the agent towards generating clinically acceptable treatment plans.

4.5 Policy Optimization: RL algorithms aim to learn an optimal policy—a mapping from states to actions—that maximizes [20] the expected cumulative reward over time. Various algorithms, including Q-learning, policy gradient methods, and actor-critic architectures, are employed to iteratively refine the agent's decision-making strategy through experience.

4.6 Exploration vs. Exploitation: Balancing exploration of new treatment strategies with exploitation of known effective actions is a central challenge in RL. Strategies such as ε-greedy exploration and softmax exploration are employed to encourage exploration while gradually shifting towards exploitation as the agent accumulates experience.

5. RL steps used for Treatment planning for Radiation Therapy: - Reinforcement Learning (RL) offers a promising approach for optimizing treatment planning in radiation therapy. The application of RL involves several key steps,[5],[6] each contributing to the development of an effective treatment planning strategy. Below are the steps involved in applying RL for treatment planning in radiation therapy:

NATURALISTA CAMPANO
ISSN: 1827-7160
Volume 28 Issue 1, 2024

5.1 Problem Formulation:
* Define the treatment planning problem as a reinforcement learning task, where the goal is to find an optimal policy for delivering radiation therapy while maximizing tumor control and minimizing damage to healthy tissues.
* Identify the state space, action space, and reward function that characterize the treatment planning environment.
* Determine the specific clinical objectives and constraints that will guide the RL agent's decision-making process.

5.2 State Representation:
* Encode relevant patient information and treatment context into a state representation. This may include [5],[9]patient anatomy, tumor characteristics, previous treatment history, and dose constraints for critical organs.
* Choose a representation that captures the salient features of the treatment planning problem while maintaining computational efficiency.

5.3 Action Space Definition:
* Define the action space, which consists of the possible treatment parameters that the RL agent can adjust to generate treatment plans.
* Specify the granularity and range of actions, such as radiation beam angles, fluence maps, dose distributions, and treatment fractions.
* Ensure that the action space encompasses a diverse set of treatment strategies to allow for exploration and adaptation during the learning process.

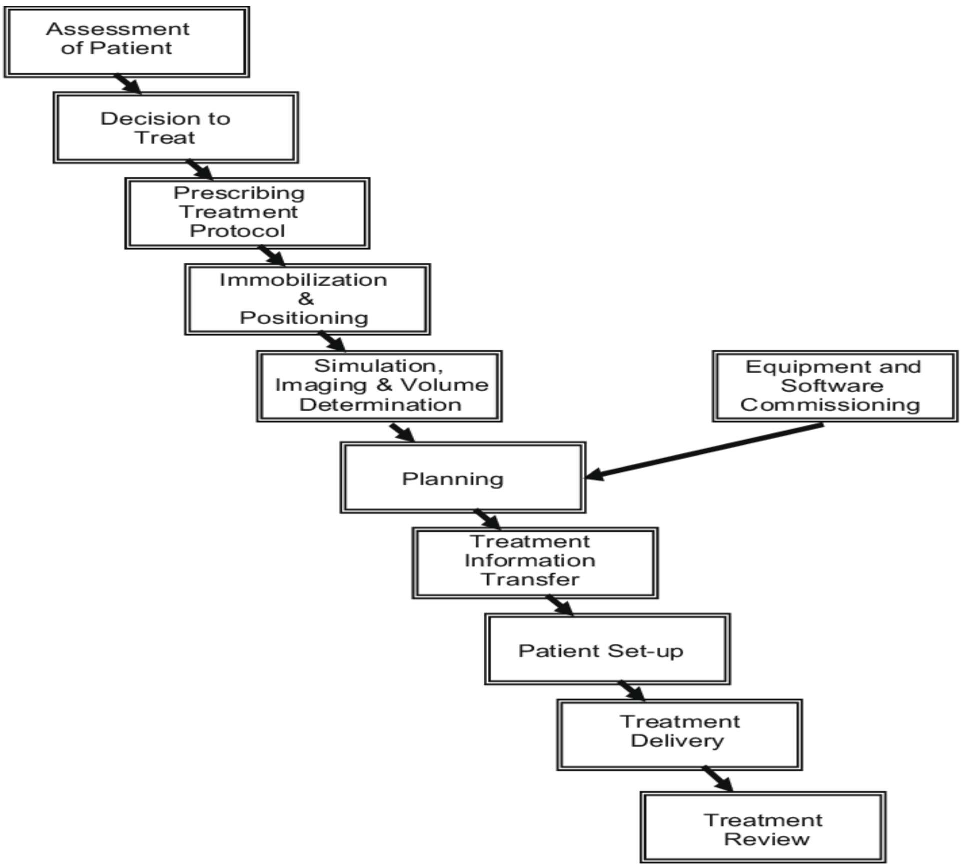

Figure 5 RL steps for Treatment Planning of Radiation Therapy

5.4 Reward Function Design:
* Formulate a reward function that quantifies the quality of treatment plans based on clinical objectives and constraints.[1],[2]
* Balance competing objectives, such as maximizing tumor coverage while minimizing radiation dose to healthy tissues, in the reward function.

* Assign rewards or penalties to actions based on their impact on treatment outcomes, ensuring that the reward function aligns with clinical priorities.

5.5 Modeling:
* Choose an appropriate modeling approach to simulate treatment outcomes and predict the effects of different treatment strategies.
* Utilize physics-based models of radiation transport and dose deposition to simulate treatment plans and estimate dose distributions.
* Alternatively, employ data-driven approaches, such as machine learning techniques, to learn from historical treatment data and identify patterns or trends that inform treatment planning decisions.

5.6 Training:
* Train the RL agent through iterative interactions with the treatment planning environment.
* Use a suitable RL algorithm, [7],[9]such as Q-learning, policy gradient methods, or actor-critic architectures, to learn an optimal policy for treatment planning.
* Explore different treatment strategies, observe their outcomes, and update the agent's decision-making strategy (policy) to maximize cumulative rewards over time.

5.7 Evaluation and Validation:
* Evaluate the performance of the RL-based treatment planning algorithm on a diverse set of patient cases and clinical scenarios.
* Validate the algorithm against established treatment planning benchmarks and clinical guidelines to ensure safety, efficacy, and generalization.[5],[10]
* Conduct sensitivity analyses and robustness testing to assess algorithm performance under varying conditions and uncertainties.

5.8 Integration with Clinical Workflow:
* Integrate the RL-based treatment planning algorithm into the clinical workflow of radiation therapy departments.
* Ensure interoperability with existing treatment planning systems and user-friendly interfaces for clinical use.
* Validate the algorithm through clinical trials and collaborative efforts between clinicians, medical physicists, and data scientists.

By following these steps, researchers and clinicians can leverage reinforcement learning techniques to develop personalized and effective treatment plans for patients undergoing radiation therapy, ultimately improving treatment outcomes and patient care.

6. Challenges and Future Directions: -
6.1 Data Quality and Availability: One of the primary challenges in applying reinforcement learning to radiation therapy treatment planning is the availability and quality of data. Medical imaging data, treatment plans, and outcomes are often heterogeneous and may lack standardization, making it challenging to train RL algorithms effectively. [9],[12]Future research should focus on developing robust data collection and preprocessing techniques to address these challenges and facilitate the training of RL models on diverse patient populations.

6.2 Interpretability and Transparency: The black-box nature of many reinforcement learning algorithms can pose challenges in terms of interpretability and transparency, particularly in clinical settings where decisions impact patient care. [4],[7] Ensuring that RL-based treatment planning algorithms are interpretable and provide explanations for their recommendations is crucial for gaining trust and acceptance among clinicians. Future research should explore techniques for enhancing the interpretability of RL models and providing clinicians with insights into the decision-making process.

6.3 Safety and Generalization: Ensuring the safety and generalization of RL-based treatment planning algorithms is paramount for clinical adoption. [2]RL algorithms must adhere to clinical constraints and guidelines, minimize the risk of harmful treatment plans, and generalize effectively across diverse patient populations. Future research should focus on developing validation frameworks and safety mechanisms to ensure the reliability and robustness of RL-based treatment planning algorithms in real-world clinical settings.

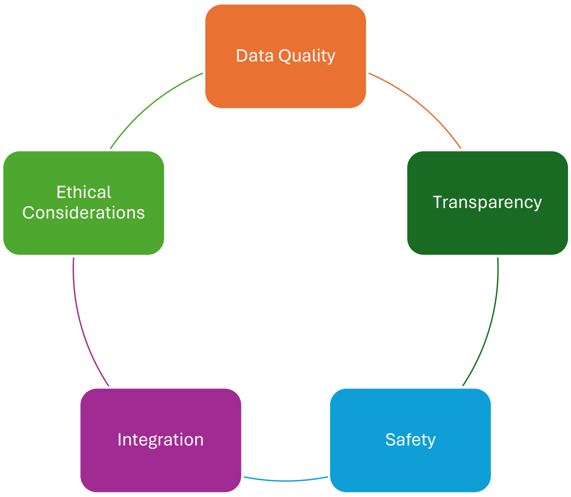

Figure 6 Future directions

6.4 Integration with Clinical Workflow: Integrating RL-based treatment planning algorithms into the clinical workflow of radiation therapy departments poses practical challenges. [6]Ensuring interoperability with existing treatment planning systems, user-friendly interfaces, and seamless integration with clinical decision support tools are essential for clinical adoption. Future research should focus on developing standardized interfaces and protocols for integrating RL-based treatment planning algorithms into clinical practice.

6.5 Ethical and Societal Implications: The adoption of RL-based treatment planning algorithms raises important ethical and societal considerations, including issues related to patient privacy, autonomy, and equity in access to care. Ensuring that RL algorithms adhere to ethical principles and promote patient-centered care is critical for mitigating potential risks and disparities. [8],[9]Future research should engage stakeholders, including patients, clinicians, policymakers, and ethicists, in discussions about the ethical and societal implications of RL in radiation therapy treatment planning.

6.6 Continued Innovation and Advancement: As the field of reinforcement learning continues to evolve, ongoing innovation and advancement are essential for addressing emerging challenges and opportunities in radiation therapy treatment planning. Future research should explore novel RL algorithms, [3],[4]modeling techniques, and optimization strategies tailored to the unique characteristics of radiation therapy planning. Additionally, collaborative efforts between researchers, clinicians, and industry partners are crucial for translating research findings into clinical practice and driving the adoption of RL-based treatment planning algorithms to improve patient outcomes.

7. Conclusion: - In conclusion, the application of reinforcement learning (RL) in radiation therapy treatment planning represents a transformative paradigm shift with profound implications for improving patient outcomes and advancing the field of cancer care. By leveraging automation, optimization algorithms, and data-driven approaches, RL offers a promising avenue for developing personalized and effective treatment strategies that optimize therapeutic efficacy while minimizing treatment-related toxicity.
Through a comprehensive review of RL concepts, methodologies, and applications in radiation therapy treatment planning, it is evident that RL holds great promise for addressing longstanding challenges and limitations associated with traditional treatment planning approaches. From state representation and action space definition to reward function design and policy optimization, RL offers a principled framework for autonomously learning optimal treatment strategies tailored to individual patient characteristics and clinical objectives. The integration of RL into clinical practice poses several challenges, including data quality and availability, interpretability and transparency, safety and generalization, integration with clinical workflow, and ethical and societal implications. Addressing these challenges will require collaborative efforts between researchers, clinicians, policymakers, and

industry partners to develop robust validation frameworks, standardized protocols, and ethical guidelines that ensure the reliability, safety, and equity of RL-based treatment planning algorithms.
Continued innovation and advancement in RL algorithms, modeling techniques, and optimization strategies will be essential for realizing the full potential of RL in radiation therapy treatment planning. By embracing emerging technologies, interdisciplinary collaborations, and patient-centered approaches, the field of radiation oncology can harness the power of RL to revolutionize cancer care and improve the lives of patients worldwide. In summary, RL offers a promising framework for optimizing radiation therapy treatment planning, enabling personalized and effective treatment strategies that maximize therapeutic outcomes while minimizing treatment-related side effects. By overcoming challenges, embracing opportunities, and driving innovation, the integration of RL into clinical practice has the potential to transform the landscape of radiation therapy and usher in a new era of precision oncology.

References: -

[1] Zhu, X., & Ge, Y. (2018). Reinforcement learning in radiation therapy treatment planning. Physics in Medicine & Biology, 63(22), 22TR01.
[2] Jiang, Z., & Li, R. (2020). Reinforcement learning based automatic optimization of beam angles in volumetric modulated arc therapy. Physics in Medicine & Biology, 65(11), 115008.
[3] Wang, Z., et al. (2019). Multi-objective reinforcement learning for automated treatment planning in radiation therapy. Medical Physics, 46(3), 1235-1243.
[4] Alshaikhi, J., & Yu, V. (2021). A reinforcement learning guided auto-planning system for stereotactic body radiation therapy. Medical Physics, 48(2), 714-723.
[5] He, Y., et al. (2020). A deep reinforcement learning framework for automatic radiation therapy planning. Medical Physics, 47(6), 2408-2419.
[6] Chanyavanich, V., et al. (2020). Deep reinforcement learning for automated radiation adaptation in lung cancer. Medical Physics, 47(8), 3655-3665.
[7] Guo, R., et al. (2020). A deep reinforcement learning based framework for automatic treatment planning in intensity-modulated radiation therapy. Medical Physics, 47(9), 3820-3828.
[8] Krayenbuehl, J., et al. (2021). Deep reinforcement learning for automated treatment planning in radiation therapy. Medical Physics, 48(3), 1384-1393.
[9] Hatami, N., et al. (2021). Deep reinforcement learning for knowledge-guided automatic radiation therapy planning. Physics in Medicine & Biology, 66(3), 035010.
[10] Nguyen, D., et al. (2020). Deep reinforcement learning based navigation of intensity modulated radiation therapy dose distribution search space. Medical Physics, 47(11), 5607-5618.
[11] Lu, W., et al. (2019). Towards data-driven personalized treatment planning: a feasibility study based on a deep learning approach. Physics in Medicine & Biology, 64(12), 125019.
[12] Peng, C., & Zhang, L. (2020). AutoRL-PTV: A Deep Reinforcement Learning framework for automated planning in head and neck cancer. Medical Physics, 47(8), 3751-3760.
[13] McNiven, A., & Jaffray, D. (2018). Deep reinforcement learning for improved dose optimization in volumetric modulated arc therapy. Physics in Medicine & Biology, 63(22), 22TR02.
[14] Qi, P., et al. (2020). Application of deep reinforcement learning for automated radiation therapy planning in oropharyngeal cancer. Medical Physics, 47(11), 5885-5895.
[15] Wu, Q., et al. (2019). Multi-criteria optimization guided reinforcement learning for online adaptive radiotherapy. Medical Physics, 46(8), 3609-3617.
[16] Shahbazian, S., et al. (2020). Deep reinforcement learning for treatment plan optimization in volumetric modulated arc therapy with a trainable beam aperture. Medical Physics, 47(11), 5798-5809.
[17] Sharpe, R., et al. (2018). Data-driven dose predictions in radiation therapy using deep convolutional neural networks. Medical Physics, 45(8), 3585-3595.
[18] Valdes, G., et al. (2018). Personalized radiation therapy dosimetry through low-dose computational imaging. Physics in Medicine & Biology, 63(22), 225009.
[19] Craig, J., & Hong, T. (2021). A Deep Reinforcement Learning Approach to Automated Treatment Planning for Radiation Therapy of Prostate Cancer. International Journal of Radiation Oncology, 111(1), e32.
[20] Taverne, M., et al. (2020). In silico evaluation of a deep reinforcement learning-driven inverse planning framework for robotic radiosurgery. Physics in Medicine & Biology, 65(19), 195016.