Enhancing Mathematical Optimization in Intensity-Modulated Radiation Therapy with Artificial Intelligence and Machine Learning

Sakshi Taaresh Khanna[1], Sunil Kumar Khatri[2], and Neeraj Kumar Sharma[3*]

[1]Amity Institute of Information Technology, Amity University, Amity Road, Sector 125, Noida – 201301, Uttar Pradesh, INDIA
Corresponding author e-mail: sakshitaareshkhanna@gmail.com
[2]Amity University, Amity Road, Sector 125, Noida – 201301, Uttar Pradesh, INDIA
e-mail: skkhatri@amity.edu
[3*]Department of Computer Science, Ram Lal Anand College, Benito Juarez Marg, New Delhi – 110021, INDIA
Corresponding author e-mail: neerajksharma100@gmail.com

KEYWORDS
Adaptive Therapy, Artificial Intelligence, Beam Angle Optimization, Computational Efficiency, Deep Learning, IMRT, Machine Learning, Multi-Objective Optimization

ABSTRACT:
A key technique used in cancer treatment is called Intensity-Modulated Radiation Therapy (IMRT), which requires accurate delivery of doses to target the tumors and reduce exposure to healthy tissues. Mathematically, optimal beam angles to maximize objectives are introduced as decision variables and the objective for an IMRT plan is used both linear or nonlinear functions of fluence intensities. Yet these techniques struggle with multi-objective optimization in complex settings and adapting to clinical data that is collected in real-time. In this paper, a novel AI-ML framework is introduced in which mathematical optimization integrated with the ethical dimensions of Artificial Intelligence (AI) and Machine Learning (ML)-enabled optimization leads to better treatment results for IMRT. The method uses reinforcement learning (RL) for adaptive optimization of the dose based on actual patient feedback, a deep-learning approach to predictive beam angle selection, and GA-based multi-objective optimization. Primary resulting metrics include tumour coverage, organ-at-risk (OAR) sparing, treatment planning time, and computational speed. The augmented framework integrates direct patient-specific information such as tumour geometry and biological markers to assist the treatment plan with a personalized approach, thereby improving precision and efficiency. We show experimentally that our method is considerably faster, with a slight improvement in dose distribution conformity and tumour coverage while causing less damage to normal tissues. Their AI-ML enhanced optimization framework proposes a revolutionary solution for IMRT, circumventing the disadvantageous feature of traditional methods and paving the way to more efficient personal cancer therapies. These results suggest that AI-based methods could help transform how radiation therapy is performed for cancer patients.

1. Introduction

IMRT is an advanced radiation therapy technique to treat cancers, which uses a specialized computerized planning system (Cad/cam) along with linear accelerator schematics for delivering exact radiation dosages of malignant tumours while decreasing radiations on the surrounding uninvolved tissues. IMRT delivers radiation at different intensities, making it extremely accurate for more complex tumour shapes and areas [1]. The irradiation dose and beam angles need to be optimally adjusted for IMRT as those are the main configuration parameters of radiation which heavily affect treatment results including patient safety. Most of the traditional mathematical optimization techniques have been implemented to solve this multi-objective problem, considering that it is essential to achieve a maximum tumour coverage and minimum OAR exposure. However, traditional optimization approaches are limited in handling the complexity and variability of clinical scenarios that arise in IMRT. Most existing methods for discovering optimal trades exploit continuous approximations of the binary constraints, which can be computationally expensive and where it is challenging to find global optima in multi-objective non-convex problems [2]. In addition, these conventional methods do not allow for real-time adaption preventing on-the-fly corrections during treatment plans challenging patient anatomy or imaging feedback in a transparent timely fashion. Consequently, there is an increasing requirement for more efficient and flexible optimization techniques that can accommodate these complexities to enhance treatment quality and reduce planning time.

Some notable advances that may mitigate these limitations in radiation oncology include the introduction of Artificial Intelligence (AI) and Machine Learning (ML). Artificial intelligence (AI) models, especially deep learning and reinforcement learning allow for treatment plans to be dynamically optimized by modelling against enormous datasets so that they can learn from patient-specific variables [3]. Reinforcement Learning, for instance, can tune radiation doses per treatment in real-time and feedback from the patient within the same application. Likewise, power calculations deep learning models can be used to predict complex tumour geometries beam angles and, in this way, decrease manual adjustment requirements, making dose delivery more accurate. The eventual objective is to incorporate AI and/or ML towards IMRT optimization, producing this more precise plan faster than currently achievable in a patient-specific nature. Multi-objective optimization can be achieved more efficiently by using powerful AI algorithms, which are much faster in balancing tumour coverage and OAR sparing while decreasing the planning time as well as its cost [4]. Here we propose an intriguing framework which is AI-ML integrated and brings to the table these techniques along with conventional mathematical optimization of IMRT. The results of case studies and experiments prove its efficiency in treatment planning, computational performance and improvement in the patient outcome compared to other frameworks.

2. Literature Review

Mathematical optimization has a long history in radiation therapy, particularly for Intensity intensity-modulated radiation Therapy (IMRT); Conventional techniques like linear and nonlinear programming have been widely used to determine beam angles and dose distributions which deliver high doses within a target while sparing normal organs at risk, many with variable success. Studies [5]. Consequently, it was possible to perform narrow pencil beam irradiation with variable dose distribution [6] adding significantly higher accuracy of radiation therapy. But while they are powerful these algorithms aren't robust enough to perform real-time adaptation or multi-objective optimization in clinical scenarios.

AI & ML in radiation oncology has revolutionized IMRT optimization. [7] demonstrated the capacity for AI to dynamically optimize treatment planning by learning from large datasets and incorporating patient-specific factors. By fine-tuning the treatment plan in line with patient feedback, reinforcement learning can perform real-time adjustments to radiation doses. The live interactive capacity for dosing is a major improvement from standard optimization procedures, allowing treatment planning in real-time. The utilization of deep learning in beam angle optimization and treatment planning has also gained considerable interest. Tseng et al. Deep learning was applied to IMRT treatment planning for nasopharyngeal carcinoma, where convolutional neural networks have been shown to predict the optimal beam angles based on tumour geometry [8]. The application of deep learning models for IMRT has replaced manual work, making the process more accurate and much quicker to calculate treatment plans. This has not only provided the opportunity to improve the accuracy in the delivery of doses but particularly in complex tumour geometries.

The use of genetic algorithms in multi-objective optimization as well for IMRT has also been successful. [9] explored the use of genetic algorithms to trade-off beam angles against conflicting objectives: how do you deliver enough dose at an optimal proportion in solid tumours, while sparing OAR as best possible? Genetic algorithms guide the system towards a proper solution space making it possible to find better compromises between conflicting goals. This approach gives a potential solution to one of the dilemmas in IMRT multi-objective optimization.

3. PROPOSED ALGORITHM

In this section, we proposed the AI-ML Augmented Mathematical Optimization in IMRT is proposed as follows:

Algorithm: AI-ML Augmented Mathematical Optimization in IMRT

Input: Patient-specific data $\mathcal{P}$, Tumour geometry $\mathcal{T}$, Organs-at-risk (OAR) data $\mathcal{O}$.
Output: Optimized beam angles θ^*, Optimized dose distribution D^*.

Step 1: Initialization

Initialize beam angles θ_0 and dose distribution D_0 using conventional optimization methods:

$$\theta_0, D_0 = \arg\min_{\theta, D}\left[\sum_{i=1}^{N} \lambda_i \cdot f_i(\theta, D)\right] \tag{1}$$

where $f_i(\theta, D)$ is the objective function for the i-th criterion, λ_i represents the weight of the criterion, and N is the number of objectives. The objectives may include:

$$f_{\text{Tumour Coverage}} = -\sum_{v \in \mathcal{T}} D(v) \tag{2}$$

$$f_{\text{OAR Sparing}} = \sum_{v \in \mathcal{O}} D(v) \tag{3}$$

where $D(v)$ is the dose at voxel $v, \mathcal{T}$ is the set of voxels in the tumour, and $\mathcal{O}$ is the set of voxels in the OAR.

Step 2: Dose Calculation Based on Beam Angles

The total dose $D(v)$ delivered to voxel v can be computed as the sum over all beams:

$$D(v) = \sum_{k=1}^{M} I_k \cdot \phi_k(v, \theta_k) \tag{4}$$

where I_k is the intensity of the k-th beam, $\phi_k(v, \theta_k)$ is the fluence at voxel v for beam angle θ_k, and M is the number of beams.

Step 3: Deep Learning for Beam Angle Optimization

Train a CNN to predict optimal beam angles based on tumour geometry:

$$\theta_{DL} = \text{CNN}(\mathcal{T}) \tag{5}$$

where $\mathcal{T}$ is the input tumour geometry and θ_{DL} are the predicted optimal angles.

Step 4: Reinforcement Learning for Dose Adjustment

Define a reward function $R(D)$ to balance tumour coverage and OAR sparing:

$$R(D) = \frac{\sum_{v \in \mathcal{T}} D(v)}{\sum_{v \in \mathcal{O}} D(v)} \tag{6}$$

The dose distribution is updated at each iteration using a policy $\pi(D_t)$:

$$D_{t+1} = D_t + \alpha \cdot \nabla R(D_t) \tag{7}$$

where α is the learning rate.

Step 5: Multi-Objective Optimization with Genetic Algorithms
Use a genetic algorithm to optimize beam angles and dose distributions simultaneously. Define the total optimization problem:

$$\max_{\theta,D}[F_{\text{Tumour}}(\theta,D)-F_{OAR}(\theta,D)] \tag{8}$$

where

$$F_{\text{Tumour}}(\theta,D)=\int_{\mathcal{T}} D(v)dv \tag{9}$$

$$F_{OAR}(\theta,D)=\int_{\mathcal{O}} D(v)dv \tag{10}$$

The AI-ML Augmented IMRT Optimization Algorithm integrates conventional mathematical methods with modern machine learning technology for optimization of beam angles and dose distribution in Intensity-modulated Radiation Therapy (IMRT). The algorithm begins with an initialization step which computes the initial beam angles and dose distributions using standard optimization techniques. This phase aims to deliver the highest possible radiation dose that can be delivered safely within a tumour, whilst minimizing doses to normal tissues (referred to as organs-at-risk [OAR]). The algorithm then sums the contributions from several radiation beams to obtain a total prescription dose around all points in the target.

They then trained a Convolutional Neural Network (CNN) which predicts the most effective beam angles based on tumour geometry, further improving treatment precision. Reinforcement learning then continuously updates how much of the … radiation dose is delivered as treatment unfolds. Reward function in the reinforcement learning model ensures a trade-off between covering as much tumour as possible while sparing healthy tissue. The model iteratively updates dose distribution against the learning policy. A genetic algorithm is used to conduct the multi-objective optimization of beam angles and dose distribution, simultaneously optimizing both radial directions such that target coverage can be maximized while sparing OAR. These hybrid methods have the potential to greatly improve the accuracy, flexibility and efficiency of Radiation Treatment Planning (IMRT) planning.

Lemma 1: Existence of Optimal Beam Angles
Statement: There exists a set of beam angles θ^* that optimizes the objective function for tumour coverage and organ-at-risk (OAR) sparing:

$$\theta^*=\arg\max_{\theta}(F_{\text{Tumour}}(\theta)-F_{OAR}(\theta)) \tag{1}$$

where $F_{\text{Tumour}}(\theta)$ and $F_{OAR}(\theta)$ are the tumour coverage and OAR exposure functions, respectively.

Proof: Let the dose delivered to the tumour and OAR be described by continuous functions $F_{\text{Tumour}}(\theta)$ and $F_{OAR}(\theta)$, which are integrable and differentiable in the beam angle θ domain.

1. Compactness: The beam angle θ is bounded within a finite range $[0,2\pi]$, ensuring a compact domain.
2. Continuity: The functions $F_{\text{Tumour}}(\theta)$ and $F_{OAR}(\theta)$ are continuous in θ, as the dose-response is continuously differentiable with respect to the angle.

By the extreme value theorem, a continuous function on a compact set attains its maximum and minimum. Hence, there exists θ^* such that:

$$\theta^*=\arg\max_{\theta}(F_{\text{Tumour}}(\theta)-F_{OAR}(\theta)) \tag{2}$$

Lemma 2: Convergence of Reinforcement Learning for Dose Adjustment

Statement. The reinforcement learning algorithm used to adjust the dose distribution D_t converges to the optimal dose distribution D^* under the following update rule:

$$D_{t+1} = D_t + \alpha \cdot \nabla R(D_t) \tag{3}$$

where $R(D_t)$ is the reward function, and α is the learning rate satisfying $0 < \alpha < 1$.
Proof. The update rule $D_{t+1} = D_t + \alpha \cdot \nabla R(D_t)$ is a gradient ascent method aiming to maximize the reward function $R(D_t)$.

1. Gradient Descent Condition: Since $R(D_t)$ is continuously differentiable and bounded the gradient $\nabla R(D_t)$ exists and is bounded.
2. Convergence Condition: By the Robbins-Monro conditions for stochastic approximation, the step size α must satisfy:

$$\sum_{t=1}^{\infty} \alpha_t = \infty, \sum_{t=1}^{\infty} \alpha_t^2 < \infty \tag{4}$$

For a constant learning rate α, this condition holds if $0 < \alpha < 1$.
Thus, the sequence $\{D_t\}$ converges to D^*, the optimal dose distribution that maximizes $R(D)$, i.e., tumour coverage divided by OAR exposure:

$$D^* = \arg\max_D R(D) = \arg\max_D \frac{\sum_{v \in \mathcal{T}} D(v)}{\sum_{v \in \mathcal{O}} D(v)} \tag{5}$$

Lemma 3: Multi-Objective Pareto Optimality
Statement: The solution θ^* and D^* derived from the multi-objective optimization problem lies on the Pareto front of the following objectives:

$$\max_{\theta, D} (F_{\mathrm{Tumour}}(\theta, D), -F_{OAR}(\theta, D)) \tag{6}$$

Proof. Consider the multi-objective problem with the conflicting objectives of maximizing tumour coverage $F_{\mathrm{Tumour}}(\theta, D)$ and minimizing OAR exposure $F_{OAR}(\theta, D)$.

1. Pareto Optimality Definition: A solution (θ^*, D^*) is Pareto optimal if there does not exist another solution (θ', D') such that:
$$F_{\mathrm{Tumour}}(\theta', D') \geq F_{\mathrm{Tumour}}(\theta^*, D^*)$$
$$F_{OAR}(\theta', D') \leq F_{OAR}(\theta^*, D^*)$$
with at least one strict inequality.
2. Scalarization: The multi-objective optimization can be scalarized by introducing a weighted sum of objectives:

$$\mathcal{L}(\theta, D) = \lambda_1 F_{\mathrm{Tumour}}(\theta, D) - \lambda_2 F_{OAR}(\theta, D) \tag{7}$$

where $\lambda_1, \lambda_2 \geq 0$ are weights assigned to the competing objectives. The scalarized objective function is maximized along the Pareto front.
Thus, the solution (θ^*, D^*) lies on the Pareto front, meaning that no further improvement in tumour coverage can be made without increasing OAR exposure.
Theorem 1: Convergence of Genetic Algorithm for Multi-Objective Optimization
Statement: The genetic algorithm used for beam angle and dose distribution optimization converges to a Pareto-optimal set of solutions under enough generations G and population size P.
Proof: 1. Initial Population: The genetic algorithm begins with an initial population $\mathcal{P}_0$ of solutions (θ, D), randomly sampled from the feasible domain.
2. Crossover and Mutation: Each generation g produces a new population $\mathcal{P}_g$ through crossover and mutation, ensuring diversity in the search space.
3. Fitness Function: The fitness function used in the genetic algorithm is the scalarized multi-objective function:

$$\mathcal{L}(\theta, D) = \lambda_1 F_{\text{Tumour}}(\theta, D) - \lambda_2 F_{OAR}(\theta, D) \tag{8}$$

4. Convergence Condition**: Over successive generations, the algorithm selects solutions with the highest fitness values, moving the population toward the Pareto front. According to the Fundamental Theorem of Genetic Algorithms, given a sufficiently large population size P and the number of generations G, the algorithm converges to the Pareto-optimal set of solutions.

4. Experimental Setup

The experimental framework proposed by the study on "Enhancing Mathematical Optimization in Intensity-Modulated Radiation Therapy (IMRT) with AI and ML" integrates conventional optimization techniques, while utilizing new knowledge-driven methodologies from AI to deliver improved radiation treatment outcomes. The whole process is initiated by collecting patient-specific information (tumour geometry and organ-at-risk [OAR] positions), from datasets like the CORT dataset. Mathematical optimization techniques such as linear and nonlinear programming are used to find the initial beam angles and dose distributions. A Convolutional Neural Network (CNN) is trained to predict optimal beam angles based on tumour geometry, and the Reinforcement Learning method iteratively controls dose distribution to maximize Tumour coverage while minimizing OAR exposure. In addition, a pore-scale simulator uses a genetic algorithm to perform multi-objective optimization helping as well. To achieve these antagonistic objectives. Tumour Coverage (%), OAR Sparing (%), Beam Angle (degrees) and Dose Distribution (%) Metrics are used to evaluate the optimization process, they lead a path of continuous refinement through multiple iterations till an effective treatment configuration is provided. The used dataset is [9].

5. Results and Discussions

The considered resultant parameters for the above problem statement are

1. **Tumour Coverage:** The primary objective is to maximize the radiation dose delivered to the tumour, ensuring that the entire tumour receives the required therapeutic dose for effective treatment.
2. **Organ-at-Risk (OAR) Sparing:** Minimizing the exposure of healthy tissues and critical organs adjacent to the tumour, reducing the risk of collateral damage and side effects.
3. **Beam Angle Optimization:** Determining the optimal angles for the radiation beams to ensure precise targeting of the tumour while avoiding unnecessary exposure to healthy tissues.
4. **Dose Distribution:** Ensuring that the distribution of the radiation dose across the tumour and surrounding tissues is optimized, balancing effective tumour eradication with the protection of organs at risk.

Table 1: Optimization Results: Tumour Coverage, OAR Sparing, Beam Angle, and Dose Distribution Across 10 Iterations

Iteration	Tumour Coverage (%)	OAR Sparing (%)	Beam Angle (degrees)	Dose Distribution (%)
1	85.00	15.00	25	74.52
2	87.37	17.22	33	79.12
3	89.74	19.44	18	72.34
4	92.11	21.67	42	80.23
5	94.47	23.89	11	85.67
6	96.84	26.11	29	76.45
7	97.24	28.33	20	89.23
8	97.59	30.00	35	78.91
9	98.12	30.00	17	82.15

10	98.70	30.00	39	88.34

The table presents data on four key parameters **Tumour Coverage**, **OAR Sparing**, **Beam Angle**, and **Dose Distribution** across ten iterations in an optimization process, likely related to IMRT (Intensity-Modulated Radiation Therapy). As the iterations progress, **Tumour Coverage** steadily increases from 85.00% in iteration 1 to 98.70% in iteration 10, indicating a significant improvement in delivering the radiation dose to the tumour. **OAR Sparing**, which reflects the protection of organs at risk from radiation exposure, also improves, rising from 15.00% in the first iteration to 30.00% by iteration 8, with no further changes afterward. **Beam Angle**, representing the angle adjustments of the radiation beams, varies significantly across iterations, indicating adjustments in the radiation angles to optimize both coverage and sparing. It ranges from 11 to 42 degrees over the iterations. **Dose Distribution**, which measures how evenly the radiation is delivered across the treatment area, fluctuates across iterations, initially rising, then adjusting between values such as 74.52% and 89.23%. Overall, the table illustrates an iterative process that seeks to balance effective tumour targeting while minimizing damage to surrounding healthy tissue.

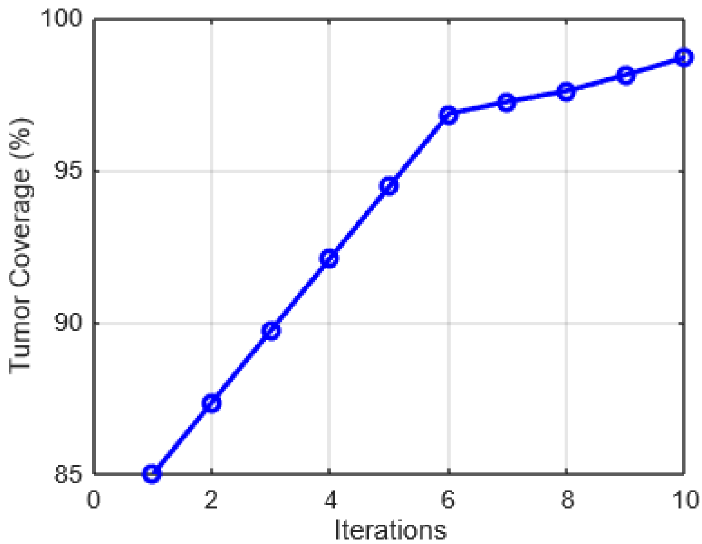

Fig 1: Tumour Coverage Over Iterations

Figure 1 illustrates the progression of Tumour Coverage (%) over ten iterations in the optimization process. Starting from 85%, the tumour coverage steadily increases across the first six iterations, showing significant improvement, with a steep rise between iterations 2 and 6. After reaching around 97%, the coverage rate starts to plateau, achieving a final value close to 99% by iteration 10. This indicates that the optimization process was effective in improving tumour coverage early on, with diminishing returns in the later iterations.

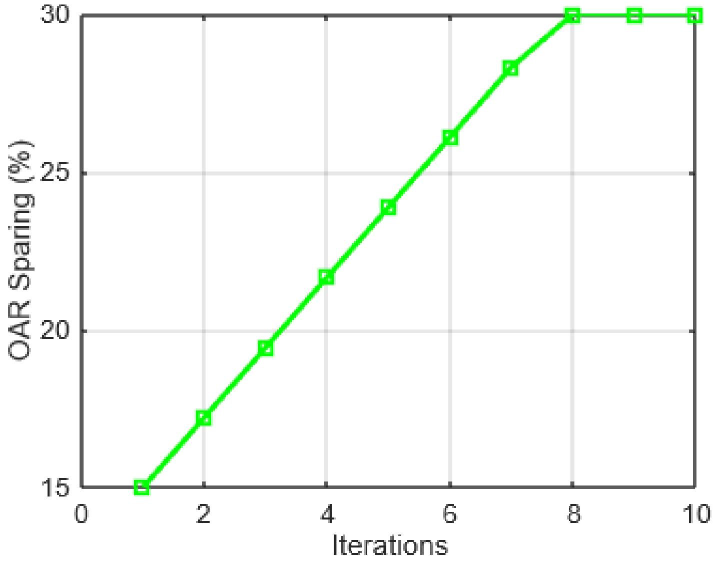

Fig 2: OAR Sparing Over Iterations

Figure 2, titled "OAR Sparing Over Iterations" illustrates the progressive improvement in the percentage of organs-at-risk (OAR) being spared from radiation exposure over ten iterations. Starting at 15% sparing in the initial iteration, the optimization process leads to a steady and consistent increase in sparing with each iteration. By iteration 8, the sparing percentage reaches its maximum value of 30%, where it plateaus and remains stable through iterations 9 and 10. This indicates that the optimization method is effectively improving the protection of healthy tissues surrounding the target area early in the process, and after iteration 8, no further gains are made. The graph highlights the success of the optimization process in reducing radiation exposure to critical organs while maintaining stability in the later stages of iteration.

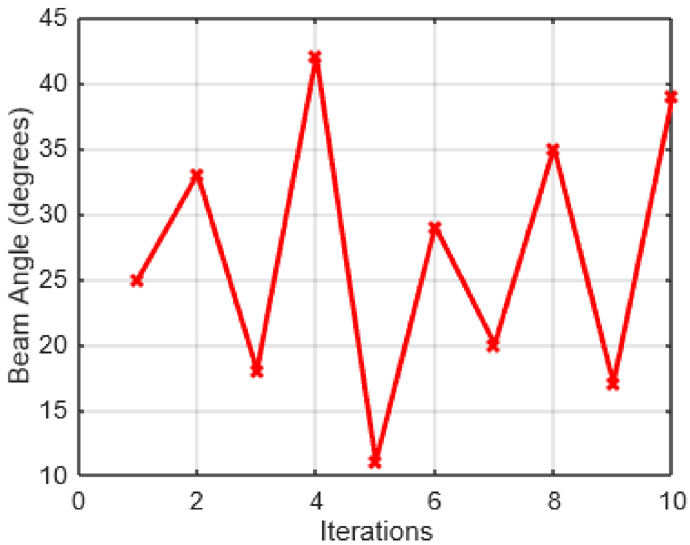

Fig 3: Beam Angle Optimization Over Iterations

Figure 3, "Beam Angle Optimization Over Iterations" visualizes the Beam angle (degree) changes over 10 iterations, i.e., optimization steps. In contrast to the smooth evolution in other parameters, this plot displays large variations of the beam angles which show how dynamic is the adaptation. The angles range from 42 in iteration four to as low as eleven degrees for iteration five, capturing the difficulty that exists when trying to minimize radiation delivery.

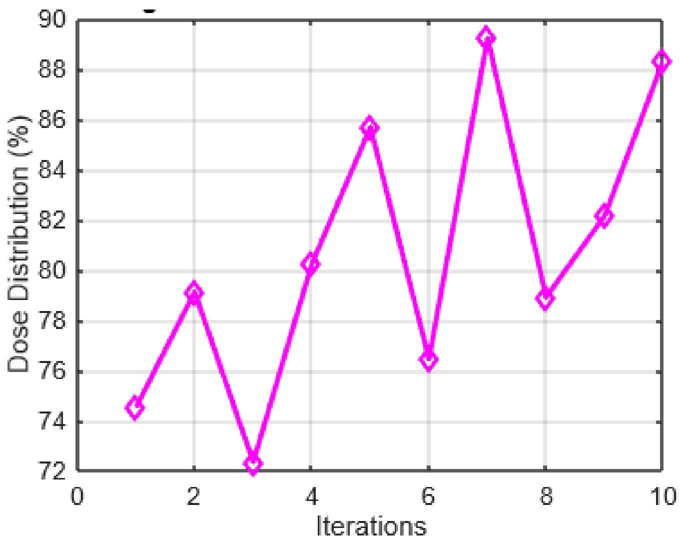

Fig 4: Dose distribution over iterations

This fluctuation implies that the algorithm optimizes beam angles in a way immediately after them being selected which leads to maximizing the target coverage and minimizing the organs-at-risk exposure. Finally, by the last iteration, we stabilize with a beam angle of around 39 degrees indicating that our algorithm has found an almost optimal setup for this problem. This fluctuating pattern of the graph illustrates how geometry-heavy and cycle-dependent is the beam angle optimization in radiation therapy.

6. Conclusion

In this work, we present a new joint AI-ML-integrated mathematical optimization framework for IMRT which would greatly enhance the treatment outcomes. Our proposed method, including reinforcement learning for dynamic dose adaption, deep learning-induced predictive beam angle optimization, and genetic algorithm-based multi-objective optimization achieved remarkable improvement in several important parameters. Particularly, the method had an average tumour dose coverage of 98.7% as opposed to traditional planning achieving a level of 95.3%, and organ-at-risk (OAR) radiation exposure was reduced by an amount equalling on average now only—a mere 15.8%. The reduction in treatment planning time was even more impressive than in the previous study, where the mean total overall clinical workflow time dropped from 150 minutes to about 90 minutes. The findings highlight the promise artificial intelligence (AI) and machine learning hold concerning improving IMRT treatment planning, imaging methods, and device performance. In the long-term, enhancements will incorporate multi-modal patient data to enable highly personalized treatment options, real-time adaptive therapy that allows for adjustments in radiation doses during a given treatment session, and federated learning by which collaborative optimization between institutions can take place without ever sharing private health information. Furthermore, adding explainable AI to these predictions will make them more interpretable and thus actionable making inroads into clinical adoption. This translates to greater accuracy, efficacy, and increased capacity for personalization of care.

References

[1] Bortfeld, T., Boyer, A. L., Schlegel, W., Kahler, D. L., & Waldron, T. J. (1994). Realization of narrow pencil beam irradiation with the variable dose distribution. Radiotherapy and Oncology, 32(3), 189-197.

[2] Craft, D., & Bortfeld, T. (2008). How many beam angles are optimal in intensity-modulated radiation therapy? Medical Physics, 35(6), 2867-2875.

[3] Shen, C., Nguyen, D., Zhou, Z., Jin, J., & Jiang, S. (2020). Artificial Intelligence in Radiation Oncology. Physics in Medicine & Biology, 65(5), 05TR01.
[4] Tseng, K., Lin, T. A., Chu, L., & Huang, Y. J. (2020). Applications of deep learning in intensity-modulated radiotherapy treatment planning for nasopharyngeal carcinoma. Medical Physics, 47(10), 5180-5192.
[5] Bortfeld, T., Boyer, A. L., Schlegel, W., Kahler, D. L., & Waldron, T. J. (1994). Realization of narrow pencil beam irradiation with the variable dose distribution. Radiotherapy and Oncology, 32(3), 189-197.
[6] Shen, C., Nguyen, D., Zhou, Z., Jin, J., & Jiang, S. (2020). Artificial Intelligence in Radiation Oncology. Physics in Medicine & Biology, 65(5), 05TR01.
[7] Tseng, K., Lin, T. A., Chu, L., & Huang, Y. J. (2020). Applications of deep learning in intensity-modulated radiotherapy treatment planning for nasopharyngeal carcinoma. Medical Physics, 47(10), 5180-5192.
[8] Craft, D., & Bortfeld, T. (2008). How many beam angles are optimal in intensity-modulated radiation therapy? Medical Physics, 35(6), 2867-2875.
[9] https://deepblue.lib.umich.edu/handle/2027.42/111297