Physics and Imaging in Radiation Oncology 33 (2025) 100715

Contents lists available at ScienceDirect

Physics and Imaging in Radiation Oncology

journal homepage: www.sciencedirect.com/journal/physics-and-imaging-in-radiation-oncology

Original Research Article

Optimal timing of organs-at-risk-sparing adaptive radiation therapy for head-and-neck cancer under re-planning resource constraints

Fatemeh Nosrat [a,*,1], Cem Dede [b,2], Lucas B. McCullum [b,c,2], Raul Garcia [a,1], Abdallah S.R. Mohamed [b,d,2], Jacob G. Scott [e], James E. Bates [f], Brigid A. McDonald [b,2], Kareem A. Wahid [b,2], Mohamed A. Naser [b,2], Renjie He [b,2], Aysenur Karagoz [a,1], Amy C. Moreno [b,2], Lisanne V. van Dijk [b,g], Kristy K. Brock [h], Jolien Heukelom [i], Seyedmohammadhossein Hosseinian [j,1], Mehdi Hemmati [k,1], Andrew J. Schaefer [a,*,1], Clifton D. Fuller [a,b,*,1,2]

[a] Department of Computational Applied Mathematics and Operations Research, Rice University Houston TX USA
[b] Department of Radiation Oncology, The University of Texas MD Anderson Cancer Center Houston TX USA
[c] The University of Texas MD Anderson Cancer Center UTHealth Houston Graduate School of Biomedical Sciences Houston TX USA
[d] Department of Radiation Oncology, Baylor College of Medicine Houston TX USA
[e] Department of Translational Hematology and Oncology Research, Lerner Research Institute Cleveland OH USA
[f] Department of Radiation Oncology, Emory University Atlanta GA USA
[g] Department of Radiation Oncology, University of Groningen University Medical Center Groningen Groningen Netherlands
[h] Department of Imaging Physics, The University of Texas MD Anderson Cancer Center Houston TX USA
[i] Department of Radiation Oncology (Maastro), GROW School for Oncology and Reproduction Maastricht University Medical Centre+ Maastricht Netherlands
[j] Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University Raleigh NC USA
[k] School of Industrial and Systems Engineering, University of Oklahoma Norman OK USA

ARTICLE INFO

Keywords:
Personalized adaptive radiation therapy
Organs at risk
Normal tissue complication probability
Markov decision process
Optimal strategy

ABSTRACT

Background and purpose: Prior work on adaptive organ-at-risk (OAR)-sparing radiation therapy has typically reported outcomes based on fixed-number or fixed-interval re-planning, which represent one-size-fits-all approaches and do not account for the variable progression of individual patients' toxicities. The purpose of this study was to determine the personalized optimal timing of re-planning in adaptive OAR-sparing radiation therapy, considering limited re-planning resources, for patients with head and neck cancer (HNC).

Materials and methods: A novel Markov decision process (MDP) model was developed to determine optimal timing of re-planning based on the patient's expected toxicity, characterized by normal tissue complication probability (NTCP), for four toxicities. The MDP parameters were derived from a dataset comprising 52 HNC patients treated between 2007 and 2013. Kernel density estimation was used to smooth the sample distributions. Optimal re-planning strategies were obtained when the permissible number of re-plans throughout the treatment was limited to 1, 2, and 3, respectively.

Results: The MDP (optimal) solution recommended re-planning when the difference between planned and actual NTCPs (ΔNTCP) was greater than or equal to 1%, 2%, 2%, and 4% at treatment fractions 10, 15, 20, and 25, respectively, exhibiting a temporally increasing pattern. The ΔNTCP thresholds remained constant across the number of re-planning allowances (1, 2, and 3).

Conclusion: In limited-resource settings that impeded high-frequency adaptations, ΔNTCP thresholds obtained from an MDP model could derive optimal timing of re-planning to minimize the likelihood of treatment toxicities.

* Corresponding authors at: Department of Computational Applied Mathematics and Operations Research, Rice University, Houston, TX, USA (C.D. Fuller).
E-mail addresses: fatemeh.nosrat@rice.edu (F. Nosrat), andrew.schaefer@rice.edu (A.J. Schaefer), cdfuller@mdanderson.org (C.D. Fuller).
[1] On behalf of the Rice/MD Anderson Center for Operations Research in Cancer (CORC).
[2] On behalf of the MD Anderson Head and Neck Cancer Symptom Working Group.

https://doi.org/10.1016/j.phro.2025.100715
Received 13 May 2024; Received in revised form 29 November 2024; Accepted 24 January 2025
Available online 27 January 2025

1. Introduction

Advancements in radiation delivery techniques, such as intensity-modulated radiation therapy (IMRT) and volumetric-modulated arc therapy, enable accurate dose delivery to tumor targets while minimizing radiation exposure of the surrounding organs at risk (OARs) [1]. However, anatomical changes during the treatment, such as weight loss or tumor shrinkage, may cause the actual delivered dose to OARs to deviate from the planned dose. This can increase the risk of treatment-induced toxicities, particularly in cases where multiple OARs are in close proximity to the target, as in head and neck cancer (HNC) [2–4]. To address this, adaptive radiation therapy (ART) has been clinically introduced, proposing on-therapy re-planning in response to anatomical changes in the target and OARs [5–10].

In practice, however, the clinical implementation of ART with daily (or even less-frequent) re-planning remains limited, in large part due to the extensive human/personnel/workflow resources required to frequently perform key tasks such as segmentation and quality assurance as well as limited device accessibility time [11,12]. Recent artificial intelligence (AI)-based algorithms (such as auto-segmentation or synthetically created CTs) [13,14] may mitigate some or all of these process level frictions; however, the integration of such AI tools within the ART workflow is still evolving [15]. With the advent of hybrid MR-Linac devices, real-time adjustment of daily radiation plans, known as on-line ART, is now a possibility. On-line ART can also be enabled with the availability of high-frequency, high-quality cone-beam CT or CT-on-Rails devices [16,17]. Regardless of ART implementation imaging inputs (MR or CT), cancer centers typically have implemented ART at fixed intervals, notably once mid-therapy [18] and often as a 'verification' of re-simulation. Most of the relevant studies also only report outcomes on fixed-number and/or fixed-interval re-planning [19–21]; see Supplementary Table S1 for a comprehensive literature review. Such predetermined schedules for treatment re-planning, however, take a one-size-fits-all approach and do not account for the uncertain trajectory of individual patients' toxicities [21], nor patient-specific tumor regression. As a result, determining the optimal timing of re-planning episodes remains a crucial unmet need, particularly for OAR-sparing adaptive approaches (whether for MR-Linac as have been implemented in MR-guided clinical trials [19] or for analogous CT-based approaches [20]).

Heukelom et al. [22] investigated the optimal implementation of ART with a single re-planning allowance (in OAR-sparing radiation therapy) using daily on-treatment CT imaging with a CT-on-rails device. Leveraging the same dataset, this paper presents a new analytical approach to derive optimal re-planning strategies based on Markov decision process (MDP) models. Our aim is to identify the optimal timing of re-planning based on changes in normal tissue complication probabilities (NTCP) of four toxicities: xerostomia, dysphagia, parotid gland dysfunction, and feeding tube dependency at 6 months post-treatment. We further include allowances in HNC treatment plan adaptations (through limiting the number of available re-plans) to enhance personalized treatment and efficacy. MDPs constitute a class of mathematical optimization models that aim to determine optimal decisions/actions in stochastic dynamic systems [23,24]. MDPs have been successfully employed to find the optimal timing for various medical interventions [25–31]; however, to our knowledge, MDPs have not been applied for triggering adaptive re-planning. We develop a generalized framework for the utilization of MDPs for evidence-based individualized radiation treatment re-planning, scalable across resource- rich and resource-limited facilities, and applicable to both CT- and MR-based platforms. Thus, rather than a class solution based on population estimates of toxicity reduction potential, we enable personalized adaptive therapy with consideration of a budget.

2. Materials and methods

2.1. 3 Data

This study used a prior dataset of CT-on-Rails image-guided radiation therapy (IGRT), detailed by Heukelom et al. [22], which comprised information from patients treated for HNC at the University of Texas MD Anderson Cancer Center between 2007 and 2013; this retrospective secondary analysis was performed under MD Anderson Cancer Center Institutional Review Board approval (MDA RCR03-0800). The patients' characteristics are summarized in Supplementary Table S2.

For these 52 HNC patients, Heukelom et al. [22] calculated the deviation of the actual dose from the planned dose for nine OARs at fractions 10 and 15 of the treatment. At each fraction, they estimated NTCP for the toxicities related to the OARs (xerostomia, dysphagia, parotid gland dysfunction, and feeding tube dependency at 6 months post-treatment) by projecting the actual dose through the remainder of the treatment period. Subsequently, they compared these findings with the planned NTCPs and determined the difference, i.e., ΔNTCP, for each toxicity. The NTCP models are presented in Supplementary Table S3. The MDP model presented in this paper used the ΔNTCP from this dataset [22], which are summarized in Table 1. For each observed ΔNTCP value, Heukelom et al. [22] reported the number of patients for whom this ΔNTCP was the highest value among the four NTCP models.

2.2. Decision model

In the MDP model, the estimates of an individual patient's toxicity outcome, as a function of the delivered radiation dose to OARs, determine the state of the system at each decision epoch during the treatment (e.g., day). Depending on the observed state, the clinician may decide between two possible actions: (1) Re-plan or (2) continue with the current plan. When the action is to continue with the current plan, the system may transition from one toxicity state to another stochastically, governed by transition probabilities. Re-planning changes the probabilistic transition towards more favorable outcomes/states. Given that a limited number of 're-planning' actions may be taken throughout the treatment, an optimal solution to the MDP model identifies the optimal timing for taking such actions, as a function of the toxicity states. The MDP model captures the stochastic evolution of post-treatment toxicity risk and identifies optimal re-planning times to mitigate the toxicities (if necessary). The components of the MDP model are as follows:

2.2.1. Re-planning allowance

Depending on available resources for plan adaptations, the model considered a maximum number of re-plans B that could be implemented throughout the treatment. The analysis was performed for $B = 1, 2, 3$.

2.2.2. Decision epochs

Given a treatment period consisting of 33–35 fractions, the decision epochs were set at fractions 10, 15, 20, and 25. Prior studies have shown that anatomical changes are unlikely to happen very early during the treatment [22]; thus fraction 5 was omitted. Fraction 30 was also excluded due to its proximity to the end of treatment, with negligible impact on the total dose to the OARs.

2.2.3. States

At each decision epoch, the state of the system was captured by the pair ($\Delta NTCP, b$), where ΔNTCP denoted the deviation of treatment toxicity from the planned value at that time, and $b \leq B$ was the number of remaining re-plans. The ΔNTCP ranged from 0 % to 12 % in the model (Table 1); in computing the number of cases for each reported ΔNTCP value, only those patients were included who experienced the change of ΔNTCP in at least one of the four aforementioned toxicities.

Table 1
Observed ΔNTCP values based on the difference between the planned dose and the actual dose, along with the number of patients associated with each ΔNTCP value. Adapted from [22] with permission.

(a) Fraction 10												
ΔNTCP	0 %	1 %	2 %	3 %	4 %	5 %	6 %	7 %	8 %	9 %	10 %	11 %	12 %
Number of Patients	26	7	4	6	2	2	0	1	0	0	1	1	2
(b) Fraction 15													
ΔNTCP	0 %	1 %	2 %	3 %	4 %	5 %	6 %	7 %	8 %	9 %	10 %	11 %	12 %
Number of Patients	23	8	9	2	2	1	1	2	0	0	2	1	1

2.2.4. Actions

At each decision epoch, two possible actions were included: (1) Re-planning, or (2) continuing with the current plan (no re-planning). The stochastic transition of the toxicity state from one decision epoch to the next was a function of the action taken and was determined by the associated transition probabilities. See Fig. 1 for an illustration of the MDP.

2.2.5. Transition probabilities

The transition probabilities, governing the stochastic evolution of toxicity under each action, were estimated using the information presented in Table 1. Under the 'no re-planning' action, the probabilities for transitions from fraction 0 to 10 and from fraction 10 to 15 were directly estimated based on the number of patients in each ΔNTCP category. The transition probability matrix from fraction 10 to 15 was smoothed out using kernel density estimation with Gaussian kernels, implemented in Python [27]. The probabilities under the 'no re-planning' action are presented in Supplementary Tables S4–S6. It was assumed that transition probabilities from fraction 10 to 15 remained constant for subsequent decision epochs, due to the absence of reported ΔNTCP information beyond fraction 15 by Heukelom et al. [22]. To model the impact of the 're-planning' action on transition probabilities, it was assumed that, at each decision epoch, the action immediately decreased ΔNTCP to a value proportional to the elapsed treatment time and then followed by a transition (from fraction 10 to 15) according to the probabilities associated with the decreased ΔNTCP value at fraction10; see Fig. 2. The probabilities under the 're-planning' action are presented in Supplementary Tables S7–S10.

2.2.6. Rewards

For each set of consecutive actions taken at the decision epoch, the model considered the expected −ΔNTCP at the end of the treatment period (with respect to the transition probabilities) as the corresponding reward. The objective of the MDP model was to maximize the expected reward by identifying an optimal set of actions, one at each decision epoch, as a function of the system's state. This is referred to as an optimal policy. Because the rewards were defined by negative values in the model, smaller end-treatment ΔNTCP values translated to higher rewards.

The MDP model was solved using the MDPtoolbox of MATLAB [32], for B = 1, 2, 3. The MATLAB code and its outputs are available at https://figshare.com/s/64bc3481737d17fc287e.

The optimal policy of an MDP model may become a single-threshold policy (also referred to as control-limit policies) [24], which refers to a class of policies that use thresholds on the state value to recommend an action. For this study, a single-threshold policy recommended re-planning when the ΔNTCP exceeded a threshold, while no re-planning was needed when it falls below the threshold. This policy reduces the complexity of decision-making to a simple rule that uses only one threshold at each fraction to trigger action and is efficient to implement.

3. Results

The analysis revealed that the optimal policy was a single-threshold policy. When only one re-plan was allowed ($B = 1$), the optimal policy at fraction 10 was to re-plan for any ΔNTCP value greater than or equal to 1 %. Subsequently, at fraction 15, this threshold increased to 2 % and remained at 2 % for fraction 20. At fraction 25, the minimum ΔNTCP required for a re-planning was 4 %. These thresholds remained the same

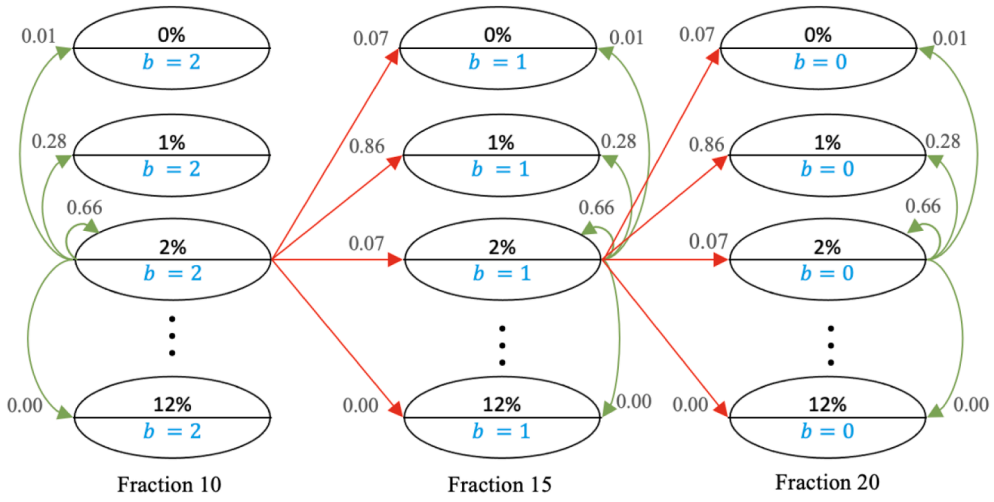

Fig. 1. Markov Decision Process Model. The permissible number of re-plans is 2 (B = 2). The states are shown by ellipses.ΔNTCP values are located in the upper half of the ellipses, while the lower halves contain the value of b (the number of remaining re-plans). Transitions between states are shown by green arrows when the action is 'no re-planning' and by red arrows when the action is 're-planning.' Not all arrows, states, and fractions are included to avoid ambiguity. For example, the smoothed transition probability from state (2 %,2) state to (1 %,1) when the action is 're-planning' is 0.86. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

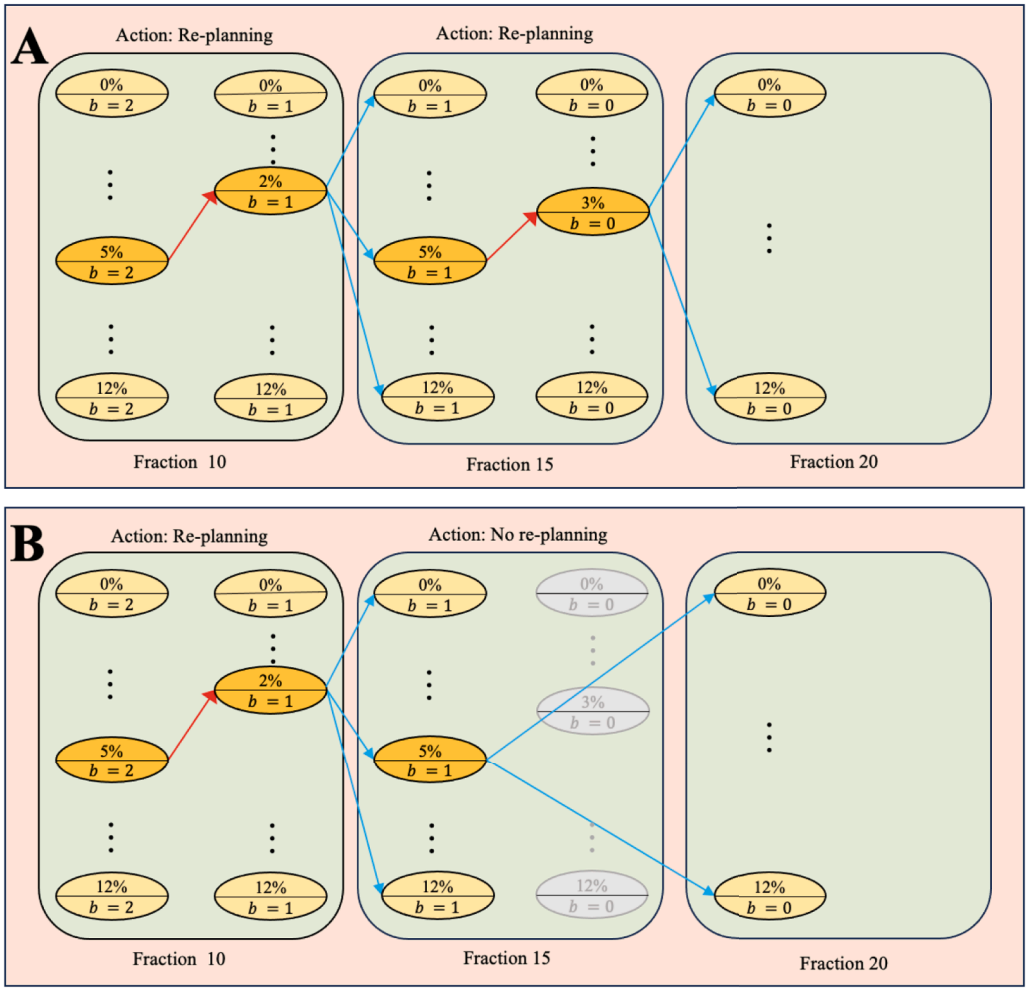

Fig. 2. Effect of actions on system transitions. The permissible number of re-plans is 2 (B = 2). The states are represented by ellipses. The ΔNTCP values are located in the upper halves of the ellipses, while the lower halves contain the value of b (the number of remaining re-plans). When the action at a fraction is 're-planning,' as in Fractions 10 and 15 in Part A and Fraction 10 in Part B, the transition to a state with a lower ΔNTCP within the same fraction is shown with red arrows. Transitions between states from one fraction to the next are shown with blue arrows. When the action is 'no re-planning,' as in Fraction 15 in Part B, no transition occurs within that fraction. Instead, the system transitions to a new state at Fraction 20. This is depicted by the gray states at Fraction 15, indicating that there is no immediate decrease in ΔNTCP. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

in the optimal policies for $B = 2, 3$. The results were summarized in Table 2 and illustrated in Supplementary Fig. S1 for $B = 3$.

4. Discussion

This study introduced the first — to our knowledge — application of the mathematically rigorous MDP methodology, to determine optimal timing of ART in HNC. The MDP model guides clinicians in determining the minimum values of ΔNTCP at fractions 10, 15, 20, and 25 for performing a re-plan, given a re-planning allowance of 1, 2, or 3 throughout the treatment.

The re-planning ΔNTCP threshold increased over time in the optimal policy, consistent with the diminishing impact of re-planning as the treatment progresses. The results suggest optimality of re-planning for any changes in NTCP (ΔNTCP ≥ 1) at fraction 10. This supports the findings of Heukelom et al. [22], who identified fraction 10 as the optimal time for a single re-plan. Importantly, at a given fraction, the ΔNTCP thresholds remained the same for different number of re-plans ($B = 1, 2, 3$); the re-planning allowance did not affect these thresholds. Furthermore, in cases where ΔNTCP at fraction 25 was below the 4 % threshold, both actions ('re-planning' or 'no re-planning') yielded the same impact on the end-treatment ΔNTCP, indicating that re-planning did not result in an improvement.

We acknowledge that minimizing the expected ΔNTCP may result in prescribing re-planning for any ΔNTCP value, potentially leading to a high number of false negatives. This arises from our modeling assumptions, which permitted a fixed number of re-plans at no cost, encouraging frequent re-planning due to low ΔNTCP thresholds. In an ongoing study, we are incorporating re-planning costs without limiting the number of re-plans to better explore the trade-off between cost and benefit; however, at a minimum, we have opted to err on the side of patient benefit, rather than cost-control; as patient NTCP benefit is potentially scalable across any health system, while cost per re-plan and acceptable cost constraints are variable across national and international health policy and reimbursement systems.

Table 2
Optimal re-planning thresholds based on ΔNTCP and re-planning allowance.

Re-planning allowance (B)	1	2		3		
Number of remaining re-plans (b)	1	2	1	3	2	1
ΔNTCP threshold at fraction 10	1 %	1 %	–	1 %	–	–
ΔNTCP threshold at fraction 15	2 %	2 %	2 %	2 %	2 %	–
ΔNTCP threshold at fraction 20	2 %	2 %	2 %	2 %	2 %	2 %
ΔNTCP threshold at fraction 25	4 %	4 %	4 %	4 %	4 %	4 %

While seemingly minor, a ΔNTCP of 4 % or below can significantly improve patient outcomes by reducing severe toxicity, such as osteoradionecrosis, and thereby enhancing quality of life and survival. Optimizing treatment at this level allows healthcare practitioners to minimize toxicity and maximize patient benefit.

It is crucial to acknowledge the limitations regarding the generalizability of findings from the single-site, retrospective, in silico CT-on-rails reference dataset [22] used in this study. For instance, the in silico daily dose accumulation was not actively applied to individual patients, but rather calculated post hoc from a high-granularity CT-on-rails daily volumetric IGRT series. The CT-on-rails platform at MD Anderson utilized an in-house custom- constructed intermediary localization and alienation software (CT-Assisted Targeting (CAT)) [33]. Consequently, there were instances where delivered geometric shifts were either unrecorded or unrecoverable, or clearly aberrant (such as extensive shift records representing an initial setup that was then revised after repositioning) during the secondary export of coordinate displacement to the commercial Record and Verify software (Mosaiq, Elekta AB). These discrepancies were subsequently omitted in the in-silico model to streamline data, leading to conceptual gaps in the resultant NTCP modeling where these missing values were not accounted for. Furthermore, it is important to note that this modeled secondary dataset did not include adaptation or daily re-optimization of the initial daily dose in vivo. Consequently, the data presented in this paper should be viewed as a clinically approximate semi-synthetic illustrative use-case, rather than a definitive rationale for the large-scale implementation of the observed idealized re-planning thresholds across distinct operational platforms. Nonetheless, we believe that the resultant MDP model could be readily scaled using higher-quality prospective or observational cohort data for secondary validation. In essence, the individualized planning parameters suggested by the MDP model should be viewed as proof of concept rather than a formal criterion barring external validation.

In the prior work, Heukelom et al. [22] exclusively reported ΔNTCP for fractions 10 and 15, consistent with internal re-planning practices at MD Anderson Cancer Center based on data from a Phase II study by Maki et al. [34]. Consequently, for model extensibility in the current application, we have explicitly assumed transition probabilities remain stable for the subsequent epochs. This assumption introduced a known level of uncertainty that warrants consideration and is an area of future research, as it has been unclear for specific OARs whether these transition states were indeed stable over therapy. Moreover, the MDP model stipulated weekly re-planning intervals on indexed fractions (e.g., fractions 10, 15, 20, 25, and 30) as a simplification for clarity of presentation reflective of our current adaptive protocols [19], but could readily be adapted to continuous daily fraction-based re-planning intervals.

The four NTCP models in Supplementary Table S3 are among the most currently used models to calculate the NTCP values for the considered toxicities [22,35]. With the advent of more recent NTCP models for HNC radiation therapy, e.g., [36], it is possible that new NTCP models could offer improved estimations. While we recognize the importance of sample size in calibrating MDP models, data collected from 52 HNC patients is considered substantial in the context of HNC research. Furthermore, our results are contingent upon the available CT-on-Rails data, and future research may benefit from incorporating higher-dimensional data (e.g., GTV/the clinical target volume (CTV) modifying approaches, MRI anatomic and/or biomarker data for TCP/NTCP) for a more extensive insight.

Nonetheless, the proposed MDP model for ART is clinically relevant, mathematically rigorous, resource-aware, and scalable, and can be adjusted based on new OAR toxicity with reference NTCP values. Necessarily, the precision of the model relies on accurate calculations of NTCP, particularly when adhering to rigorous criteria that determine whether patients are suitable for or excluded from ART [37]. Despite the challenges and limitations, our study introduces novel contributions to the field of ART. Unlike previous works [22], which primarily considered a single re-planning allowance and only reported results for fractions 10 and 15, our optimization model extends its applicability to scenarios with multiple re-plans, and our policy spans across fractions 10, 15, 20, and 25. Furthermore, our methodology is different from Gan et al. [38], which simulated different ART scenarios and evaluated accumulated dose differences before and after re-planning to determine the optimal timing for re-planning.

An aspect not explored in this study is the adaptation based on GTV or CTV modification, either for shrinking GTV/CTVs [19] or isotoxic boost approaches [39,40]; we concentrated solely on OAR-based adaptation. Adapting based on GTV could open avenues for optimal re-plans, potentially influencing NTCP and extending into scenarios such as Stereotactic Body RT [41]. This introduces a distinctive problem and solution space beyond the scope of our current investigation. Furthermore, we exclusively focused on optimizing the ART workflow within the context of photon therapy. Similar optimization methodologies could prove advantageous when exploring ART in the context of proton therapy, particularly in addressing setup variability reduction [42]. For this study, NTCP calculations were based on the 'plan of the day.' While our findings may vary with deformable dose registration, the operational implementation remains consistent. Future efforts should consider incorporating deformable dose registration to enhance the model's generalization.

Several surveys in both low/middle-income countries [43,44] and high-income economies [44] have identified resource constraints as an impediment to ART implementation; the use of models such as our MDP provides a potential avenue for stratification of resource allocation. Put simply, with one re-plan allowed, almost all patients would be best served via re-planning early during treatment; however, as the budget of re-planning staff/technical/time resources expand, evidence-based personalized re-planning is potentiated by our MDP approach.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the National Institutes of Health (NIH) National Cancer Institute (NCI) Research Grant (R01CA257814), the NCI Cancer Center Support Grant Program in Image-Driven Biologically-informed Therapy (IDBT) Program (P30CA016672), and the Image Guided Cancer Therapy Research Program at The University of Texas MD Anderson Cancer Center, and the MD Anderson Charles and Daneen Stiefel Center for Head and Neck Cancer Oropharyngeal Cancer Research Program. Lucas B. McCullum and Raul Garcia were supported by the NCI Supplement program under R01CA257814-02S2, and R01CA257814-03S1, respectively. Drs. McDonald and Wahid are both supported by the Image-guided Cancer Therapy T32 Fellowship (T32CA261856). Dr. Heukelom received related support from the Netherlands Rene Vogels Fond Fellowship. Drs. Heukelom and van Dijk received relevant support from the Dutch Cancer Society/KWF Kankerbestrijding. Dr. Fuller received related support from the National Institute of Dental and Craniofacial Research (NIDCR) (R01DE028290). Dr. Fuller has received related direct industry grant/in-kind support, honoraria, and travel funding from Elekta AB. Dr. Fuller has served in an unrelated consulting capacity for Varian/Siemens Healthineers. Philips Medical Systems, and Oncospace, Inc. Dr. Brock received unrelated support from the Helen Black Image Guided Fund, RaySearch Laboratories AB, the Apache Corporation, and the Tumor Measurement Initiative through the MD Anderson Strategic Initiative Development Program (STRIDE) and various NIH mechanisms. Dr. Scott received relevant support under the NCI Case Comprehensive Cancer Center Support Grant (P30CA043703), with additional unrelated NIH support. Dr. Moreno received related support from NIDCR (K01DE030524) and

NCI (K12CA088084) during the project period, with additional unrelated NIH support.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.phro.2025.100715.

References

[1] Iacovelli NA, Cicchetti A, Cavallo A, Alfieri S, Locati L, Ivaldi E, et al. Role of IMRT/VMAT-based dose and volume parameters in predicting 5-year local control and survival in nasopharyngeal cancer patients. Front Oncol 2020;10:e518110. https://doi.org/10.3389/fonc.2020.518110.

[2] Barker Jr JL, Garden AS, Ang KK, O'Daniel JC, Wang H, Court LE, et al. Quantification of volumetric and geometric changes occurring during fractionated radiotherapy for head-and-neck cancer using an integrated CT/linear accelerator system. Int J Radiat Oncol Biol Phys 2004;59:960–70. https://doi.org/10.1016/j.ijrobp.2003.12.024.

[3] Beltran M, Ramos M, Rovira JJ, Perez-Hoyos S, Sancho M, Puertas E, et al. Dose variations in tumor volumes and organs at risk during IMRT for head-and-neck cancer. J Appl Clin Med Phys 2012;13:e3723.

[4] Chen W, Bai P, Pan J, Xu Y, Chen K. Changes in tumor volumes and spatial locations relative to normal tissues during cervical cancer radiotherapy assessed by cone beam computed tomography. Technol Cancer Res Treat 2017;16:246–52. https://doi.org/10.1177/1533034616685942.

[5] Beadle BM, Chan AW. The potential of adaptive radiotherapy for patients with head and neck cancer-too much or not enough? JAMA Oncol 2023;9:1064–5. https://doi.org/10.1001/jamaoncol.2023.1306.

[6] Castelli J, Simon A, Lafond C, Perichon N, Rigaud B, Chajon E, et al. Adaptive radiotherapy for head and neck cancer. Acta Oncol 2018;57:1284–92. https://doi.org/10.1080/0284186X.2018.1505053.

[7] Figen M, Öksüz DÇ, Duman E, Prestwich R, Dyker K, Cardale K, et al. Radiotherapy for head and neck cancer: evaluation of triggered adaptive replanning in routine practice. Front Oncol 2020;59:960–70. https://doi.org/10.3389/fonc.2020.579917.

[8] Håkansson K, Giannoulis E, Lindegaard A, Friborg J, Vogelius I. CBCT-based online adaptive radiotherapy for head and neck cancer - dosimetric evaluation of first clinical. Acta Oncol 2023;62:1369–74. https://doi.org/10.1080/0284186X.2023.2256966.

[9] Lavrova E, Garrett MD, Wang YF, Chin C, Elliston C, Savacool M, et al. Adaptive radiation therapy: a review of CT-based techniques. Radiol Imaging Cancer 2023;5:e230011. https://doi.org/10.1148/rycan.230011.

[10] O'Hara CJ, Bird D, Al-Qaisieh B, Speight R. Assessment of CBCT-based synthetic CT generation accuracy for adaptive radiotherapy planning. J Appl Clin Med Phys 2022;23:e13737. https://doi.org/10.1002/acm2.13737.

[11] Brock KK. Adaptive radiotherapy: moving into the future. Semin Radiat Oncol 2019;29:181–4. https://doi.org/10.1016/j.semradonc.2019.02.011.

[12] Wahid KA, Sahlsten J, Jaskari J, Dohopolski MJ, Kaski K, He R, et al. Harnessing uncertainty in radiotherapy auto-segmentation quality assurance. Phys Imaging Radiat Oncol 2023;29:100526. https://doi.org/10.1016/j.phro.2023.100526.

[13] Allen C, Yeo AU, Hardcastle N, Franich RD. Evaluating synthetic computed tomography images for adaptive radiotherapy decision making in head and neck cancer. Phys Imaging Radiat Oncol 2023;27:100478. https://doi.org/10.1016/j.phro.2023.100478.

[14] Taasti VT, Hattu D, Peeters S, van der Salm A, van Loon J, de Ruysscher D, et al. Clinical evaluation of synthetic computed tomography methods in adaptive proton therapy of lung cancer patients. Phys Imaging Radiat Oncol 2023;27:e100459. https://doi.org/10.1016/j.phro.2023.100459.

[15] Huynh E, Hosny A, Guthier C, Bitterman SF, Hass-Kogan DA, Kann B, et al. Artificial intelligence in radiation oncology. Nat Rev Clin Oncol 2020;17:771–81. https://doi.org/10.1038/s41571-020-0417-8.

[16] van de Schoot AJ, Hoffmans D, van Ingen KM, Simons MJ, Wiersma J. Characterization of Ethos therapy systems for adaptive radiation therapy: A multi-machine comparison. J Appl Clin Med Phys 2023;24:e13905. https://doi.org/10.1002/acm2.13905.

[17] Varian Ethos, https://www.varian.com; 2024 [accessed 24 February 2024].

[18] Avgousti R, Antypas C, Armpilia C, Simopoulou F, Liakouli Z, Karaiskos P, et al. Adaptive radiation therapy: when, how and what are the benefits that literature provides? Cancer Radiother 2022;26:622–36. https://doi.org/10.1016/j.canrad.2021.08.023.

[19] Bahig H, Yuan Y, Mohamed ASR, Brock KK, Ng SP, Wang J, et al. Magnetic resonance-based response assessment and dose adaptation in human papilloma virus positive tumors of the oropharynx treated with radiotherapy (MR-ADAPTOR): An R-IDEAL stage 2a-2b/Bayesian phase II trial. Clin Transl Radiat Oncol 2018;13:19–23. https://doi.org/10.1016/j.ctro.2020.11.008.

[20] Castelli J, Thariat J, Benezery K, Hasbini A, Gery B, Berger A, et al. Weekly adaptive radiotherapy vs standard intensity-modulated radiotherapy for improving salivary function in patients with head and neck cancer: A phase 3 randomized clinical trial. JAMA Oncol 2023;9:1056–64. https://doi.org/10.1001/jamaoncol.2023.1352.

[21] Westerhoff JM, Daamen LA, Christodouleas JP, Blezer ELA, Choudhury A, Westley RL, et al. Safety and tolerability of online adaptive high-field magnetic resonance–guided radiotherapy. JAMA Netw Open 2024;7:e2410819. https://doi.org/10.1001/jamanetworkopen.2024.10819.

[22] Heukelom J, Kantor ME, Mohamed ASR, Elhalawani H, Kocak-Uzel E, Lin T, et al. Differences between planned and delivered dose for head and neck cancer, and their consequences for normal tissue complication probability and treatment adaptation. Radiother Oncol 2020;142:100–6. https://doi.org/10.1016/j.radonc.2019.07.034.

[23] Feinberg EA, Shwartz A. Handbook of Markov decision processes: methods and applications. New York: Springer; 2002.

[24] Puterman ML. Markov Decision processes: discrete stochastic dynamic programming. New York: Wiley-Interscience; 2005.

[25] Beck JR, Pauker SG. The Markov process in medical prognosis. Med Decis Making 1983;3:419–58. https://doi.org/10.1177/0272989X8300300403.

[26] Kuntz KM, Russell LB, Owens DK, Sanders GD, Trikalinos TA, Salomon JA. Decision models in cost-effectiveness analysis. In: Neumann PJ, editors. Cost-Effectiveness in Health and Medicine, New York: Oxford Academic; 2016, p. 105–36. https://doi.org/10.1093/acprof:oso/9780190492939.003.0005.

[27] Ng SP, Ajayi T, Schaefer AJ, Pollard 3rd C, Bahig H, Garden AS, et al. Surveillance imaging for patients with head and neck cancer treated with definitive radiotherapy: a partially observed Markov decision process model. Cancer 2020;126:749–56. https://doi.org/10.1002/cncr.32597.

[28] Alagoz O, Hsu H, Schaefer AJ, Roberts MS. Markov decision processes: A tool for sequential decision making under uncertainty. Med Decis Making 2010;30:474–83. https://doi.org/10.1177/0272989X09353194.

[29] Alagoz O, Maillart LM, Schaefer AJ, Roberts MS. The optimal timing of living-donor liver transplantation. Manage Sci 2004;50:1420–30. https://doi.org/10.1287/mnsc.1040.0287.

[30] Denton BT, Kurt M, Shah ND, Bryant SC, Smith SA. Optimizing the start time of statin therapy for patients with diabetes. Med Decis Making 2009;29:351–67. https://doi.org/10.1177/0272989X08329462.

[31] Shechter SM, Bailey MD, Schaefer AJ, Roberts MS. The optimal time to initiate HIV therapy under ordered health states. Oper Res 2008;56:20–33. https://doi.org/10.1287/opre.1070.0480.

[32] MATLAB version: 9.13.0 (R2022b), Natick, Massachusetts: The MathWorks Inc.; 2022.

[33] Zhang L, Dong L, Court L, Wang H, Gillin M, Mohan R. TU-EE-A4-05: Validation of CT-assisted targeting (CAT) software for soft tissue and bony target localization. Med Phys 2005;32:2106. https://doi.org/10.1118/1.1998440.

[34] Maki RG, D'Adamo DR, Keohan ML, Saulle M, Schuetze SM, Undevia SD, et al. Phase II study of sorafenib in patients with metastatic or recurrent sarcomas. J Clin Oncol 2009;27:3133–40. https://doi.org/10.1200/JCO.2008.20.4495.

[35] Stieb S, Lee A, van Dijk LV, Frank S, Fuller CD, Blanchard P. NTCP modeling of late effects for head and neck cancer: A systematic review. Int J Part Ther 2021;8:95–107. https://doi.org/10.14338/20-00092.

[36] Hosseinian S, Hemmati M, Dede C, Salzillo TC, van Dijk LV, Mohamed ASR, et al. Cluster-based toxicity estimation of osteoradionecrosis via unsupervised machine learning: Moving beyond single dose-parameter normal tissue complication probability by using whole dose-volume histograms for cohort risk stratification. Int J Radiat Oncol Biol Phys 2024;119:1569–78. https://doi.org/10.1016/j.ijrobp.2024.02.021.

[37] Gan Y, Langendijk JA, van der Schaaf A, van den Bosch L, Oldehinkel E, Lin Z, et al. An efficient strategy to select head and neck cancer patients for adaptive radiotherapy. Radiother Oncol 2023;186:e109763. https://doi.org/10.1016/j.radonc.2023.109763.

[38] Gan Y, Langendijk JA, Oldehinkel E, Lin Z, Both S, Brouwer CL. Optimal timing of re-planning for head and neck adaptive radiotherapy. Radiother Oncol 2024;194:110145. https://doi.org/10.1016/j.radonc.2024.110145.

[39] Heukelom J, Hamming O, Bartelink H, Hoebers F, Giralt J, Herlestam T, et al. Adaptive and innovative radiation treatment for improving cancer treatment outcome (ARTFORCE); a randomized controlled phase II trial for individualized treatment of head and neck cancer. BMC Cancer 2013;13:84. https://doi.org/10.1186/1471-2407-13-84.

[40] Leeuw AD, Giralt J, Tao Y, Benavente S, Nguyen TVF, Hoebers FJP, et al. Acute toxicity in ARTFORCE: A randomized phase III dose-painting trial in head and neck cancer. Int J Radiat Oncol Biol Phys 2022;114:S98. https://doi.org/10.1016/J.IJROBP.2022.07.519.

[41] Mohamad I, Karam I, El-Sehemy A, Abu-Gheida I, Al-Ibraheem A, Al-Assaf H, et al. The evolving role of stereotactic body radiation therapy for head and neck cancer: Where do we stand? Cancers (Basel) 2023;15:e5010.

[42] Lalonde A, Bobić M, Sharp GC, Chamseddine I, Winey B, Paganetti H. Evaluating the effect of setup uncertainty reduction and adaptation to geometric changes on normal tissue complication probability using online adaptive head and neck intensity modulated proton therapy. Phys Med Biol 2023;68:e115018. https://doi.org/10.1088/1361-6560/acd433.

[43] Yap LM, Jamalludin Z, Ng AH, Ung NM. A multi-center survey on adaptive radiation therapy for head and neck cancer in Malaysia. Phys Eng Sci Med 2023;46:1331–40. https://doi.org/10.1007/s13246-023-01303-x.

[44] Bertholet J, Anastasi G, Noble D, Bel A, van Leeuwen R, Roggen T, et al. Patterns of practice for adaptive and real-time radiation therapy (POP-ART RT) part II: Offline and online plan adaption for interfractional changes. Radiother Oncol 2020;153:88–96. https://doi.org/10.1016/j.radonc.2020.06.017.