Radiotherapy and Oncology 134 (2019) 96–100

Contents lists available at ScienceDirect

Radiotherapy and Oncology

journal homepage: www.thegreenjournal.com

Original Article

Towards optimal stopping in radiation therapy

Ali Ajdari [a], Maximilian Niyazi [b], Nils Henrik Nicolay [c], Christian Thieke [b], Robert Jeraj [d], Thomas Bortfeld [a,*]

[a] Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, USA; [b] Department of Radiation Oncology, Ludwig-Maximilians University, Munich, Germany; [c] Department of Radiation Oncology, University of Freiburg Medical Center; and [d] Department of Medical Physics, University of Wisconsin, Madison, USA

ARTICLE INFO

Article history:
Received 5 October 2018
Received in revised form 20 December 2018
Accepted 10 January 2019
Available online 6 February 2019

Keywords:
Optimal stopping
OSRT
Predictive biomarker
Bayesian learning
Mathematical optimization
Adaptive radiotherapy

ABSTRACT

A typical fractionated radiotherapy (RT) course is a long and arduous process, demanding significant financial, physical, and mental commitments from patients. Each additional session of RT significantly increases the physical and psychological burden on patients and leads to higher radiation exposure in organs-at-risk (OAR), while, in some cases, the therapeutic benefits might not be high enough to justify the risks. Today, through technological advancements in molecular biology, imaging, and genetics more information is gathered about individual patient response before, during, and after the treatment. we highlight some of the ways that mathematical tools can help assess treatment efficacy on the fly, adapt the treatment plan based on individual biological response, and optimally stop the treatment, if necessary. We term this "Optimal Stopping in RT (OSRT)", after a similar concept in the fields of dynamic programming and Markov decision processes. In short, OSRT can dynamically determine "whether, when and how" to stop the treatment to improve therapeutic ratios.

A typical fractionated radiotherapy (RT) course is a long and arduous process, demanding significant financial, physical, and mental commitments from patients. RT plans are usually delivered over the course of several weeks, and in most cases have long-term and even life-long consequences on patients' quality of life after treatment. Each additional session of RT significantly increases the physical and psychological burden on patients and leads to higher radiation exposure in organs-at-risk (OAR), while, in some cases, the therapeutic benefits might not be high enough to justify the risks. Despite this, most RT plans remain "temporally frozen", that is, once initialized, they are administered until treatment completion with minimal modifications to the original plan. Moreover, plan efficacy is currently assessed in *a posterior* manner during regular follow-ups after the plan has been delivered. Today, through technological advancements in molecular biology, imaging, and genetics more information is gathered about individual patient response before, during, and after the treatment. Although many studies [2,15] have discussed the potential of such information for a more personalized RT planning, discussions about *how* such information should be incorporated in the decision-making for plan adaptation have been extremely rare.

In this editorial, we highlight some of the ways that mathematical tools can help assess treatment efficacy *on the fly*, adapt the treatment plan based on individual biological response, and optimally stop the treatment, if necessary. In particular, we view the planning problem not as a static, but a dynamic sequential decision making problem, with the option of stopping the treatment at any given fraction during the therapy course. We term this "Optimal Stopping in RT (OSRT)", after a similar concept in the fields of dynamic programming (DP) and Markov decision processes (MDP). In short, OSRT can dynamically determine *whether, when* and *how* to stop the treatment to improve therapeutic ratios.

Optimal stopping: classical definition

Optimal stopping has its roots in the fields of stochastic processes and dynamic programming, with a wide range of real-world applications such as asset selling, gambling, job searching (also known as secretary problem), and sequential hypothesis testing. In a classical Markovian definition of the optimal stopping [20], the system evolves stochastically over a set of "states" during the "stages" of the planning horizon. At each stage, the decision-maker receives a stochastic signal/quote from outside which signifies which state the system is in. This signal could be the random offer he receives from the buyer for his asset, the resume' from a job candidate, or a card drawn from the deck in blackjack. Based on this signal, the decision-maker decides whether to stop or continue the decision-making process. If he decides to stop, he

* Corresponding author at: Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, USA.
E-mail addresses: aajdari@mgh.harvard.edu (A. Ajdari), Maximilian.Niyazi@med.uni-muenchen.de (M. Niyazi), nils.nicolay@uniklinik-freiburg.de (N.H. Nicolay), Christian.Thieke@med.uni-muenchen.de (C. Thieke), rjeraj@wisc.edu (R. Jeraj), tbortfeld@mgh.harvard.edu (T. Bortfeld).

https://doi.org/10.1016/j.radonc.2019.01.010
0167-8140/

receives an immediate reward and the process is concluded, otherwise the system stochastically evolves into the next state and the process is repeated at the subsequent stage. The termination criteria is met when he decides to stop, or when the system reaches its terminal stage (assuming a finite horizon model). The goal in optimal stopping is to maximize the overall expected payoff or minimize expected cost.

Tumor/OAR kinetics and their response to radiation is in essence a stochastic process which is governed by a wide range of extrinsic and intrinsic factors, including, but not limited to, clinicopathological factors (e.g., age, sex, tumor histology, lymphovascular invasion, etc.), intrinsic and acquired radiosensitivity of tumor cells [10], and immune system response to radiation, including the attraction of both cytotoxic and regulatory T cells into the tumor microenvironment [33]. This view of the optimal stopping lends itself well to a dynamic version of treatment planning in which at each fraction, the planner has the option to stop or continue the treatment based on some measure of tumor/OAR response. Although to this date no strong and consistently validated biomarkers of individual treatment response have been identified, using carefully chosen biomarkers that have been shown to carry 'some' predictive ability - keeping in mind their inherent uncertainty - could provide opportunity to adapt the treatment as tumor biomarker expressions evolve [4,18,27]. In an OSRT view of the problem, the changing status of such biomarkers send regular signals to the planner. These signals can then be interpreted as 'rough' estimates of tumor/OAR states during treatment. Depending on each estimated state, the planner decides whether to stop/adapt the plan, or to continue with the original one.

In OSRT, the decision is based on two integral metrics: tumor control probability (TCP) and normal tissue complication probability (NTCP). Over the course of the treatment, the planner receives information (signals) about tumor/OAR response through serial biomarkers and evaluates the "updated" values of TCP and NTCP. The plan might be adapted or even stopped, if the probability of tumor control cannot justify the excessive damage to OAR (i.e., if TCP-to-NTCP ratio is very low).

OSRT in plan adaptation

Several clinical trials and research groups are currently investigating the efficacy of 'selective' dose (de) escalation approaches (see for examples clinical trials NCT03323463 and NCT03396718 in ClinicalTrials.gov). The goal in these studies is to diverge from the population-average "one-size-fits-all" guidelines by modifying prescribed dose for the more radiosensitive/radioresistant patients in the cohort. However, several important hurdles still remain on the way of a true plan personalization. First, these studies rely on an *a priori* patient stratification based on "prognostic" biomarkers of outcome, which does not allow for interfractional plan adaptation, although it has been shown that, for instance, tumor radioresistance is subject to temporal change during treatment [2]. Moreover, they mostly use a binary dose stratification to distinguish potential candidates from the rest of the population; that is, plans are dichotomized into two *pre-determined* groups: (de) escalated and baseline, without further optimizing the dose level for patients in each group. Furthermore, rather than a true individualized scheme, such (de) escalation approaches are group-based (poor/favorable responders and baseline group), and therefore ignore inter-patient heterogeneity that exist within each group. Lastly, patient stratification and subsequent adaptation is often based on a single biomarker, which given the inherent uncertainty associated with it, is unlikely to give an accurate and well-rounded picture of the state of the system.

There are few exceptions that have performed dose-escalation based on mid-treatment response evaluation [13,29]. However, the results remain inconclusive and often contradictory, which, at least partially, could be due to the sub-optimal nature of dose adaptation in these studies.

In OSRT, patient stratification is not temporally frozen, but can be re-evaluated and re-done mid-treatment as new information is revealed. In particular, the planner starts off by assuming a prior belief about parameters of TCP and NTCP in each cohort, which could be based on population-wide estimates or some pre-treatment biomarker expressions. Bayesian learning [14] can be used to update the prior and obtain posterior belief on these parameters. These updated values are then fed into the OSRT framework to update the values of TCP and NTCP, which in turn are used to re-evaluate the stratification method and re-optimize the plan (Fig. 1).

Dose modification in OSRT could be done over a *dose continuum* rather than in a binary stratification. The field of radiation therapy has long been used to spatial and temporal optimization of dose using models of dose–response such as the linear-quadratic (LQ) model. However, such models are mostly being used for relatively minor adjustments of the dose fractionation scheme, based on rough estimates of the linear-quadratic parameter α/β for the patient population. Augmenting such models with patient-specific information will lead to a finer dose adaptation and give the planner the ability to further optimize the dose for each individual patient, depending on his/her specific response [1,21,25]. However, before such bio-mechanistic models are realized in practice, reliable *predictive* biomarkers with interpretable mid-treatment signals should be available (see Forker et al. [9] for an overview of some of the potentially relevant predictive biomarkers that could be of use in OSRT). Moreover, given the fact that the

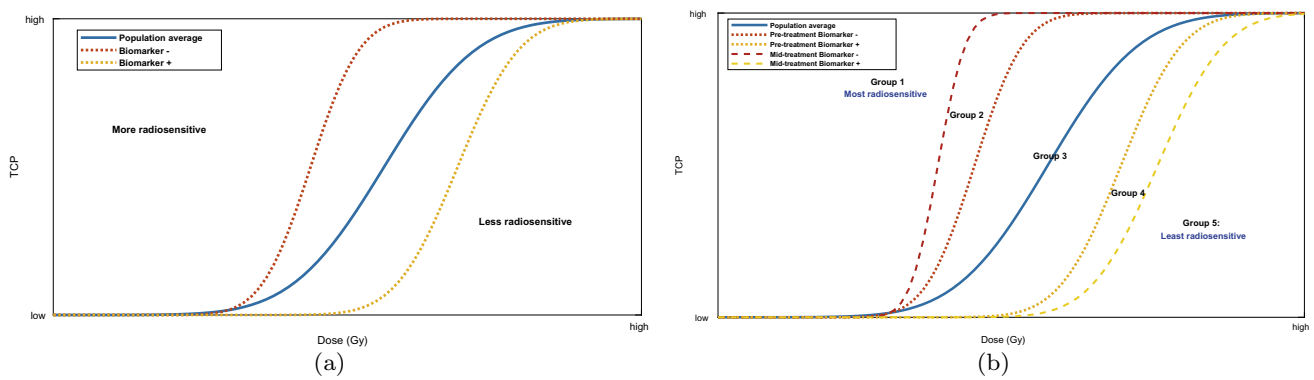

Fig. 1. Schematic illustration of OSRT for plan adaptation. (a) a priori patient stratification based on a pre-treatment biomarker; (b) combining pre- and mid-treatment biomarker for a finer patient stratification. Denoted groups represent the radiosensitivity spectrum across patient population, from the most (Group 1) to the least (Group 5) radiosensitive group.

majority of the historical data comes from a rather narrow spectrum of prescribed dose and fractionation schemes that have not been adapted mid-treatment, one has to be careful when assessing the potential benefit of the OSRT-guided adapted schemes that might fall outside of that spectrum. To what extent the current bio-mechanical models can represent the actual dose–response relationship and patients' outcome has to be validated in prospective clinical trials. In any case, using such models in pre-clinical animal studies would be an appropriate first step.

In plan adaptation, one important component is the timing of the intervention. Premature adaptation, before the true effect of radiation can reveal itself in the biomarker expressions might jeopardize the quality of the information, and therefore, adaptation itself. Similarly, waiting too long before adapting the treatment will most likely result in limited payoff, as the modified plan would not have enough 'room' to make an impact. Therefore, finding the optimal time for plan adaptation is a key factor in determining its effectiveness. Modeling OSRT as a *machine repair problem* (MRP) can be a potential solution. MRP is an instance of dynamic programming in which the decision maker at each stage has to decide whether to *continue* operating a given machine which is suspected to be in a 'poor' condition (and risk machine failure), or to *stop* and perform an extensive diagnostics inspection of the machine. Similarly, in dynamic adaptation of RT, the planner receives a signal about the tumor/OAR state and decides to stop and assess the risk of continuing the treatment more carefully, and possibly adapt the treatment if the risk is deemed too high.

OSRT in treatment stopping

It is important to note what is meant by "stopping" in this context. From a therapeutic standpoint, stopping refers to the discontinuation of the primary mode of therapy which has been deemed ineffective, and transitioning to a more effective mode. In this sense, OSRT could also be regarded as a 'radical' form of treatment adaptation in which not only the RT plan, but the primary therapy mode itself could be changed or adapted. The changes in the original RT plan could be in the form of switching to another modality such as surgery or concurrent chemo-radiation, using radiosensitizers to overcome acquired hypoxia/radioresistance [31], and administering radioprotectants or immuno-modulatory agents [11] to fight toxicity or enhance tumor control. An early example of OSRT can be seen in a phase II trial of esophageal cancer, in which mid-treatment fludeoxyglucose ([18]F-FDG) PET scans were used for patient stratification. Subsequently, RT treatment of good responders was continued as planned, while the plan was stopped in the poor responder group in favor of adjuvant surgery [16].

In OSRT, the planner compares the immediate reward of stopping in terms of sparing normal tissues, with the expected payoff of treatment continuation in terms of probability of tumor control. Therefore, the natural competition of NTCP vs. TCP that underlies any RT planning plays a key role in determining when and how to stop the treatment. In short, treatment is stopped if the projected damage to NTCP using current plan is too high to justify the expected TCP; i.e., when the *updated* therapeutic window [12] (i.e., TCP/NTCP ratio) is too narrow, or is shrinking too rapidly, to justify continuation of the treatment. Dose de-escalation in order to spare OAR can also be contemplated as a proper intervention. In the case of widening therapeutic window, a dose-escalation with the goal of improving TCP might be preferable. If it is believed that the desirable level of TCP can be achieved with less dose than initially planned, early stopping might also be entertained (Fig. 2). Thompson sampling [26], which is a method for deciding between exploiting immediate performance and investing to accumulate new information, could be a potential solution.

OSRT could also be used in palliative treatment, in which the decision to stop can be interpreted as finding the time to transition from a curative treatment to a palliative one. Receiving intensive curative treatment near the end of life is generally considered as an indicator of poor quality care [17]. Unnecessarily protracting treatment course with little to no effect in tumor control will lead to losing invaluable time that could have otherwise been spent on providing better end-of-life care, pain relief, and symptom control. OSRT could be a valuable tool to decide when the probability of tumor control no longer justifies the long-term RT-induced toxicities.

Another aspect is the patients' preferences before and during the radiation treatment course. Although patients are typically consulted by their treatment team on high-level decisions, such as the choice of surgery or concurrent chemotherapy, they usually are not involved in the decision making once the radiation treatment begins. However, several studies have shown that physicians consistently under-report and under-recognize patients' toxicities compared to those reported by the patients themselves [6,30]. In another study, patient-reported outcome (PRO) was shown to be a valid early indicator of late radiation-induced toxicities [34]. Taking into account patients' preferences might also increase the compliance rate and decrease the risk of premature treatment termination due to patients' decreasing morale, especially in a palliative setting. Therefore, finding ways to quantifiably assess and incorporate these preferences is a goal in OSRT. One simple way to incorporate such information is through PRO-enhanced NTCP models. Currently, most NTCP models only include dosimetric factors in predicting toxicity outcomes for a given plan. Including PRO

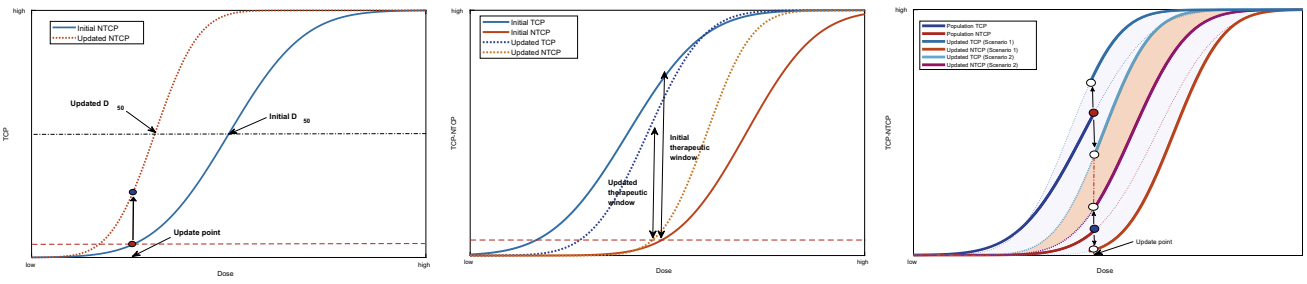

Fig. 2. A schematic illustration of optimal stopping in radiotherapy. (a) Updated NTCP (red) after observing some measure of tumor response at the "Update Point" and updating NTCP parameter D_{50}. (b) Shrinking therapeutic window after updating both TCP and NTCP (dotted lines) via mid-treatment biomarkers. The role of OSRT is to find the optimal threshold for the therapeutic window, beyond which the treatment should be stopped. (c) Change in the therapeutic window in two scenarios: Sensitive tumor cells and resistant OAR result in widening therapeutic window (Scenario 1), while the reverse leads to narrowing therapeutic window (Scenario 2).

scores as an additional factor in these models and updating the NTCP based on dynamically changing PRO scores assessed throughout the treatment course (e.g. using electronic devices like smartphones [8]) could be one way of doing so. Setting PRO-based constraints such as maximum number of fractions in the planning phase [24] can also be a viable option, particularly in palliative cases.

Information quality

The field of radiation oncology has long been used to addressing different sources of uncertainties during patient planning, from range and setup uncertainties [28] to organ movements [5] and clinical target volume (CTV) definition [22]. Since the implementation of OSRT in routine planning crucially depends on the quality of data obtained via biomarkers, and given that the interpretation of such data inevitably carries some level of uncertainty, a brief overview of some potential mathematical tools for addressing information uncertainties might be warranted.

Partially observable Markov decision processes (POMDP): These models take a "model-free" Markovian view of the RT problem in which tumor/OAR kinetics are assumed to form a Markovian process (i.e., current state only depends on the previous state and the action taken at that state). However, since the actual states (e.g., number of clonogenic tumor cells) are extremely hard to be directly observed via biomarkers, only partial information is available to treatment planner at each stage. POMDP models could help with relating this partial information to the true (unknown) state of the tumor, and cushioning against uncertainties that arise from imperfect biomarker information.

Imperfect state information (ISI): Another method for tackling the imperfect nature of the information is to model the problem as an ISI problem, in which instead of assuming a deterministic (and unknown) state for the tumor/OAR at each stage, the planner maintains a *probabilistic belief* about the state of the system. As new information becomes available, he then updates this belief using Bayesian learning, Q-learning [32], or other reinforcement learning methods, and continue the planning process until termination.

Adjustable robust optimization (ARO): A common and well-accepted tool to address model uncertainty during planning is to resort to robust optimization (RO) techniques. These methods, although effective in protecting against uncertainty, are often too conservative, that is, their "robustness guarantees" often come with large sacrifices in therapeutic objectives. ARO [3] is an extension of RO for multi-period robust optimization problems in which information about the uncertain parameters is revealed to the planner as the system progresses. In ARO, the robust model "knows" about the incoming information in future stages, and thus takes into account this knowledge when planning the "here-and-now" decisions. In doing so, ARO solutions tend to be less conservative and more data-driven. Such an approach could be particularly useful when dealing with multi-stage information acquisition through serial biomarkers.

Bayesian networks (BN): As any single biomarker is extremely unlikely to give a clear picture of the entire underlying state of the system, there might be value in combining multiple 'complimentary' biomarkers, each shedding light on different mechanisms of treatment response. This will lead to finer patient stratification and more tailored plan adaptation. Multi-variable prediction models such as dynamic Bayesian networks that are able to model the inter-dependence of these biomarkers could help combine different sources of information and potentially lead to models that are more robust under information uncertainty.

Toward clinical implementation of OSRT

Before OSRT can be safely and reliably implemented in a clinical setting, a number of steps have to be taken. One major step concerns biomarkers acquisition as well as interpretation and implementation of the information they provide. The biomarkers should provide maximal information combined with minimal associated uncertainty and noise. To achieve this, one can strive to meet the following basic characteristics as much as possible:

- The biomarker has to reflect the underlying biological state and clinically relevant effects without substantial delays. This means in particular that its magnitude should have the ability to change during the course of radiotherapy. It must be able to reflect changes in the order of days or weeks; otherwise, a simple pre-stratification of the patients based on the biomarker value would suffice and there would be no need for mid-treatment interventions.
- It has to be practical; i.e., relatively cheap to obtain and evaluate, and minimally invasive. This is especially important in cases where we are uncertain about the ideal timepoint for assessment, or wish to re-evaluate OSRT at multiple time points during treatment. Furthermore, the time from acquisition to interpretation must be relatively short to allow for timely mid-treatment intervention.
- It has to be reliable. This means that the biomarker must exhibit minimal 'random' error, and its measurements be reproducible across time and patients. These can be assessed e.g. through the so-called "test-retest stability" and"intra-rater reliability" tests.
- It has to have high validity, requiring a good mixture of specificity and sensitivity that must be validated in large clinical trials. This ensures confidence in interpreting the information acquired by the biomarker, and help justify its clinical role in treatment adaptation.

In addition to the mathematical tools mentioned before, OSRT heavily relies on advanced technologies for data acquisition and processing. Most of these technologies such as advanced imaging techniques and liquid biopsy assays currently exist; and some are actively being used in clinical settings, at least in clinical trials. However, the biomarkers need better validation and qualification before their information can be safely used in practice. Moreover, high associated cost and low information quality are among the major hindrances that have to be resolved before the implementation of these techniques become part of the treatment routine.

There is also an urgent need for developing more dynamic and biologically-oriented TCP and NTCP models capable of reflecting the temporal and biological evolution of tumor and OAR states throughout the treatment course and beyond. In the past decade there have been some efforts to include more biological factors in these models [7,19,23]; however, more works need to be done to fully explore this potentially fruitful area. Lastly, to truly evaluate the efficacy of OSRT, and any adaptive therapy, specific clinical trials need to be designed to prospectively collect data and compare treatment outcomes of patients receiving these interventions. Before this can be realized, pre-clinical studies, either using animal models or retrospective data could be used to demonstrate a proof-of-concept. Currently, a wealth of data are already available on different treatment adaptation schemes using different tumor response markers. Our team is currently pursuing two such retrospective studies of treatment adaptation: one study focuses on adapting treatment for non-small cell lung cancer (NSCLC) patients using mid-treatment ^{18}F-FDG PET, and the other on treatment modification on a head-and-neck canine dataset for which $^{18}F-FDG, ^{18}F-FLT, ^{64}Cu-ATSM$ PET images are available as

biomarkers. Such retrospective data would also be helpful, and indeed necessary, in developing the aforementioned 'predictive' TCP/NTCP models.[1]

Conclusions

Mathematical and optimization techniques such as the ones mentioned in this paper should be viewed as powerful tools with great potential to guide personalized treatment strategies in radiation therapy and other cancer therapies. Without such tools for optimal implementation of patient-specific information, even the most sophisticated technologies in response assessment and plan delivery would fall short of fulfilling their true potentials. Therefore, it appears to be prudent to bring mathematicians and cancer researchers and clinicians together to explore this largely untapped resource.

References

[1] Ajdari Ali, Ghate Archis, Kim Minsun. Adaptive treatment-length optimization in spatiobiologically integrated radiotherapy. Phys Med Biol 2018;63: 075009.
[2] Baumann M, Krause M, Overgaard J, et al. Radiation oncology in the era of precision medicine. Nat Rev Cancer 2016;16:234–50.
[3] Ben-Tal A, Goryashko A, Guslitzer E, et al. Adjustable robust solutions of uncertain linear programs. Math Program 2004;99:351–76.
[4] Bogowicz M, Riesterer O, Stark LS, et al. Comparison of PET and CT radiomics for prediction of local tumor control in head and neck squamous cell carcinoma. Acta Oncologica 2017;56:1531–6.
[5] Colgan R, McClelland J, McQuaid D, et al. Planning lung radiotherapy using 4D CT data and a motion model. Phys Med Biol 2008;53:5815–30.
[6] Di Maio Massimo, Gallo Ciro, Leighl Natasha B, Piccirillo Maria Carmela, Daniele Gennaro, Nuzzo Francesco, et al. Symptomatic Toxicities Experienced During Anticancer Treatment: Agreement Between Patient and Physician Reporting in Three Randomized Trials. J Clin Oncol 2015;33:910–5.
[7] El Naqa Issam, Johansson Adam, Owen Dawn, Cuneo Kyle, Cao Yue, Matuszak Martha, et al. Modeling of normal tissue complications using imaging and biomarkers after radiation therapy for hepatocellular carcinoma. Int J Radiation Oncol Biol Phys 2018;100:335–43.
[8] Falchook Aaron D, Tracton Gregg, Stravers Lori, Fleming Mary E, Snavely Anna C, Noe Jeanne F, et al. Use of mobile device technology to continuously collect patient-reported symptoms during radiation therapy for head and neck cancer: a prospective feasibility study. Adv Radiat Oncol 2016;1:115–21.
[9] Forker LJ, Choudhury A, Kiltie AE. Biomarkers of tumour radiosensitivity and predicting benefit from radiotherapy. Clin Oncol 2015;27:561–9.
[10] Gerweck Leo E, Vijayappa Shashirekha, Kurimasa Akihiro, et al. Tumor cell radiosensitivity is a major determinant of tumor response to radiation. Cancer Res 2006;66:8352–5.
[11] Hall Susan, Rudrawar Santosh, Zunk Matthew, et al. Protection against radiotherapy-induced toxicity. Antioxidants (Basel, Switzerland) 2016;5(3).
[12] Holthusen H. Erfahrungen über die Verträglichkeitsgrenze für Röntgenstrahlen und deren Nutzanwendung zur Verhütung von Schäden. Strahlentherapie 1936;57:254–69.
[13] Kong Feng-Ming, Ten Haken Randall K, Schipper Matthew, et al. Effect of midtreatment PET/CT-adapted radiation therapy with concurrent chemotherapy in patients with locally advanced non-small-cell lung cancer. JAMA Oncol 2017;3:1358–65.
[14] Kononenko I. Inductive and Bayesian learning in medical diagnosis. Appl Artif Intell 1993;7:317–37.
[15] Krause M, Dubrovska A, Linge A, et al. Cancer stem cells: radioresistance, prediction of radiotherapy outcome and specific targets for combined treatments. Adv Drug Deliv Rev 2017;109:63–73.
[16] Lordick Florian, Ott Katja, Krause B, et al. PET to assess early metabolic response and to guide treatment of adenocarcinoma of the oesophagogastric junction: the MUNICON phase II trial. Lancet Oncol 2007;8:797–805.
[17] Lutz S, Berk L, Chang E, et al. Palliative radiotherapy for bone metastases: an ASTRO evidence-based guideline. Int J Rad Oncol Biol Phys 2011;79:965–76.
[18] Madani Indira, De Ruyck Kim, Goeminne Hannelore, et al. Predicting risk of radiation-induced lung injury. J Thoracic Oncol 2007;2:864–74.
[19] Palmer Joshua D, Zaorsky Nicholas G, Witek Matthew, Bo Lu. Molecular markers to predict clinical outcome and radiation induced toxicity in lung cancer. J Thoracic Disease 2014;6:387.
[20] Puterman ML. Markov decision processes: discrete stochastic dynamic programming. NJ, Hoboken: Wiley-Interscience; 2014.
[21] Saberian Fatemeh, Ghate Archis, Kim Minsun. A theoretical stochastic control framework for adapting radiotherapy to hypoxia. Phys Med Biol 2016;61:7136–61.
[22] Shusharina N, Craft DL, Chen Y-L, et al. The clinical target distribution: a probabilistic alternative to the clinical target volume. Phys Med Biol 2018.
[23] Stenmark Matthew H, Cai Xu-Wei, Shedden Kerby, Hayman James A, Yuan Shuanghu, Ritter Timothy, et al. Combining Physical and Biologic Parameters to Predict Radiation-Induced Lung Toxicity in Patients With Non-Small-Cell Lung Cancer Treated With Definitive Radiation Therapy. Int J Rad Oncol, Biol, Phys 2012;84: e217.
[24] Tang JI, Shakespeare TP, Lu JJ, Chan YH, Lee KM, Wong LC, et al. Patients' preference for radiotherapy fractionation schedule in the palliation of symptomatic unresectable lung cancer. J Med Imaging Radiat Oncol 2008;52:497–502.
[25] ten Eikelder SCM. Adjustable robust treatment-length optimization in radiation therapy. Research master thesis. Tilburg University; 2018.
[26] Thompson WR. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika 1933;25:285.
[27] Tucker Susan L, Li Minghuan, Ting Xu, et al. Incorporating single-nucleotide polymorphisms into the lyman model to improve prediction of radiation pneumonitis. Int J Rad Oncol Biol Phys 2013;85:251–7.
[28] Unkelbach J, Chan TCY, Bortfeld T. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. Phys Med Biol 2007;52:2755–73.
[29] Vera Pierre, Thureau Sébastien, Chaumet-Riffaud P, et al. Phase II study of a radiotherapy total dose increase in hypoxic lesions identified by 18 F-misonidazole PET/CT in patients with non-small cell lung carcinoma (RTEP5 Study). J Nucl Med 2017;58:1045–53.
[30] Vera-Badillo FE, Shapiro R, Ocana A, Amir E, Tannock IF. Bias in reporting of end points of efficacy and toxicity in randomized, clinical trials for women with breast cancer. Ann Oncol 2013;24:1238–44.
[31] Wardman P. Chemical radiosensitizers for use in radiotherapy. Clin Oncol 2007;19:397–417.
[32] Watkins C, Dayan P. Technical Note Q.-Learning. Mach Learn 1992;8:279–92.
[33] Weichselbaum Ralph R, Liang Hua, Deng L, et al. Radiotherapy and immunotherapy: a beneficial liaison? Nat Rev Clin Oncol 2017;14(6):365–79.
[34] Yue Jinbo, Shi Qiuling, Ting Xu, Jeter Melenda, Chen Ting Yu, Komaki Ritsuko, et al. Patient-reported lung symptoms as an early signal of impending radiation pneumonitis in patients with non-small cell lung cancer treated with chemoradiation: an observational study. Qual Life Res 2018;27:1–8.

[1] There are a few ongoing clinical trials, mostly focusing on dose painting and (de) escalation schemes, that might partially meet these criteria; including clinical trials NCT01341535, NCT03376386, NCT01190527, NCT00406289, NCT02773238, and NCT01507428.