



Adjustable robust treatment-length optimization in radiation therapy

S. C. M. ten Eikelder¹ · A. Ajdari² · T. Bortfeld² · D. den Hertog³

Received: 12 February 2021 / Revised: 28 November 2021 / Accepted: 23 December 2021 /
Published online: 7 May 2022
© The Author(s) 2022

Abstract

Traditionally, optimization of radiation therapy (RT) treatment plans has been done before the initiation of RT course, using population-wide estimates for patients' response to therapy. However, recent technological advancements have enabled monitoring individual patient response during the RT course, in the form of biomarkers. Although biomarker data remains subject to substantial uncertainties, information extracted from this data may allow the RT plan to be adapted in a biologically informative way. We present a mathematical framework that optimally adapts the treatment-length of an RT plan based on the acquired mid-treatment biomarker information, while accounting for the inexact nature of this information. We formulate the *adaptive treatment-length optimization problem* as a 2-stage problem, wherein the information about the model parameters gathered during the first stage influences the decisions in the second stage. Using Adjustable Robust Optimization (ARO) techniques we derive explicit optimal decision rules for the stage-2 decisions and solve the optimization problem. The problem allows for multiple worst-case optimal solutions. To discriminate between these, we introduce the concept of Pareto Adjustable Robustly Optimal solutions. In numerical experiments using lung cancer patient data, the ARO method is benchmarked against several other static and adaptive methods. In the case of exact biomarker information, there is sufficient space to adapt, and numerical results show that taking into account both robustness and adaptability is not necessary. In the case of inexact biomarker information, accounting for adaptability and inexactness of biomarker information is particularly beneficial when robustness (w.r.t. organ-at-risk (OAR) constraint violations) is of high importance. If minor OAR violations are allowed, a nominal folding horizon approach (NOM-FH) is a good performing alternative, which can outperform ARO. Both the difference in performance and the magnitude of OAR violations of NOM-FH are highly influenced by the biomarker information quality.

Keywords Radiation therapy · Adjustable robust optimization · Decision rules · Optimal stopping

1 Introduction

In radiation therapy (RT), the goal is to deliver a curative amount of radiation dose to the target volume(s), while keeping the dose to all organs-at-risk (OARs) within tolerable limits. As the radiation beam delivers energy to all tissues that are on its path, the OARs will (often) inevitably receive some dose as well. Treatment planning has a spatial component, where the optimal dose distribution is determined, and a temporal component, where the optimal number of treatment sessions, or fractions, is determined. Whereas the former is predominantly a geometric problem, the latter involves radiobiological effects.

Technological advances in treatment monitoring through imaging and other forms of data acquisition allow for a more accurate assessment of an individual's radiation response (Baumann et al. 2016). Biologically-based adaptive treatments aim to monitor the treatment, acquire mid-treatment biomarker information, and adapt the remainder of the treatment course accordingly. Many approaches to adaptive treatment planning have been studied in the literature. To the best of our knowledge, all existing approaches assume that all information acquired mid-treatment is exact, i.e., gives a perfect representation of the current state of treatment response. Unfortunately, the limited availability and accuracy of required biomarkers pose a primary challenge for adaptive treatments (Baumann et al. 2016). Any information from biomarker data acquired during treatment remains subject to uncertainties, stemming from both measurement errors and the inexactness in the translation of measured data to usable model parameters. Therefore, it is crucial that any adaptive treatment planning approach takes this into account. Ajdari et al. (2019) provide an overview of relevant mathematical (optimization) tools. We present an approach to optimally adapt the treatment length of RT using inexact mid-treatment biomarker information.

Specifically, we take an adjustable robust optimization (ARO, Ben-Tal et al. 2004; Yanikoğlu et al. 2019) approach that accounts for the inexactness of biomarker information. ARO is an extension of robust optimization (RO) that takes into account the flow of information over time and exploits the fact that some decisions need to be taken only after the data has (partially) revealed itself. By using ARO, we ensure that the treatment plan delivered in the initial treatment stage (prior to obtaining biomarker information) is 'adaptation aware'. That is, it is designed with adaptation in mind, which may yield more flexibility at the time of treatment adaptation. In the standard paradigm, ARO assumes that the revealed information is exact; we employ the ARO methodology developed in De Ruiter et al. (2017) for the case when revealed information is not exact, but provides only an estimate of the true parameters.

1.1 Contributions

We consider a stylized two-stage ARO model to optimally adapt the treatment-length based on inexact biomarker information acquired mid-treatment. Although the stylized model makes several simplifying assumptions to aid the analysis, we

believe it captures several important characteristics of realistic adaptive treatment planning, and it enables a precise analysis of the influence of uncertainty in biomarker information. Our main contributions are:

- We develop mathematical tools based on ARO that enable us to (i) optimally adapt the dose per fraction and treatment length after acquiring mid-treatment biomarker information, (ii) analyze the influence of biomarker information uncertainty.
- We present explicit optimal decision rules for a difficult (nonconvex, mixed-integer) yet practically relevant ARO problem.
- We show that there are multiple optimal solutions for the worst-case scenario, and that these differ in performance in non-worst-case scenarios. To handle this, we introduce the concept of Pareto Adjustable Robustly Optimal (*PARO*) solutions, a generalization of Pareto Robustly Optimal (*PRO*) solutions (Iancu and Trichakis 2014) to two-stage robust optimization problems. In case the acquired biomarker information is exact, *PARO* solutions are obtained.
- We perform a computational study using real lung cancer patient data to determine the optimal timing of acquiring biomarker information in case biomarker quality improves over time. Later biomarker acquisition also limits adaptation possibilities, and the optimal balance depends on the improvement rate.

1.2 Literature review

There is a large body of adaptive treatment planning research in RT, the majority of which focuses not on biologically-based uncertainties but on geometric uncertainties and anatomical changes. Chan and Mišić (2013), Mar and Chan (2015), Böck et al. (2017), Böck et al. (2019) and Lim et al. (2020) present adaptive treatment planning approaches that start with delivery of the original treatment plan, often derived using RO methods. At given adaptation moments, the ‘state’ of the patient (e.g., anatomical changes, tumor shrinkage, breathing motion pattern) is observed, and the treatment is re-optimized for the remainder of the treatment plan. In RO terminology, these approaches are known as folding horizon (FH) methods. They disregard adaptation possibilities initially, and re-optimize the updated model once mid-treatment information is acquired.

Several treatment plan adaptation approaches have been proposed for adapting to biological information, differing in considered biomarker information, adapted treatment plan decisions and used methodology. Kim et al. (2012) and Ghate (2014) propose a theoretical stochastic control framework to optimally adapt the dose distribution over a fixed number of fractions, based on hypothetically-observed tumor states. Saberian et al. (2016b) concretize this theoretical framework, using simulated hypoxia (insufficient oxygen supply at cellular level) status as biomarker. Long (2015) consider a model with a constraint on the probability of radiation-induced lung toxicity, which depends on an a priori unknown model parameter. The problem is formulated as a two-stage model (before and after parameter observation), and the optimal dose distribution is determined for each stage. They consider a finite scenario set for the parameter, and

the lung toxicity constraint either has to hold in expectation or has to hold for all considered scenarios.

Nohadani and Roy (2017) consider a two-stage model to adapt to hypoxia information, where the uncertainty is time-dependent. As the hypoxia information ages the uncertainty grows, until it resets at the observation moment. In each stage the dose distribution is optimized; for the second stage a finite adaptability approach is taken. It is shown that total information degradation is minimized if the observation moment is set mid-treatment. Dabadghao and Roy (2020) consider a similar time-dependent uncertainty set, for adapting to hypoxia information in a multi-stage setting. In each stage the mean tumor dose per fraction is optimized using an FH approach. It is shown that total information degradation is minimized if the observation moments are set equidistant. They introduce a cost of observation (additional dose due to mid-treatment positron emission tomography (PET) scans), and determine the optimal number of observations. Both papers assume that the hypoxia state varies over time, and is exactly observed at the observation moment(s). In contrast, we assume that uncertain parameters are constant in time, and consider inexact biomarker information. Moreover, they solely adapt the dose, whereas we also adapt the total number of treatment fractions itself.

Adapting the treatment length (i.e., the fractionation schedule) based on observed radiation response has been studied before in the literature. Saka et al. (2014) consider a two-stage model where after a fixed number of treatment fractions hypoxia information is acquired. Based on this information both the remaining number of treatment fractions and the dose distribution are re-optimized, in order to maximize average hypoxia-corrected tumor dose. They focus on maintaining hypoxia-corrected fraction size requirements. A similar approach is taken by Ajdari et al. (2018), where after each treatment fraction the tumor cell density in each voxel is observed, and adaptations are made after each treatment fractions instead of only once. The objective is to minimize the total number of tumor cells remaining (TNTCR) at the end of the treatment course. Both approaches can be considered FH methods. In contrast to our approach, it is assumed that any information acquired mid-treatment is exact.

Iancu et al. (2021) propose a conceptual robust monitoring and stopping model. They consider a system with a state $x(t)$, and after each observation moment the uncertainty in the system state $x(t)$ grows as t increases. At a new observation moment, the uncertainty reduces to zero, i.e., an exact observation is made. They consider multiple observation moments, and the goal is to time these optimally. At each observation moment, the (state-dependent) direct stopping reward is compared to the worst-case continuation reward, and the according action is taken. Their model does not allow for controls that influence state variables, i.e., applying their model to RT optimization problems would not allow to adjust the dose distribution or the mean dose per fraction.

1.3 Notation and organization

All variables and constants are 1-dimensional (belong to \mathbb{R} or \mathbb{N}) unless indicated otherwise. In functions, a semicolon (;) is used to separate variables and constant arguments from uncertain parameters. Optimal solutions to optimization problems are indicated with an asterisk (*). Properties of optimal solutions to optimization

problems have calligraphic font (e.g., \mathcal{ARO}) to distinguish them from methods with the same or similar abbreviations.

The remainder of this paper is organized as follows. Section 2 introduces the used biological models, background information on biomarkers and states modeling choices. Section 3 introduces the adjustable treatment-length optimization problem under the assumption of exact biomarker information and solves this using ARO techniques. Section 4 generalizes this to inexact biomarker information. Section 5 presents and discusses results of numerical experiments on a lung cancer data set. Finally, Sect. 6 concludes the paper.

2 Adaptive fractionation

2.1 The fractionation problem

Spatial optimization exploits the fact that, by mounting the beam head on a gantry, the tumor can be targeted from various angles. It aims to find the combination of beam angles and weights that gives the best trade-off between tumor dose conformity and healthy tissue sparing. There is a large body of literature on this topic, see for example Shepard et al. (1999) and Ehr Gott et al. (2008) for reviews. The resulting dose distribution gives the dose (in Gray (Gy)) to each voxel (3-dimensional subvolume) of the tumor and OARs. Figure 1 gives an example of a slice of a dose distribution. The beam angles and weights are chosen such that the target (contoured in black) receives a high dose and a nearby organ-at-risk (OAR, contoured in red) receives a low dose.

Typically, this dose is not delivered in a single treatment session, but spread out over multiple treatment fractions (fx). The underlying idea is that compared to tumor cells, healthy tissues often have better repair capabilities between fractions (Fowler 1989; Withers 1985). On the other hand, a treatment spread out over a large number of treatment fractions may not deliver sufficient damage to the target volume, and increases the risk of tumor proliferation. The effect of fractionation differs per healthy tissue type and per tumor site, see, e.g., Hall and Giaccia (2012) for further details. Determining the optimal number of treatment fractions is known as the fractionation problem. Treatments with a higher number of fractions and a lower dose per fraction than the conventional regimen are known as *hyperfractionated* treatments. Treatments with a lower number of fractions and a higher dose per fraction than conventional are known as *hypofractionated* treatments.

In each treatment fraction, a scaled version of the dose distribution is administered. Scaling the dose distribution influences the absolute dose delivered to each voxel, but the relative dose remains unchanged. That is, for each voxel we can define a *dose sparing factor*: the dose of this voxel as a fraction of the mean target dose. In Fig. 1, the target volume receives a uniform dose of 45 Gy over the entire treatment course. The indicated OAR voxel receives 18 Gy and thus has a dose sparing factor of 0.4. Consequently, if in an individual treatment fraction we administer a mean target dose of 3 Gy, the OAR voxel receives 1.2 Gy. Thus, with a fixed dose distribution, the sole decision in treatment fraction t is d_t , the mean tumor dose in that




Fig. 1 Slice of a 3D dose distribution (illustration). Target is contoured in black, OAR is contoured in red. The target volume has a mean dose of 45 Gy. The indicated OAR voxel with dose 18 Gy has a dose sparing factor of 0.4.

fraction. The corresponding dose to any voxel i with dose sparing factor s_i is $s_i d_i$. Altogether, with a fixed dose distribution fractionation problems can be described using a low number of decision variables. Typically, the target dose is homogeneous, so for simplicity we assume that each target voxel receives dose d_t (i.e., has dose sparing factor 1). Nevertheless, we emphasize that this modeling approach allows for both heterogeneous dose distributions in target and OAR volumes.

The Biologically Effective Dose (BED) model (Fowler 1989, 2010) states that the biological effect of an N -fraction dose sequence $d = (d_1, \dots, d_N)$ delivered to a tumor volume is given by

$$\text{BED}_T = \sum_{t=1}^N d_t + \frac{1}{\alpha/\beta} \sum_{t=1}^N d_t^2, \quad (1)$$

which is a model governed by a single parameter, the α/β ratio, which signifies the fractionation sensitivity of the tumor tissue. The BED to the OAR can be described by

$$\text{BED}_{\text{OAR}} = \sum_{t=1}^N \sigma d_t + \frac{1}{\alpha/\beta} \sum_{t=1}^N \sigma^2 d_t^2, \quad (2)$$

where σ is the generalized dose-sparing factor. For the maximum BED in the OAR volume, σ is the maximum of the individual dose sparing factors s_i . In order to describe a mean BED constraint or dose-volume BED constraint other choices for σ can be used (Saberian et al. 2016a; Perkó et al. 2018).

For notational convenience, let τ (for tumor) and ρ (for risk) denote the inverse α/β ratio of the tumor and OAR volume, respectively. Mizuta et al. (2012) consider the problem of minimizing OAR BED subject to a lower bound $\text{BED}_T^{\text{pres}}$ on tumor BED. The number of fractions N is restricted to be at most N^{\max} . The problem reads

$$\min_{d,N} \sigma \sum_{t=1}^N d_t + \rho \sigma^2 \sum_{t=1}^N d_t^2 \quad (3a)$$

$$\text{s.t.} \quad \sum_{t=1}^N d_t + \tau \sum_{t=1}^N d_t^2 \geq \text{BED}_T^{\text{pres}} \quad (3b)$$

$$d_1, \dots, d_N \geq 0, \quad N \in \{1, \dots, N^{\max}\}. \quad (3c)$$

Let (d^*, N^*) denote the optimal solution to (3). A simple analysis in Mizuta et al. (2012) reveals the following important result:

$$N^* = \begin{cases} 1 & \text{if } \tau \geq \sigma \rho \\ N^{\max} \text{ and } d_1^* = \dots = d_{N^{\max}}^* & \text{otherwise.} \end{cases} \quad (4)$$

In both cases the optimal dose d^* is such that (3b) is active. The same result holds if we maximize tumor BED subject to an upper bound on OAR BED (Bortfeld et al. 2015), and a similar result has been derived for the case with multiple OARs (Sabirian et al. 2016a). There is a large body of research that optimizes the number of treatment fractions for different model formulations (see Sabirian et al. (2017) and references therein).

In the current paper, we restrict to one dose-limiting OAR. For many tumor sites, there is a single OAR that restricts the doses that can be delivered, and other OARs are much less restrictive. For example, for lung cancer the mean lung dose is an important indicator of toxicity. On the other hand, for head and neck cancer many OARs must be accounted for. We emphasize that other OARs are not completely disregarded. They are taken into account implicitly, because the original dose distribution was planned with all relevant OARs taken into consideration. Moreover, by restricting the minimum and maximum (mean target) dose per fraction, extreme deviations from the standard fractionation schedule are avoided, which is also designed whilst taking multiple OARs into account.

2.2 Adaptive fractionation using biomarkers

Most fractionation optimization methods assume the tumor and OAR fractionation sensitivity parameters τ and ρ are known exactly. There is much research on the α/β ratios for different tumor sites (Van Leeuwen et al. 2018) and OAR sites (Kehwar 2005), but they remain subject to considerable uncertainties. We assume box uncertainty of the form:

$$Z = \{(\rho, \tau) : \rho_L \leq \rho \leq \rho_U, \tau_L \leq \tau \leq \tau_U\}, \quad (5)$$

with $0 < \rho_L < \rho_U$ and $0 < \tau_L < \tau_U$. It is assumed that there is a nominal scenario, e.g., parameter values $\bar{\tau}$ and $\bar{\rho}$ derived from literature. There are two reasons for assuming a box uncertainty set. First, to the best of our knowledge there is little evidence for any correlation between the α/β ratios of target volumes and normal tissues. Second, box uncertainty leads to simpler models, which allow for a more detailed analysis of optimal fractionation decisions.

Ajdari and Ghate (2016) also consider a box uncertainty set, and determine a *robustly optimal* fractionation scheme, i.e., one that is that is feasible for all possible realizations and that is optimal for the worst-case realization. If biomarker information acquired during treatment provides more accurate information on fractionation sensitivity than what was available at the start of the treatment, such a static RO approach may be overly conservative.

Somaiah et al. (2015) give an overview of various mechanisms for determining fractionation sensitivity. Using blood samples, one can quantify the involvement of non-homologous end-joining (NHEJ) and homologous recombination (HR) in cells. For details on how to measure these, we refer to Bindra et al. (2013) and Barker and Powell (2010), respectively. Change in the expression of epidermal growth factor receptor (EGFR) genes can also give some hints regarding the fractionation sensitivity (Somaiah et al. 2015), which can also be measured mid-treatment. Lastly, Somaiah et al. (2015) mention that there is a close link between proliferation index and hypoxia, both of which can be measured during RT using different PET tracers. We note that there is evidence that some of these mechanisms could be subject to change during RT, depending on, amongst others, the delivered dose, hypoxia, and immune system interaction. However, as a first study to adapting to inexact biomarker information, we make the assumption that fractionation sensitivity is static throughout treatment, i.e., there is a static ‘true’ (ρ, τ) . In Section 3 we assume to observe (measure) the true (ρ, τ) exactly, and in Sect. 4 we assume to observe only an estimation/approximation $(\hat{\rho}, \hat{\tau})$.

The quality of the observed parameter estimates depends amongst others on the suitability of the biomarker itself, the measurement accuracy and when the biomarker measurement is taken during the treatment course. The relationship between the data quality and the moment of biomarker observation is complex, and it is impossible to exactly quantify this. For some biomarkers the data quality may greatly improve in the first few fractions, with a diminishing improvement in later fractions¹. For others (e.g., functional imaging such as PET and magnetic resonance imaging (MRI)), the data quality is poor at the first couple of fractions and

¹ This is especially true in the case of certain blood biomarkers of innate immune status (such as interleukin (IL)-6 or tumor necrosis factor (TNF)- α) which are also the markers of inflammation. In these biomarkers, as the biomarker acquisition is shifted towards later in the RT course, the information regarding the immune status gets mixed with the RT-induced inflammation and loses its specificity to immune system condition.

only increases substantially in later fractions². In practice, some biomarkers, e.g., radiographic information, may also exhibit a decreasing data quality for very late observation moments (due to, for instance, interference from acute inflammation in the lung). Such behavior is rare, and as such not considered here. We impose a minimum dose per fraction, and make the assumption that biomarker data quality increases in the number of treatment fractions. In this way, the change in biomarker quality is influenced by both the dose delivered and the time passed. We will investigate several functional forms for this relationship in the numerical experiments.

2.3 Modeling choices

In order to establish a meaningful model for the adjustable robust optimization approach, we restrict the dose sequence $d = (d_1, \dots, d_{N_{\max}})$ in several ways. In addition to a maximum number of fractions, we also set a minimum N^{\min} . Furthermore, we assume there is a single moment N_1 where we can adapt the treatment. Under the assumption that the uncertain parameters remain constant over time, more than one observation moment is not useful if the parameter is observed exactly. With inexact observations, there can be value in multiple observations, but given the patient burden³ and financial cost this is not considered here.

The dose per fraction in the first N_1 fractions is assumed to be the same, denote this by d_1 . Variable N_2 denotes the number of fractions after adaptation; also these fractions have equal dose, denoted by d_2 . In current clinical practice, uniform fractionated treatments are the standard. By restricting to only two different dose levels, extreme deviations from standard protocols are prevented. The above implies

$$N_2 \in \{N_2^{\min}, \dots, N_2^{\max}\}, \quad (6)$$

with $N_2^{\min} = \max\{1, N^{\min} - N_1\}$ and $N_2^{\max} = N^{\max} - N_1$. We additionally set the constraint that $d_1, d_2 \geq d^{\min}$, for some predetermined value d^{\min} . Aside of preventing an unrealistically low dose per fraction, the minimum dose serves a modeling purpose for stage 1. As noted in Sect. 2.2, the biomarker quality can depend on both dose and time. The current model implicitly makes the assumption that an early response can only be observed via biomarkers once N_1 fractions of dose at least d_{\min} have been delivered. Thus, in our models, this can be interpreted as a threshold. In the numerical experiments we investigate several temporal relationships between N_1 and biomarker quality. Lastly, we set a maximum dose per fraction d_1^{\max} in stage 1, to avoid delivering dose levels in stage 1 that severely restrict adaptation possibilities in stage 2. We will later impose some restrictions on the allowed combinations of d^{\min} , d_1^{\max} and N_2^{\max} .

Figure 2 provides a schematic overview of the situation.

² This is because the effect of RT on tissue is cumulative and is morphologically manifested only after a certain amount of dose (which depends on the underlying tissue threshold) is delivered.

³ Next to additional logistical hurdles for patients, imaging biomarkers such as PET/computed tomography (CT) scans also deliver small amounts of dose to the patient.




Fig. 2 Schematic overview of the considered model. There are 3 variables: d_1 , d_2 and N_2 . First, we deliver N_1 fractions of dose d_1 per fraction. After this, we observe (ρ, τ) or $(\hat{\rho}, \hat{\tau})$. Subsequently, we deliver N_2 fractions of dose d_2 per fraction

We wish to maximize the tumor BED, subject to the constraint that the BED to the OAR is below the generalized tolerance level BED_{tol} , given by

$$BED_{tol}(\rho) = \varphi D \left(1 + \frac{\varphi D}{T} \rho \right), \tag{7}$$

i.e., the OAR is known to tolerate a total dose of D Gy if delivered in T fractions under dose shape factor φ . The dose shape factor is a parameter characterizing the spatial heterogeneity of a dose distribution, for more details see Saberian et al. (2016a) and Perkó et al. (2018). Note that $BED_{tol}(\rho)$ is a function of uncertain parameter ρ , the inverse α/β ratio of the OAR.

We emphasize that the model resulting from our modeling choices and assumptions does not directly represent a realistic decision-making problem in radiation therapy treatment planning. Nevertheless, it captures several important aspects of fractionation optimization. Moreover, using ARO on this stylized model, we gain insight into optimal decision rules and the role of uncertainty in adaptive fractionation optimization.

3 ARO: biomarkers provide exact information

We present an adjustable robust optimization approach that optimally adjusts the remainder of the treatment once biomarker information has provided the true value of parameters τ and ρ . This serves as a stepping stone to the inexact data model.

3.1 Problem formulation

The Exact Data Problem EDP reads:

$$\max_{d_1, d_2(\rho, \tau), N_2(\rho, \tau)} \min_{(\rho, \tau) \in Z} N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau (N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \tag{8a}$$

$$\text{s.t.} \quad \sigma (N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau)) + \rho \sigma^2 (N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \leq BED_{tol}(\rho), \quad \forall (\rho, \tau) \in Z \tag{8b}$$

$$N_2(\rho, \tau) \in \{N_2^{\min}, \dots, N_2^{\max}\}, \quad \forall (\rho, \tau) \in Z \tag{8c}$$

$$d_2(\rho, \tau) \geq d^{\min}, \quad \forall(\rho, \tau) \in Z \quad (8d)$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \quad (8e)$$

The value for the stage-1 dose d_1 has to be decided before the value of (ρ, τ) is revealed; in ARO this is also commonly referred to as a *here-and-now* variable or decision. The values for the stage-2 dose d_2 and stage-2 number of fractions N_2 need to be decided only after (ρ, τ) is revealed, as they may depend on the values of these parameters. Hence, they are written as functions $d_2(\rho, \tau)$ and $N_2(\rho, \tau)$ of the uncertain parameters (ρ, τ) . In ARO such variables are also called *wait-and-see* variables or decisions. In this paper, we will adhere to the terms *stage 1* and *stage 2*, however.

Before we solve (8), we need some definitions and assumptions. The remaining BED tolerance level of the OAR, if N' fractions with dose d' have been administered, is given by

$$B(d', N'; \rho) = \text{BED}_{\text{tol}}(\rho) - \sigma d' N' - \rho \sigma^2 (d')^2 N'. \quad (9)$$

Subsequently, define the function

$$g(d', N', N''; \rho) = \frac{-1 + \sqrt{1 + \frac{4\rho}{N''} B(d', N'; \rho)}}{2\sigma\rho}. \quad (10)$$

The value of g can be interpreted as the maximum dose that can be delivered in N'' fractions if already N' fractions of dose d' are (scheduled to be) delivered. It is obtained by solving the equality version of (8b) for d_1 or d_2 . Functions of this form will be used frequently throughout the remainder of this paper.

The following assumption on the relation between d^{\min} , d_1^{\max} and the bounds on N_2 makes sure that for a given optimal number of fractions, it is feasible to deliver that number of fractions with minimum dose.

Assumption 1 It holds that

$$d^{\min} \leq d_1^{\max} \leq \min \left\{ g(d^{\min}, N_2^{\min}, N_1; \rho_L), g(d^{\min}, N_2^{\max}, N_1; \frac{\tau_L}{\sigma}), g(d^{\min}, N_2^{\max}, N_1; \rho_U) \right\}. \quad (11)$$

The particular form of the upper bound on d_1^{\max} will become clear later. Numerical experiments indicate that results are not very sensitive to the choices of d^{\min} and d_1^{\max} .

We continue by formally defining several properties of solutions. Let $X(\rho, \tau)$ denote the feasible region defined by constraints (8b)–(8e) for fixed (ρ, τ) .

Definition 1 (*Adjustable robustly feasible*) A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is *adjustable robustly feasible* (ARF) to (8) if $(d_1, d_2(\rho, \tau), N_2(\rho, \tau)) \in X(\rho, \tau)$ for all $(\rho, \tau) \in Z$.

Definition 2 (*Adjustable robustly optimal*) A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is *adjustable robustly optimal* ($\mathcal{AR}\mathcal{O}$) to (8) if it is $\mathcal{AR}\mathcal{F}$ and there does not exist an $\mathcal{AR}\mathcal{F}$ tuple $(\bar{d}_1, \bar{d}_2(\cdot), \bar{N}_2)$ such that

$$\begin{aligned} & \min_{(\rho, \tau) \in Z} N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau(N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \\ & < \min_{(\rho, \tau) \in Z} N_1 \bar{d}_1 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau) + \tau(N_1 \bar{d}_1^2 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau)^2). \end{aligned} \tag{12}$$

We also define the $\mathcal{AR}\mathcal{O}$ property for the stage-1 decisions d_1 individually.

Definition 3 (*Adjustable robustly optimal d_1*) A stage-1 decision d_1 is $\mathcal{AR}\mathcal{O}$ to (8) if there exist decision rules $d_2(\cdot)$ and $N_2(\cdot)$ such that $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{AR}\mathcal{O}$ to (8).

Lastly, we define optimality of a decision rule.

Definition 4 (*Optimal decision rule*) For a given d_1 , a decision rule pair $(d_2(\cdot), N_2(\cdot))$ is *optimal* to (8) if $(d_1, d_2(\cdot), N_2(\cdot))$ is $\mathcal{AR}\mathcal{F}$ and for any $(\rho, \tau) \in Z$ it holds that

$$\begin{aligned} & N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau(N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \\ & \geq N_1 d_1 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau) + \tau(N_1 d_1^2 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau)^2), \end{aligned} \tag{13}$$

for every $(\bar{d}_2(\cdot), \bar{N}_2(\cdot))$ such that $(d_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ is $\mathcal{AR}\mathcal{F}$.

The first observation we make in (8) is that if $\epsilon > 0$, any fixed solution (d_1, d_2, N_2) feasible for scenario $(\rho, \tau) \in Z$ is also feasible for $(\rho, \tau + \epsilon)$ with a higher objective value. Therefore, in any worst-case realization it will hold that $\tau = \tau_L$ (see (5)). This observation has consequences for what uncertainty sets Z need to be considered. Due to the result (4), one can in general distinguish three cases for uncertainty set Z and parameter σ :

Case (1) $\sigma\rho_U \leq \tau_L$: According to (4), for any realization (with $\tau = \tau_L$) it is optimal to deliver the minimum number of fractions in stage 2.

Case (2) $\sigma\rho_L \geq \tau_L$: According to (4), for any realization (with $\tau = \tau_L$) it is optimal to deliver the maximum number of fractions in stage 2.

Case (3) $\sigma\rho_L < \tau_L < \sigma\rho_U$: In the scenario (ρ_L, τ_L) , it is optimal to deliver the maximum number of fractions in stage 2 according to (4). In the scenario (ρ_U, τ_L) , it is optimal to deliver the minimum number of fractions in stage 2 according to (4).

In Cases 1 and 2, (8) is easily solved by plugging in the (worst-case) optimal value for N_2 , and solving the resulting 2-variable optimization problem. Therefore, only Case 3 is of interest and in the remainder of this paper we make the following assumption.

Assumption 2 It holds that $\sigma\rho_L < \tau_L < \sigma\rho_U$.

In our numerical experiments, we use a lung cancer data set. Recent evidence suggests that for lung cancer Assumption 2 can indeed hold, i.e., the optimal number of treatment fractions is not always known prior to treatment. Further details are provided in Sect. 5.2. For other tumor sites, such as liver cancer, the tumor α/β is generally assumed to be 10 or higher, whereas the α/β of normal liver tissue is typically assumed to be 3 or 4. Thus, for such tumor sites Assumption 2 generally does not hold, and hyperfractionation is optimal.

3.2 Optimal decision rules and worst-case solution

Problem (8) is a 2-stage non-convex mixed-integer ARO problem, which are generally hard to solve. Nevertheless, due to the small number of variables the problem can be solved to optimality. In order to solve (8), we take two steps:

- Step (1)* Determine optimal decision rules $d_2(\cdot)$ and $N_2(\cdot)$ for fixed d_1 .
Step (2) Plug in optimal decision rules and solve for d_1 .

In what follows, we give a detailed explanation of both steps. Let $(d_1^*, d_2^*(\rho, \tau), N_2^*(\rho, \tau))$ denote an \mathcal{ARO} solution to (8).

Step 1: Determine optimal decision rules $d_2(\cdot)$ and $N_2(\cdot)$ for fixed d_1

Fix stage-1 variable d_1 . Similar to the result (4), we will show that it is optimal to deliver either the minimum or the maximum number of fractions in stage-2. Moreover, (8b) is the only OAR dose-limiting constraint, so it will hold with equality if this does not violate variable bounds (8d) and (8e). We will show that the latter is not the case. The theorem below summarizes the result.

Theorem 1 *Let d_1 be the stage-1 decision of (8). The decision rules*

$$N_2^*(\rho, \tau) = \begin{cases} N_2^{\min} & \text{if } \tau \geq \sigma\rho \\ N_2^{\max} & \text{otherwise,} \end{cases} \quad (14)$$

and


$$d_2^*(d_1; \rho, \tau) = \begin{cases} g(d_1, N_1, N_2^{\min}; \rho) & \text{if } \tau \geq \sigma\rho \\ g(d_1, N_1, N_2^{\max}; \rho) & \text{otherwise} \end{cases} \quad (15)$$

are optimal to (8) for the given d_1 . These provide the unique optimal decisions unless $\tau = \sigma\rho$.

Proof See Supplementary Material B.1. □

Clearly, these decision rules are nonlinear, and in fact split the uncertainty region in two parts: one where it is optimal to deliver the minimum number of fractions N_2^{\min} in stage 2, and one where it is optimal to deliver the maximum number of fractions N_2^{\max} in stage 2. This suggests splitting the uncertainty set as follows:

Fig. 3 Split of uncertainty set Z according to (16). The circles indicate the locations of the candidate worst-case scenarios for (18)



$$Z^{\min} = Z \cap \{\tau \geq \sigma\rho\} \tag{16a}$$

$$Z^{\max} = Z \cap \{\tau < \sigma\rho\}. \tag{16b}$$

An illustration is provided in Fig. 3.

Step 2: Plug in optimal decision rules and solve for d_1

In order to find an \mathcal{ARCO} d_1 , we introduce the following objective function for given (ρ, τ) :

$$f(d_1, N_2; \rho, \tau) = \begin{cases} N_1 d_1 + N_2 g(d_1, N_1, N_2; \rho) & \text{if } d_1 \in [0, g(0, 0, N_1; \rho)] \\ +\tau(N_1 d_1^2 + N_2 g(d_1, N_1, N_2; \rho)^2) & \\ -\infty & \text{otherwise,} \end{cases} \tag{17}$$

where, for given ρ , the value $g(0, 0, N_1; \rho)$ is the maximum dose that can be delivered in stage 1 due to the nonnegativity restriction on the stage-2 dose. From Assumption 1 it follows that $g(0, 0, N_1; \rho) \geq d_1^{\max}$ for all $(\rho, \tau) \in Z$, so f is finite for all feasible d_1 . According to Lemma 1 in Supplementary Material C, function f is either convex, concave or constant in d_1 .

Plugging in (14) and (15) and using definition (17) allows us to rewrite (8) to a problem of only variable d_1 :

$$\max_{d_1} \min_{(\rho, \tau) \in Z} f(d_1, N_2^*(\rho, \tau); \rho, \tau) \tag{18a}$$

$$\text{s.t. } d_1^{\min} \leq d_1 \leq d_1^{\max}. \tag{18b}$$

As noted in Sect. 3.1, in any worst-case realization it will hold that $\tau = \tau_L$, so it is sufficient to consider only those observations with $\tau = \tau_L$.

In order to reformulate (18), we make use of the properties of g and f in Lemma 3 in Supplementary Material C. In particular, Lemma 3b states that

function f is either increasing or decreasing in ρ for fixed d_1 . Hence, if we move (18a) to a constraint and split according to (16), for both Z^{\min} and Z^{\max} it is sufficient to consider the constraint for the highest and lowest value of ρ in the uncertainty set. With $\tau = \tau_L$, this results in the scenarios (ρ_L, τ_L) and $(\frac{\tau_L}{\sigma}, \tau_L)$ for Z^{\min} and $(\frac{\tau_L}{\sigma}, \tau_L)$ and (ρ_U, τ_L) for Z^{\max} .

Therefore, the three candidate worst-case scenarios are (ρ_L, τ_L) , (ρ_U, τ_L) and $(\frac{\tau_L}{\sigma}, \tau_L)$; their locations are indicated in Fig. 3. By Lemma 1, the objective value in the third scenario is equal to

$$K = \frac{1}{\sigma} B(0, 0, \frac{\tau_L}{\sigma}). \quad (19)$$

This is the maximum target BED that can be attained if radiation sensitivity parameters are exactly such that fractionation has no influence on the optimal objective value. It is equal to the maximum tolerable OAR BED for these radiation sensitivity parameters, divided by the generalized OAR dose sparing factor σ .

Putting everything together, we conclude that if $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is \mathcal{ARCO} to the EDP (8) then there exists a $q^* \in \mathbb{R}_+$ such that (d_1^*, q^*) is an optimal solution to

$$\max_{d_1, q} q \quad (20a)$$

$$\text{s.t. } q \leq f(d_1, N_2^{\min}; \rho_L, \tau_L) \quad (20b)$$

$$q \leq f(d_1, N_2^{\max}; \rho_U, \tau_L) \quad (20c)$$

$$q \leq K \quad (20d)$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \quad (20e)$$

Conversely, if (d_1^*, q^*) is an optimal solution to (20) and $N_2^*(\cdot)$ and $d_2^*(\cdot)$ are given by (14) and (15), respectively, then $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is \mathcal{ARCO} to (8). Hence, (20) and EDP (8) are equivalent.

According to Lemma 1, the RHS of (20b) and (20c) is convex and concave in d_1 , respectively. Hence, (20) asks to find the value of d_1 that maximizes the minimum of a univariate convex (20b), concave (20c) and constant (20d) function on a closed interval (20e). Lemma 2 in Supplementary Material C provides information on the intersection points of the functions (20b)-(20d). Consequently, the optimal solution(s) to (20) is/are easily found.

Figure 4 illustrates a possible instance of (20), displaying constraints (20b)-(20d). In this case, the set of optimal solutions is the union of the two intervals for d_1 where constraint (20d) is active. This is indicated in red. Dose constraints (20e) may cut off part of these intervals. If due to constraint (20e) both these intervals are infeasible, the optimum is at one of the boundaries for d_1 .




Fig. 4 Schematic illustration of (20). The solid and dashed curves represent constraints (20b) and (20c), respectively, and the dotted line represent constraint (20d). The optimal intervals are indicated in red

3.3 Pareto adjustable robustly optimal solutions

Figure 4 illustrates that it is possible that there are multiple optimal solutions to (20). These solutions are \mathcal{ARO} stage-1 solutions to the EDP (8). In general, in case there are multiple \mathcal{ARO} solutions these may perform vastly different if a non-worst-case scenario realizes (De Ruiter et al. 2016). Iancu and Trichakis (2014) study static robust optimization problems with multiple robustly optimal solutions, and introduce the concept of Pareto robustly optimal (\mathcal{PRO}) solutions. A robustly optimal solution is called \mathcal{PRO} if there is no other robustly feasible solution that has equal or better objective value for all scenarios in the uncertainty set, while being strictly better for at least one scenario. Non- \mathcal{PRO} solutions are dominated by at least one \mathcal{PRO} solution and are therefore not desired⁴. The concept has previously been applied to RT planning for breast cancer by Mahmoudzadeh (2015).

Iancu and Trichakis (2014) study \mathcal{PRO} solutions solely for static RO problems; we generalize the concept to 2-stage adjustable robust optimization problems.

Definition 5 (*Pareto adjustable robustly optimal*) An \mathcal{ARO} tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is *Pareto adjustable robustly optimal* (\mathcal{PARO}) to (8) if there is no tuple $(\bar{d}_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ that is \mathcal{ARO} to (8) and satisfies

⁴ The concept of Pareto robust optimality closely resembles the concept of Pareto efficiency in multi-criteria optimization (MCO). In MCO, Pareto efficient solutions can only be improved in one criteria at the cost of a deterioration in another criteria. Only Pareto efficient solutions are of interest, and the overall goal in MCO is to compute this set of solutions (known as the Pareto surface).

$$\begin{aligned} N_1 d_1 + N_2(\rho, \tau) d_2(\rho, \tau) + \tau(N_1 d_1^2 + N_2(\rho, \tau) d_2(\rho, \tau)^2) \\ \leq N_1 \bar{d}_1 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau) + \tau(N_1 \bar{d}_1^2 + \bar{N}_2(\rho, \tau) \bar{d}_2(\rho, \tau)^2) \quad \forall (\rho, \tau) \in Z \end{aligned} \quad (21a)$$

$$\begin{aligned} N_1 d_1 + N_2(\bar{\rho}, \bar{\tau}) d_2(\bar{\rho}, \bar{\tau}) + \tau(N_1 d_1^2 + N_2(\bar{\rho}, \bar{\tau}) d_2(\bar{\rho}, \bar{\tau})^2) \\ < N_1 \bar{d}_1 + \bar{N}_2(\bar{\rho}, \bar{\tau}) \bar{d}_2(\bar{\rho}, \bar{\tau}) + \bar{\tau}(N_1 \bar{d}_1^2 + \bar{N}_2(\bar{\rho}, \bar{\tau}) \bar{d}_2(\bar{\rho}, \bar{\tau})^2) \quad \text{for some } (\bar{\rho}, \bar{\tau}) \in Z. \end{aligned} \quad (21b)$$

We also define the concept \mathcal{PARO} for the stage-1 decision d_1 individually.

Definition 6 (*Pareto adjustable robustly optimal d_1*) A stage-1 decision d_1 is \mathcal{PARO} to (8) if there exist decision rules $N_2(\cdot)$ and $d_2(\cdot)$ such that $(d_1, d_2(\cdot), N_2(\cdot))$ is \mathcal{PARO} to (8).

If there are multiple \mathcal{ARO} solutions, we wish to pick one that is \mathcal{PARO} . In general, finding \mathcal{PARO} solutions is hard, because it requires comparing the performance of both stage-1 decisions and stage-2 decision rules on multiple scenarios simultaneously. However, for the current problem (14) and (15) are optimal decision rules. Plugging these in conditions (21) reduces the problem of finding a \mathcal{PARO} solution to solely comparing the performance of \mathcal{ARO} stage-1 decisions d_1 in non-worst-case scenarios.

In Iancu and Trichakis (2014) it is shown for linear optimization that, if we optimize over the robustly optimal solutions for a second criterion (scenario) that is in the relative interior of the uncertainty set, the resulting solution(s) are \mathcal{PRO} . In a similar fashion \mathcal{PARO} solutions to the current problem can be found. Let $X^{\mathcal{ARO}}$ denote the set of \mathcal{ARO} stage-1 solutions to (8). It turns out that consecutively optimizing over an auxiliary scenario where hyperfractionation is optimal and an auxiliary scenario where hypofractionation is optimal yields a set of \mathcal{PARO} solutions. Let $(\rho^{\text{aux-min}}, \tau^{\text{aux-min}}) \in \text{int}(Z^{\text{min}})$, where $\text{int}(\cdot)$ is the interior operator. Define the auxiliary optimization problem for the hypofractionation scenario:

$$\max_{d_1 \in X^{\mathcal{ARO}}} f(d_1, N_2^{\text{min}}; \rho^{\text{aux-min}}, \tau^{\text{aux-min}}). \quad (22)$$

Denote the set of optimal solutions to (22) by $X^{\text{aux-min}}$. Similarly, let $(\rho^{\text{aux-max}}, \tau^{\text{aux-max}}) \in \text{int}(Z^{\text{max}})$. Define the auxiliary optimization problem for the hyperfractionation scenario:

$$\max_{d_1 \in X^{\text{aux-min}}} f(d_1, N_2^{\text{max}}; \rho^{\text{aux-max}}, \tau^{\text{aux-max}}). \quad (23)$$

Note that it uses $X^{\text{aux-min}}$ as input, i.e., we solve the auxiliary problems consecutively. Denote the set of optimal solutions to (23) by $X^{\mathcal{PARO}}$.

Theorem 2 All solutions in $X^{\mathcal{PARO}}$ are \mathcal{PARO} to (8).

Proof See Supplementary Material B.2. □

Solving (22) or (23) entails maximizing a strictly convex or strictly concave function over a feasible set consisting of a small number of intervals or points. Hence, these auxiliary problems are easily solved. Note that the second auxiliary problem is only relevant if the first auxiliary problem has multiple optimal solutions. The two-step approach is necessary; numerical results show that optimizing over only one auxiliary scenario may indeed result in non-*PARO* solutions.

Switching their order, and optimizing (22) over the set $X^{\text{aux-max}}$ may lead to different solutions, and these are also *PARO*. Thus, in general X^{PARO} does not contain all *PARO* solutions. Thus, the used auxiliary scenarios, and the order in which they are optimized for, determine which *PARO* solution is obtained. By optimizing for likely or relevant scenarios auxiliary scenarios, a *PARO* solution is obtained with the best possible performance in those scenarios.

4 ARO: biomarkers provide inexact information

In this section we present an adjustable robust optimization approach to solve a more realistic version of the adaptive treatment-length problem. Because in practice it is impossible to exactly determine the α/β parameters from biomarker data, any values for the α/β parameters obtained during treatment are inexact. This section presents a model that accounts for uncertainty in biomarker information.

4.1 Problem formulation

The setup for the ARO problem with inexact data is based on De Ruiter et al. (2017). After N_1 fractions we obtain an estimate $(\hat{\rho}, \hat{\tau})$ for (ρ, τ) , the inverse α/β parameters for the OAR and the tumor. It is still assumed that Assumption 2 holds for uncertainty set Z . Furthermore, we assume that $(\rho, \tau), (\hat{\rho}, \hat{\tau}) \in Z$ (as defined in (5)), i.e., both the observation and the true realization are in Z . For the accuracy of the observation, we assume that $(\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}$, with

$$\hat{Z} = \{(\varepsilon^\rho, \varepsilon^\tau) \in \mathbb{R}^2 : |\varepsilon^\rho| \leq r^\rho, |\varepsilon^\tau| \leq r^\tau\}. \tag{24}$$

In practice, parameters ρ and τ would be estimated using different biomarkers, with potentially different accuracies. Parameters r^ρ and r^τ define the accuracy of the observations $\hat{\rho}$ and $\hat{\tau}$. Set \hat{Z} is the uncertainty set around the inexact observation. This can also be written as $(\rho, \tau) \in \{(\hat{\rho}, \hat{\tau})\} + \hat{Z}$, which is the Minkowski sum of a singleton and a set. Of this new set, we will only use the part that is contained in the original uncertainty set Z , i.e., the observation cannot cause the true realization to be outside of Z . Define

$$U = \{(\rho, \tau, \hat{\rho}, \hat{\tau}) : (\rho, \tau), (\hat{\rho}, \hat{\tau}) \in Z, (\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}\}, \tag{25}$$

and

$$Z_{(\hat{\rho}, \hat{\tau})} = (\{(\hat{\rho}, \hat{\tau})\} + \hat{Z}) \cap Z. \tag{26}$$

The set U contains all possible observation-realization pairs and $Z_{(\hat{\rho}, \hat{\tau})}$ contains all possible realizations after observation of $(\hat{\rho}, \hat{\tau})$. For given observation $(\hat{\rho}, \hat{\tau})$, the new upper and lower bounds for (ρ, τ) are given by

$$\hat{\tau}_L = \max\{\tau_L, \hat{\tau} - r^r\}, \quad \hat{\tau}_U = \min\{\tau_U, \hat{\tau} + r^r\} \quad (27a)$$

$$\hat{\rho}_L = \max\{\rho_L, \hat{\rho} - r^p\}, \quad \hat{\rho}_U = \min\{\rho_U, \hat{\rho} + r^p\}. \quad (27b)$$

Compared to Sect. 3, we remove Assumption 1 and impose a different (slightly stricter) assumption on the relation between d^{\min} , d_1^{\max} and the bounds on N_2 .

Assumption 3 It holds that $d^{\min} \leq d_1^{\max}$ and

$$d_1^{\max} \leq \min\left\{g(d_1^{\min}, N_2^{\min}, N_1; \rho_L), g(d_1^{\min}, N_2^{\max}, N_1; \max\{\rho_L, \frac{\tau_L}{\sigma} - 2r^p\}), g(d_1^{\min}, N_2^{\max}, N_1; \rho_U)\right\}. \quad (28)$$

The inexact data problem (IDP) analogous to (8) is given by

$$\max_{d_1, d_2(\hat{\rho}, \hat{\tau}), N_2(\hat{\rho}, \hat{\tau})} \min_{(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U} N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2), \quad (29a)$$

$$\text{s.t.} \quad \sigma(N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})) + \rho \sigma^2(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2) \leq \text{BED}_{\text{tol}}(\rho), \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U \quad (29b)$$

$$N_2(\hat{\rho}, \hat{\tau}) \in \{N_2^{\min}, \dots, N_2^{\max}\}, \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U \quad (29c)$$

$$d_2(\hat{\rho}, \hat{\tau}) \geq d^{\min}, \quad \forall(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U \quad (29d)$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \quad (29e)$$

For stage-2 variables $d_2(\hat{\rho}, \hat{\tau})$ and $N_2(\hat{\rho}, \hat{\tau})$ it is indicated that they are a function of the observations $(\hat{\rho}, \hat{\tau})$ instead of the uncertain parameters (ρ, τ) . Similar to Sect. 3, we formally define several properties of solutions. Let $X(\rho, \tau, \hat{\rho}, \hat{\tau})$ denote the feasible region defined by constraints (29b)–(29e) for fixed $(\rho, \tau, \hat{\rho}, \hat{\tau})$.

Definition 7 (*Adjustable robust feasibility*) A tuple $(d_1, d_2(\cdot), N_2(\cdot))$ is *adjustable robustly feasible* (\mathcal{ARF}) to (29) if $(d_1, d_2(\hat{\rho}, \hat{\tau}), N_2(\hat{\rho}, \hat{\tau})) \in X(\rho, \tau, \hat{\rho}, \hat{\tau})$ for all $(\rho, \tau, \hat{\rho}, \hat{\tau}) \in U$.

Optimality of a decision rule is defined as follows.

Definition 8 (*Optimal decision rule*)

For a given d_1 , a decision rule pair $(d_2(\cdot), N_2(\cdot))$ is *optimal* to (29) if $(d_1, d_2(\cdot), N_2(\cdot))$ is \mathcal{ARF} and for any $(\hat{\rho}, \hat{\tau}) \in Z$ it holds that

$$\begin{aligned} & \min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} N_1 d_1 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + N_2(\hat{\rho}, \hat{\tau}) d_2(\hat{\rho}, \hat{\tau})^2) \\ & \geq \min_{(\rho, \tau) \in Z_{(\hat{\rho}, \hat{\tau})}} N_1 d_1 + \bar{N}_2(\hat{\rho}, \hat{\tau}) \bar{d}_2(\hat{\rho}, \hat{\tau}) + \tau(N_1 d_1^2 + \bar{N}_2(\hat{\rho}, \hat{\tau}) \bar{d}_2(\hat{\rho}, \hat{\tau})^2), \end{aligned} \tag{30}$$

for every $(\bar{d}_2(\cdot), \bar{N}_2(\cdot))$ such that $(d_1, \bar{d}_2(\cdot), \bar{N}_2(\cdot))$ is \mathcal{ARF} .

Note that for exact data, an optimal decision rule gives the optimal decision for any *realization* in the uncertainty set Z (given d_1). For inexact data, we call a decision rule optimal if it yields the maximum worst-case (guaranteed) objective value for any *observation* in the uncertainty set Z .

4.2 Optimal decision rules and conservative approximation

Depending on both the observed $(\hat{\rho}, \hat{\tau})$ and the quality of the biomarker information (i.e., r^p and r^τ), we may be able to immediately determine the optimal value for N_2 . Therefore, we split the uncertainty set for the observations $(\hat{\rho}, \hat{\tau})$. Define

$$Z_{ID}^{\min} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \geq \sigma \hat{\rho}_U\} \tag{31a}$$

$$Z_{ID}^{\text{int}} = \{(\hat{\rho}, \hat{\tau}) \in Z : \sigma \hat{\rho}_L < \hat{\tau}_L < \sigma \hat{\rho}_U\} \tag{31b}$$

$$Z_{ID}^{\max} = \{(\hat{\rho}, \hat{\tau}) \in Z : \hat{\tau}_L \leq \sigma \hat{\rho}_L\}, \tag{31c}$$

so that $Z = Z_{ID}^{\min} \cup Z_{ID}^{\text{int}} \cup Z_{ID}^{\max}$. Figure 5 provides an illustration. For an observation $(\hat{\rho}, \hat{\tau})$ there are two candidate worst-case scenarios: $(\hat{\rho}_L, \hat{\tau}_L)$ and $(\hat{\rho}_U, \hat{\tau}_L)$. If both are on one side of the line $\tau = \sigma\rho$, the optimal fractionation decision is known. For an observation $(\hat{\rho}, \hat{\tau})$ near the lower boundary of Z it holds that $\hat{\tau}_L = \tau_L$, so a further decrease in $\hat{\tau}$ has no influence on $(\hat{\rho}_L, \hat{\tau}_L)$ and $(\hat{\rho}_U, \hat{\tau}_L)$ (see (27)), similar for the other boundaries. This leads to the nonlinearity in the subset boundaries of Fig. 5.

The split is such that if $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\min}$ or $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\max}$ only N_2^{\min} resp. N_2^{\max} fractions can be optimal in stage 2. Subset Z_{ID}^{int} is the area between the dash-dotted lines. If $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\text{int}}$ both N_2^{\min} and N_2^{\max} fractions in stage 2 may be optimal for the true (ρ, τ) . The following theorem states the optimal stage-2 decision rules for a given value of d_1 .

Theorem 3 *Let d_1 be the stage-1 decision of (29). The decision rules*

$$N_2^*(d_1; \hat{\rho}, \hat{\tau}) = \begin{cases} N_2^{\min} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\min} \\ \arg \max_{N_2 \in \{N_2^{\min}, \dots, N_2^{\max}\}} \min\{f(d_1, N_2; \hat{\rho}_L, \hat{\tau}_L), f(d_1, N_2; \hat{\rho}_L, \hat{\tau}_L)\} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\text{int}} \\ N_2^{\max} & \text{if } (\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\max}, \end{cases} \tag{32}$$

and

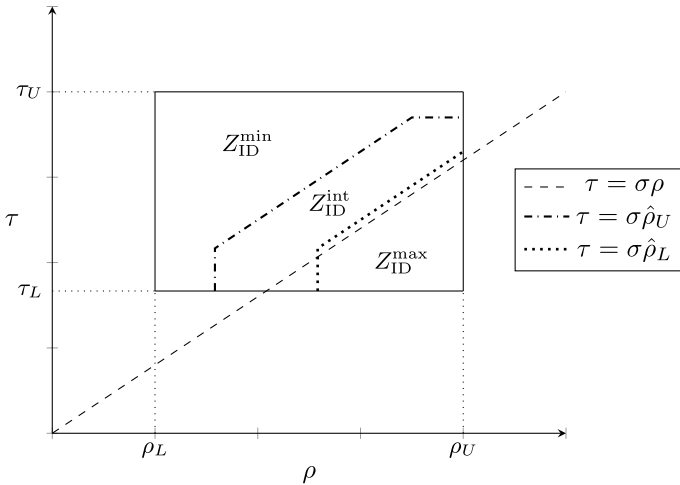


Fig. 5 The uncertainty set Z (solid lines) for $(\hat{\rho}, \hat{\tau})$ is split into Z_{ID}^{\min} , Z_{ID}^{int} , Z_{ID}^{\max} , according to (31). Subset Z^{int} is the area between the dotted and dash-dotted curves. If $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\text{int}}$ both N_2^{\min} and N_2^{\max} fractions in stage 2 may be optimal for the true (ρ, τ) . If $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\min}$ or $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\max}$ only N_2^{\min} resp. N_2^{\max} fractions can be optimal in stage 2

$$d_2^*(d_1; \hat{\rho}, \hat{\tau}) = \min\{g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_L), g(d_1, N_1, N_2^*(d_1; \hat{\rho}, \hat{\tau}); \hat{\rho}_U)\}, \quad (33)$$

are optimal to (29) for the given d_1 .

Proof See e Supplementary Material B.3. □

The worst-case optimal decision rule (32) may give a value unequal to N_2^{\min} and N_2^{\max} if $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\text{int}}$.

If r^ρ and r^τ are zero, i.e., we have exact data, then it holds that $\hat{\tau}_L = \tau$ and $\hat{\rho}_L = \hat{\rho}_U = \rho$. Hence, the two functions f in the RHS of (32) are equal, and the optimal N_2^* is the one that maximizes the resulting function. One can verify that this does not depend on d_1 . Hence, in case of exact data Theorem 3 reduces to Theorem 1.

It turns out that, after plugging in (32) and (33), and splitting the uncertainty set according to (31), it is not apparent how to determine the optimal stage-1 decision d_1^* for (29). In Supplementary Material B.4 the following lower bound problem to (29) is derived, named the Approximate Inexact Data Problem (AIDP):

$$\max_{d_1, q} q \quad (34a)$$

$$\text{s.t. } q \leq f(d_1, N_2^{\min}; \rho_L, \tau_L) \quad (34b)$$

$$q \leq f(d_1, N_2^{\max}; \rho_U, \tau_L) \quad (34c)$$




Fig. 6 Schematic illustration of (34). Compared to the case with exact data (Fig. 4), the thick black curve (constraint (34e)) is extra. The solid, dashed and dotted lines/curves represent constraints (34b), (34c) and (34d), respectively. Optimal solutions are indicated by red circles

$$q \leq K \tag{34d}$$

$$q \leq p(d_1) \tag{34e}$$

$$d^{\min} \leq d_1 \leq d_1^{\max}. \tag{34f}$$

The AIDP is best explained using an example. Figure 6 illustrates a possible instance of (34), displaying constraints (34b)–(34e). Compared to (20) for exact biomarker information (see Fig. 4), problem (34) has the added constraint (34e); a piecewise convex-concave function $p(d_1)$ defined by (C.32) in Lemma 5 (Supplementary Material C). It can be interpreted as follows. If $(\hat{\rho}, \hat{\tau}) \in Z_{ID}^{\text{int}}$, the optimal number of stage-2 fractions can be in between N_2^{\min} and N_2^{\max} , as shown in Theorem 3. In those cases, the optimal number of fractions also depends on the already delivered stage-1 dose d_1 . The upper kinks (red circles) in the piecewise convex-concave function $p(d_1)$ in Fig. 6 indicate values of d_1 where the worst-case scenario changes. The lower kinks indicate values where the optimal number of stage-2 fractions changes. The exact expression for $p(d_1)$ does not provide additional insight and is therefore omitted here.

In Figure 6, optimal solutions are locations where (34d) and (34e) are both binding, indicated by red circles. Dose constraints (34f) may cut off (some of) these points. If due to constraint (34f) none of the circles are feasible, the optimum is at one of the boundaries for d_1 . Constraint (34e) is the only conservative constraint in (34). Hence, only if the feasible values for d_1 are such that none of

the circles in Fig. 6 are feasible and constraint (34e) is binding, it is possible that the optimal objective value of (34) is strictly worse than that of (29).

Lemmas 1, 2 and 5 in Supplementary Material C provide information on the shape and intersection points of constraint functions (34b)–(34e). Consequently, the optimal solution(s) of (29) is/are easily obtained. If (d_1^*, q^*) is optimal to AIDP (34), and $N_2^*(\cdot)$ and $d_2^*(\cdot)$ are given by (32) and (33) then $(d_1^*, d_2^*(\cdot), N_2^*(\cdot))$ is \mathcal{ARF} to the original IDP (29). It is \mathcal{ARF} because AIDP is a conservative approximation of IDP.

4.3 Pareto robustly optimal solutions to conservative approximation

Figure 6 also illustrates that it is possible that there are multiple optimal solutions to the AIDP (34). Because the AIDP provides a conservative approximation to IDP (29), optimizing over auxiliary scenario(s) as in Sect. 3 does not necessarily produce a stage-1 decision d_1 that is \mathcal{PARO} to the original IDP. It turns out that a \mathcal{PRO} solution to AIDP is obtained from the set of robustly optimal solutions to AIDP if we consecutively optimize for two auxiliary observations such that any worst-case realization is in the interior of set Z^{\min} resp. Z^{\max} . Two important remarks are in place here. First, a \mathcal{PRO} solution to AIDP need not be a \mathcal{PARO} solution to IDP, even if it is \mathcal{ARO} to IDP. Second, the required auxiliary scenarios need not exist; their existence depends on the values of r^ρ and r^τ . Hence, further details are omitted.

5 Numerical results

This section presents numerical results of the methods presented in Sects. 3 and 4. First, Sect. 5.1 describes the benchmark methods against which we compare the ARO method for EDP and IDP, and Sect. 5.2 describes the setup of the numerical experiments.

5.1 Benchmark static and folding horizon methods

We analyze the performance of the static and folding horizon nominal method (NOM and NOM-FH), the static and folding horizon robust optimization method (RO and RO-FH) and the adjustable robust optimization method (ARO). In the folding horizon approaches only the stage-1 decisions are implemented, and the model is re-optimized for the second stage once the biomarker information is revealed.

The static method NOM optimizes for the nominal parameter values $(\bar{\rho}, \bar{\tau})$ and disregards any uncertainty and adaptability. This method is the same for both EDP and IDP. In stage 2, NOM-FH solves the nominal problem under the assumption that the obtained biomarker estimate is exact (which is an invalid assumption for IDP). This method does not guarantee robustly feasible solution (feasible for all $(\rho, \tau) \in Z$) nor a robustly optimal solution (\mathcal{RO} ; (static) optimal for the worst-case realization $(\rho, \tau) \in Z$). The static method RO optimizes for the worst-case realization of (ρ, τ) in the uncertainty set Z , and disregards adaptability. For EDP the method RO-FH solves the same nominal problem as NOM-FH in stage 2; for IDP it solves a static

Table 1 Guaranteed solution properties of the five methods

Problem	Method				
	NOM	NOM-FH	RO	RO-FH	ARO
EDP	–	–	\mathcal{RO}	\mathcal{RO}	\mathcal{PARO}
IDP	–	–	\mathcal{RO}	\mathcal{RO}	\mathcal{ARF}

robust optimization problem in stage 2, for which the uncertainty set is determined by the accuracy of the biomarker information. RO and RO-FH both guarantee an \mathcal{RO} solution.

One may add a folding horizon component to ARO (for either EDP or AIDP). This may improve the results in case a suboptimal stage-2 decision rule is used. However, as shown in Sects. 3.2 and 4.2, the used stage-2 decision rules are optimal for any realized scenario (and for given stage-1 decision d_1 in case of inexact information). Hence, adding a folding horizon component will not change results.

Table 1 provides an overview of the guaranteed solution properties of the methods. It is important to note that in case of inexact biomarker information (IDP) the methods RO and RO-FH guarantee an \mathcal{RO} solution, whereas ARO guarantees only an \mathcal{ARF} solution via solving the approximate problem AIDP. Depending on the approximation quality, the \mathcal{ARF} solution may be close or equal to an \mathcal{ARO} solution.

Next to these five methods, we also report the results for the perfect information optimum (PI). This is the attainable optimum if from the start of the first fraction the true (ρ, τ) is exactly known. It can be formulated by taking the nominal problem and replacing the nominal parameter values by their true values. While in practice not a viable strategy, PI provides information on the value of perfect information, and allows us to put the performance of and differences between the other methods in perspective.

We note that all instances (both EDP and AIDP) are low dimensional, so the computation time is negligible for any of the six solution methods; in the numerical experiments all computation times are within 10 ms.

5.2 Study setup

We use a data set of 30 non-small cell lung cancer (NSCLC) patients, treated with either photon or proton therapy. The mathematical models in Sects. 3 and 4 are based on the assumption that there is a single dose restricting OAR. We assume that the single dose restricting OAR is the normal lung itself⁵. For the models in Sects. 3 and 4, an instance is defined by a tuple $(\sigma, \varphi, D, T, N_1, N^{\min}, N^{\max}, d^{\min}, d_1^{\max})$ and the relevant uncertainty sets.

⁵ This is in line with clinical practice wherein normal lung is treated as the most important normal tissue and the treatment is designed as to minimize the radiation exposure to normal lung.

Clinically, the number of treatment fractions varied between 33 and 37 fractions, with the majority of patients receiving 37 fractions. We set $N^{\min} = 30$ and $N^{\max} = 40$, to allow for slight deviations from the clinical standard. We assume the biomarker acquisition is made once $N_1 = 10$ fractions have been administered. This implies $N_2^{\min} = 20$ and $N_2^{\max} = 30$. Mean lung dose tolerance is $D = 20$ Gy, and we set $T = 37$ as that is the clinically standard regimen. The patients differ in (σ, φ) , which characterize the spatial dose distribution. Using the clinically delivered dose distribution, we derive for each normal lung voxel its dose sparing factor s_i (see Sect. 2.1). The dose shape factor φ and the generalized dose sparing factor σ for mean OAR BED are given by

$$\varphi = \frac{n \sum_{i=1}^n s_i^2}{\left(\sum_{i=1}^n s_i\right)^2}, \quad (35a)$$

$$\sigma = \frac{\sum_{i=1}^n s_i^2}{\sum_{i=1}^n s_i}, \quad (35b)$$

see Perkó et al. (2018) for details.

Cox (1986) estimate normal lung tissue α/β to be between 2.4 and 6.3. We set the nominal value at the midpoint 4.35. The α/β of NSCLC lung tumors has traditionally been assumed to be above 10 Gy. However, recent NSCLC hypofractionation trials show promising results, indicating that NSCLC cells are more sensitive to fraction size than previously assumed, i.e., have a lower α/β than 10. Santiago et al. (2016) find values between 2.2 and 9.0. We set the nominal value at the midpoint 5.6. Put together, we get the following uncertainty set for the inverse α/β ratios:

$$Z = \{(\rho, \tau) : 1/6.3 \leq \rho \leq 1/2.4, 1/9.0 \leq \tau \leq 1/2.2\}, \quad (36)$$

and the nominal scenario is $(\bar{\rho}, \bar{\tau}) = (1/4.35, 1/5.6)$. With this uncertainty set, 20 out of 30 patient cases satisfy Assumption 2: these are used in the numerical experiments. For the remaining ten patients the optimal number of treatment fractions can be determined prior to treatment, so these are removed. Preliminary numerical experiments indicate that results are not sensitive to variations in the bounds of the uncertainty set.

To discriminate between multiple \mathcal{ARO} solutions, we follow the procedure detailed in Sect. 3.3 in the case of exact biomarker information. The auxiliary scenarios are sampled uniformly from $\text{int}(Z^{\min})$ and $\text{int}(Z^{\max})$. In the case of inexact biomarker information, the procedure discussed in Sect. 4.3 is followed if the required auxiliary observations exist. If such observations exist, we sample uniformly from Z until we have found two auxiliary observations for which any worst-case realization is in $\text{int}(Z^{\min})$ resp. $\text{int}(Z^{\max})$. If such observations do not exist, the robustly optimal solution to AIDP with lowest stage-1 dose is selected. The method RO (and therefore also RO-FH) may also find multiple robustly optimal solutions. For the obtained set of robustly optimal solutions we again follow the procedure detailed in Sect. 3.3. It turns out that for RO, the robustly

Table 2 Results for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z (200 scenarios)

	Method					
	NOM	NOM-FH	RO	RO-FH	ARO	PI
Tumor BED - sample mean (Gy)	162.75	161.44	156.57	160.14	161.40	161.49
Tumor BED - sample 5% quantile (Gy)	151.98	150.94	147.57	150.04	150.90	151.04
Tumor BED - sample wc (Gy)	145.98	146.33	142.53	145.58	146.32	146.39
Tumor BED - wc over Z (Gy)	114.72	116.19	116.19	116.19	116.19	116.19
OAR violation - mean (%)	1.25	0	0	0	0	0
OAR violation - max (%)	4.22	0	0	0	0	0
Stage-1 dose d_1 (Gy)	1.50	1.50	2.29	2.29	1.51	1.66
Stage-2 dose d_2 (Gy)	3.45	3.24	2.48	2.95	3.24	3.19
Stage-2 fractions N_2	20.0	22.2	27.2	22.2	22.2	22.2

For each scenario, results are averaged over 20 patients*. All methods optimize for worst-case tumor BED in Z , which is displayed in bold.

*: the maximum OAR violation is computed over all patients and scenarios

optimal solutions often perform identical in non-worst-case scenarios. We optimize over the auxiliary scenarios consecutively; the first auxiliary scenario is the scenario corresponding to $\text{int}(Z^{\min})$.

The minimum dose per fraction is $d^{\min} = 1.5$ Gy and the maximum stage-1 dose per fraction is $d_1^{\max} = 3$ Gy. This satisfies Assumption 1 (for EDP) and 3 (for IDP). Using these parameter values, it is feasible to deliver N_2^{\max} fractions with dose d^{\min} in *all* scenarios in Z . This means that stage-1 decisions cannot render stage 2 infeasible for RO, NOM (and their FH counterparts) or PI. Numerical results indicate that results are not sensitive to the choice of d^{\min} and d_1^{\max} .

We consider a sample of 200 scenarios for (ρ, τ) from Z . For each scenario, the average tumor BED over 20 patients is computed, thus creating a tumor BED distribution for the ‘average’ patient. For this tumor BED distribution we report the mean, 5% quantile and worst-case value. Next to this, we report the true worst-case tumor BED over Z (averaged over 20 patients). Note that the true worst-case scenario can differ per patient, so the true worst-case BED is typically not attained in the sample. For OAR violations, we report the percentage by which the OAR BED tolerance is exceeded (i.e., percentage overdose). The maximum violation is the maximum value found over all patients and scenarios. All reported decision variable statistics are averaged over all patients and scenarios.

First, Sect. 5.3 presents and discusses the results for the problem with exact biomarker information (EDP) of Sect. 3. After that, Sect. 5.4 presents and discusses the results for the problem with inexact biomarker information (IDP) of Sect. 4. Lastly, Sect. 5.5 again considers the inexact biomarker information case, and varies parameter N_1 in order to determine the optimal moment of biomarker acquisition.

5.3 Results exact biomarker information

Table 2 presents the results. Altogether, the results indicate that the value of exact information is high. NOM-FH performs very similar to ARO. This illustrates that ignoring uncertainty and adaptability in stage 1 neither compromises worst-case or mean performance, nor does it lead to OAR constraint violations if treatment can be adapted based on exact biomarker information. In fact, NOM-FH outperforms RO-FH, indicating that accounting for uncertainty in stage-1 is overly conservative.

NOM is the only method that is not worst-case optimal, but yields the highest mean tumor BED across the sample. However, it is the only method that results in OAR constraint violations. In the nominal scenario $(\bar{\rho}, \bar{\tau})$ it is optimal to hypofractionate for all patients, so the mean N_2 equals 20 for NOM. The other static method, RO, is worst-case optimal, but yields lower tumor BED across the entire sample. The other static method, RO, is worst-case optimal, but yields low tumor BED across the sample. This is not due to one poor (patient, scenario) pair, but consistent throughout the entire sample. It delivers significantly more fractions on average, i.e., it decides to hyperfractionate more often.

NOM-FH adds a folding horizon component to NOM, and this results in zero violations and worst-case optimality. It does have slightly lower sample mean tumor BED. RO-FH adds a folding horizon component to RO, and this results in improved performance across the entire sample. It chooses to hypofractionate more often than RO. ARO is worst-case optimal and performs very similar to NOM-FH. Excluding NOM (for OAR constraint violations) and PI (not implementable), NOM-FH, ARO, RO-FH and RO yield the (possibly joint) highest objective value in 83.4%, 76.4%, 22.1% and 0.6% of all (scenario, patient) instances, respectively.

The results of Table 2 show that the methods have different stage-1 decisions d_1 ; this indicates the existence of multiple worst-case optimal stage-1 solutions. As indicated in Sect. 5.2, RO, RO-FH and ARO optimize over auxiliary scenarios in this case. According to Theorem 2, ARO finds a \mathcal{PARO} solution this way. Overall, methods that deliver a relatively low dose in stage-1 perform better than the methods that deliver a higher dose. This may be data set-specific. From PI we see that for the majority of patients and scenarios hypofractionation is optimal (average $N_2 = 22.2$), whereas the RO results indicate that for the majority of patients it is worst-case optimal to hyperfractionate (average $N_2 = 27.2$). We emphasize that for different data sets, where for the majority of scenarios and patients hyperfractionation is optimal, a higher stage-1 dose (which allows for N^{\max} constant-dose fractions) may perform better, such as the result of RO and RO-FH.

Figure 7 shows the complete cumulative scenario-tumor BED graph. A point (x, y) in Fig. 7 can be interpreted as follows: for the average patient, in $y\%$ of scenarios the tumor BED is at least x Gy. The results clearly demonstrate that RO and RO-FH are outperformed by the other methods. Both NOM-FH and ARO are visually almost indistinguishable from PI. NOM performs even better across the entire sample (except the first percent of the sample), at the cost of OAR BED violations.

To see the difference in mean performance between the multiple worst-case optimal solutions, we compare \mathcal{PARO} solution found by the ARO method to the \mathcal{ARO} solution that performs worst in the two auxiliary scenarios. Table 3 shows the




Fig. 7 Cumulative scenario-tumor BED graph for experiments with exact biomarker information and uniform sampling of (ρ, τ) over Z (200 scenarios). A point (x, y) indicates that in $y\%$ of scenarios the tumor BED (averaged over 20 patients) is at least x Gy. ARO and NOM-FH are very close to PI

Table 3 Comparison between the best ($PARO$) and worst performing ARO solutions, for uniform sampling of (ρ, τ) over Z (200 scenarios)

	Method	
	ARO_{worst}	ARO_{best}
Tumor BED - sample mean (Gy)	159.87	161.40
Tumor BED - sample 5% quantile (Gy)	149.86	150.90
Tumor BED - sample worst-case (Gy)	145.34	146.32
Tumor BED - wc over Z (Gy)	116.19	116.19
Stage-1 dose d_1 (Gy)	2.67	1.51
Stage-2 dose d_2 (Gy)	2.79	3.24
Stage-2 fractions N_2	22.2	22.2

For each scenario, results are averaged over 20 patients. All methods optimize for worst-case tumor BED in Z , which is displayed in bold. OAR constraint violations are zero in all cases

results, OAR constraint violations are zero in all cases. The worst-performing ARO solution has a considerably higher stage-1 dose. This implies that (for the current parameter settings) delivering a high stage-1 dose does not allow as much adjustment possibilities in stage 2 as a low stage-1 dose, but it does allow for adjustments to reach the worst-case optimum. Relative to the results of Table 2, the difference between the best and worst ARO solution is considerable: the worst-performing ARO solution performs worse than the RO-FH solution.

Supplementary Material A reports the results of an auxiliary experiment where the (ρ, τ) samples is drawn from a superset of Z , to compare the out-of-sample

performance of the methods. NOM remains the only method with OAR constraint violations. Compared to Table 2, static methods NOM and RO have poor performance. The relative performance of the adaptive methods remains mostly unchanged.

Altogether, the results of Sect. 5.3 demonstrate that if exact biomarker information is available mid-treatment, most stage-1 decisions allow for sufficient adaptation space in stage 2, also with realizations outside of Z . Different stage-1 decisions yield the worst-case optimum, have good performance on the scenario sample and have no OAR BED violations. We note that all presented differences in tumor BED are of relatively small magnitude. One reason for this is that the number of stage-2 fractions is restricted to $[N_2^{\min}, N_2^{\max}] = [20, 30]$. If the minimum number of fractions represents a ‘true’ hypofractionation case, the dose per fraction can vary more, and the difference in performance between hypo- and hyperfractionation strategies is amplified.

5.4 Results inexact biomarker information

In case of inexact biomarker information (IDP), we do not obtain the true parameter values (ρ, τ) after $N_1 = 10$ fractions, but only an estimate $(\hat{\rho}, \hat{\tau})$. As discussed in Sect. 4, we specify a new uncertainty set \hat{Z} such that $(\hat{\rho}, \hat{\tau}) - (\rho, \tau) \in \hat{Z}$. Let $DQ \in [0, 1]$ indicate the data quality. Then we set \hat{Z} such that the width of the new uncertainty intervals for τ and ρ is $(1 - DQ)$ times the width of the original intervals $[\tau_L, \tau_U]$ and $[\rho_L, \rho_U]$. That is, $DQ \cdot 100\%$ can be interpreted as the percentage by which the uncertainty intervals can be reduced due to the observation. The relation with the accuracy parameter r^ρ (or similarly r^τ) is given by

$$r^\rho = \frac{1}{2}(\rho_U - \rho_L)(1 - DQ). \quad (37)$$

Note that even $DQ = 0$ has some value as the new interval is centered around the observation, which already cuts off part of the original uncertainty set Z . We pick $DQ = 2/3$, so the obtained information after fraction N_1 reduces the size of the interval by 66.7% around the new observation. Variations for DQ are considered in Sect. 5.5. For all 20 patients the required auxiliary scenarios for the method of Sect. 4.3 can be found.

Table 4 shows the results. The robust methods RO, RO-FH and ARO are all worst-case optimal. This indicates that, although not theoretically guaranteed, ARO finds an ARO solution in all considered scenarios. The mean performance of RO-FH and ARO is further away from PI than in the case with exact biomarker information (Table 2). This is as expected, as due to inexact observations the possibility for ARO and RO-FH to make adjustments is less valuable, whereas PI is not influenced by this. On the other hand, NOM and NOM-FH are not worst-case optimal, but have better performance on the sample of scenarios, at the cost of OAR violations.

ARO is the only method (together with PI) that has a different stage-1 decision than in the case with exact biomarker information. This is because it is the only method that takes inexactness of biomarker information into account at the start of

Table 4 Results for experiments with inexact biomarker information (data quality $DQ = 2/3$) and uniform sampling of (ρ, τ) over Z (200 scenarios)

	Method					
	NOM	NOM-FH	RO	RO-FH	ARO	PI
Tumor BED - sample mean (Gy)	162.52	161.03	156.38	158.76	159.46	161.16
Tumor BED - sample 5% quantile (Gy)	151.36	150.12	147.04	148.51	148.94	150.21
Tumor BED - sample wc (Gy)	147.79	146.04	144.01	145.02	145.33	146.18
Tumor BED - wc over Z (Gy)	114.72	115.96	116.19	116.19	116.19	116.19
OAR violation - mean (%)	1.25	0.16	0	0	0	0
OAR violation - max (%)	4.23	1.49	0	0	0	0
Stage-1 dose d_1 (Gy)	1.50	1.50	2.29	2.29	1.79	1.65
Stage-2 dose d_2 (Gy)	3.45	3.25	2.48	2.78	2.92	3.20
Stage-2 fractions N_2	20.0	22.0	27.2	24	24.3	22.1

All results are averages over a sample of 20 patients. For each scenario, results are averaged over 20 patients*. All methods optimize for worst-case tumor BED in Z , which is displayed in bold.

*: the maximum OAR violation is computed over all patients and scenarios

stage 1. The average stage-1 dose d_1 differs considerably between ARO and RO-FH, whereas their worst-case performance is equal on average (and equal to PI). This demonstrates the existence of multiple worst-case optimal solutions. Whereas optimizing worst-case optimal solutions for ARO over two auxiliary scenarios does not guarantee a PARO solution (Sect. 4.3), results in Table 4 indicate that it does produce solutions that perform slightly better on average than RO-FH.

For ARO, it is noteworthy that the average number of stage-2 fractions (24.3 fx) differs from that of PI (22.1 fx). Although ARO uses optimal decision rules for stage 2, these are optimal for the worst-case scenario in the new uncertainty set $Z_{(\hat{\rho}, \hat{\tau})}$, and need not be optimal for the ‘true’ realization in this set. In fact, NOM-FH treats the inexact biomarker information as the ‘true’ parameter values, and administers 22.0 fractions, on average, which is closer than that of PI. Although the fractionation decision of NOM-FH is not worst-case optimal, Table 4 shows that it performs better on the sample of scenarios.

Figure 8 shows the complete cumulative scenario-tumor BED graph for the ‘average patient’. Whereas in case of exact biomarker information (Fig. 7), the ARO line was very close to PI, here a clear difference can be observed. NOM and NOM-FH outperform ARO (and RO and RO-FH) over the entire distribution.

The good performance of NOM and NOM-FH in terms of sample mean tumor BED does come at the cost of OAR violations. However, these are relatively minor. The reason for this is that the number of stage-2 fractions is relatively high (between 20 and 30fx), so any method delivers reasonably low dose per fraction in stage 2. Consequently, the quadratic term in the BED model is smaller, and so is the influence of the α/β parameters. With higher dose per fraction, the use of incorrect (e.g., nominal) α/β parameter values may result in higher OAR constraint violations. Preliminary experiments for stereotactic body radiation therapy (SBRT, an RT modality that uses around five high dose fractions) indeed result in slightly higher OAR constraint violations for




Fig. 8 Cumulative scenario-tumor BED graph for experiments with inexact biomarker information (data quality $DQ = 2/3$) and uniform sampling of (ρ, τ) over Z (200 scenarios). A point (x, y) indicates that in $y\%$ of scenarios the tumor BED (averaged over 20 patients) is at least x Gy. NOM-FH is very close to PI

NOM and NOM-FH. In any case, a trade-off can be observed between higher tumor BED attained by NOM and NOM-FH and associated OAR constraint violations.

5.5 Optimal moment of biomarker acquisition

The moment of biomarker observation need not be fixed. Part of the decision-making process then involves choosing this observation moment such that it maximally improves treatment quality. Late observation may result in limited possibilities for treatment adaptation, whereas with too early observation one cannot yet reliably observe the true individual patient response. Although one can incorporate N_1 as a decision variable in the mathematical model, the small decision space allows to simply vary its value in numerical experiments. We assume a (hypothetical) mathematical relationship between information point N_1 and the data quality parameter DQ . With N^{\max} the maximum number of fractions, we consider the following three data quality functions:

$$DQ_1(N_1) = \left(\frac{N_1}{N^{\max}} \right)^4 \tag{38a}$$

$$DQ_2(N_1) = \frac{N_1}{N^{\max}} \tag{38b}$$

Fig. 9 The biomarker data quality is a function of the number of treatment fractions N_1 after which it is acquired. We consider three functions $DQ_i(N_1)$, $i = 1, 2, 3$.





Fig. 10 Change in stage-1 dose d_1 (averaged over all patients and scenarios) when varying the information point N_1 from 1 to $N^{\max} - 1$, for the linear data quality function $DQ_2(N_1)$.

$$DQ_3(N_1) = \left(\frac{N_1}{N^{\max}} \right)^{1/4} . \tag{38c}$$

Hence, DQ_1 , DQ_2 and DQ_3 decalibe a convex, linear and concave relationship between observation moment and data quality, respectively. Figure 9 shows the graphs of the three functions. Whether DQ_1 , DQ_2 or DQ_3 is most realistic depends on the specific biomarker(s) that is/are used, see Sect. 2.2 for details.

We vary the information point N_1 from 0 to $N^{\max} - 1$. Figure 10 shows the change in stage-1 dose d_1 (averaged over all patients and scenarios) for linear data quality function DQ_2 , for methods PI, ARO, RO/RO-FH and NOM/NOM-FH. Results are

very similar for DQ_1 and DQ_3 . Recall from Sect. 5.3 that for the majority of (patient, scenario) pairs hypofractionation is optimal. For those cases, if N_1 is very low PI sometimes delivers a high dose boost in stage 1, and a low dose/tx in stage 2. For higher values of N_1 this is not possible anymore, leading to a lower average dose in stage 1 (and a dose boost in stage 2). For NOM/NOM-FH the same holds, because the nominal scenario is a hypofractionation scenario.

Also for ARO it is optimal to start with a higher dose per fraction in stage-1 if N_1 is very low. For slightly higher N_1 , starting with a dose boost is no longer possible. For most patients it is worst-case optimal to hyperfractionate, i.e., deliver an equal dose per fraction for N^{\max} fractions. For N_1 around 10 – 20, worst-case optimality can be attained also with lower dose per fraction in stage 1. This tends to be the PARO solution, because it enables a higher dose boost in stage 2 for hypofractionation scenarios. For higher N_1 , it is often not possible to achieve worst-case optimality (in hyperfractionation scenarios) if we deliver a low dose per fraction in stage 1. This leads to a gradual shift from low dose to medium dose per fraction in stage 1 as N_1 increases.

Figure 11 shows the mean tumor BED values and OAR constraint violations for varying N_1 , for data quality functions $DQ_i(N_1)$, $i = 1, 2, 3$. The left vertical axis indicates the mean tumor BED (averaged over all patients and scenarios), the right vertical axis indicates the maximum OAR tolerance violation for NOM and NOM-FH. It is important to note that as N_1 increases past $N^{\min} = 30$, this also increases the minimum number of fractions correspondingly. Moreover, the dose per fraction is constant per treatment stage, so the choice of N_1 also influences the types of treatments that can be delivered.

For these reasons the curve for PI is not constant, even though it does not actually use biomarker information. The optimal moment of biomarker acquisition for PI is $N_1 = 29$. This is because the minimum number of fractions is $N^{\min} = 30$. Hence, if hypofractionation is optimal we can deliver one more fraction with high dose, and deliver a low dose in stage 1. If hyperfractionation is optimal we can deliver 11 more fractions (and get the total maximum of 40) with low dose. Having $N_1 > 29$ forces the use of more than $N_2^{\min} = 30$ fractions, which is disadvantageous for those (patient, scenario) cases where hypofractionation is optimal.

For the latter reason the NOM curve is also not constant. It results in a higher tumor BED than PI for any value of N_1 , at the cost of OAR violations of up to 20%. NOM-FH yields a sample mean tumor BED close to PI for all three DQ functions, and the OAR violations depend on the DQ function. With poor data quality (Fig. 11a) and an observation moment close to $N_1 = 29$, OAR violations over 10% are possible, despite the fact that NOM-FH is an adaptive method. On the other hand, with good data quality (Fig. 11c), the violations remain below 2%. The OAR tolerance violations are highest near $N_1 = 29$. This is because in case of hypofractionation in stage 2 the influence of the uncertain α/β parameters is highest, as was noted in Sect. 5.4.

The robust methods RO, RO-FH and ARO do not result in any OAR violations, by construction. The better the data quality, the larger the differences between RO, RO-FH and ARO. This implies that, if robustness is required, there is value in (i) adapting based on inexact information, (ii) taking adaptability into account when

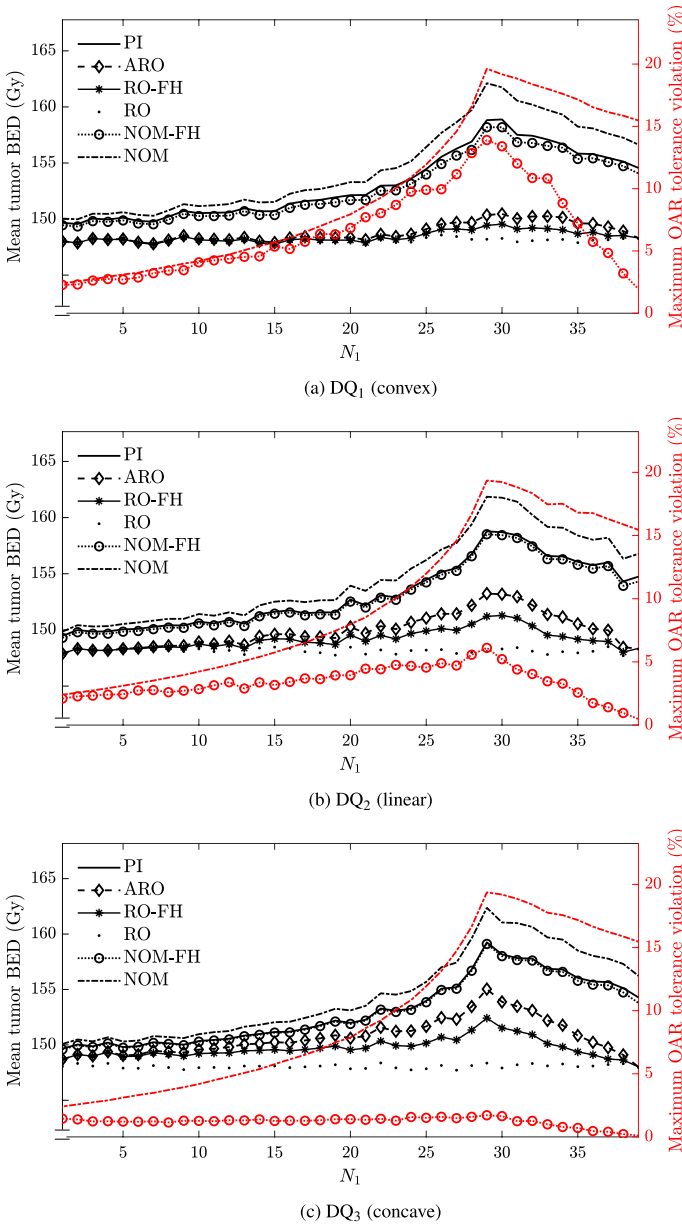


Fig. 11 Change in mean tumor BED and OAR constraint violation when varying the information point N_1 from 1 to $N^{\max} - 1$, for data quality functions $DQ_i(N_1)$, $i = 1, 2, 3$. The maximum OAR BED constraint violation (%) of NOM (dash-dotted) and NOM-FH (dotted, circle marker) is measured against the right vertical axis. Note that the left vertical axis measures displays the mean tumor BED (averaged over all patients and scenarios), while the methods maximize the worst-case tumor BED per patient

planning the stage-1 dose. The good performance of NOM-FH shows that this value diminishes if OAR violations are allowed. NOM-FH does not account for adaptability, and does not take inexactness of biomarker information into account. Nevertheless, it results in higher sample mean tumor BED than the robust methods for any N_1 , and the difference increases from good (concave) to poor (convex) data quality. Thus, Fig. 11 illustrates the trade-off between higher sample mean tumor BED and possible OAR violations that was also observed in Sect. 5.4 (for the entire sample distribution).

The shape of the data quality function influences the optimal moment of biomarker observation only slightly. For all adaptive methods, we find that the peak is more pronounced for high data quality (concave) than low data quality (convex), but it is centered around $N_1 = 29$. In case of convex data quality the peak is relatively flat, indicating a trade-off between observing at $N_1 = 29$ (giving maximum adaptation flexibility) and postponing (waiting for higher data accuracy).

6 Concluding remarks

In this paper we have presented an ARO approach to optimally adapt the treatment length of radiation therapy treatments, using mid-treatment biomarker information. Using an ARO approach, adaptability is taken into account prior to treatment and it provides insight into the optimal stage-2 decisions.

In the case of exact biomarker information, there is sufficient space to adapt, and numerical results show that taking into account both robustness and adaptability is not necessary. In the case of inexact biomarker information, adaptive strategies can use only parameter estimates instead of true parameter values, and may still result in violations if this uncertainty is not accounted for. Accounting for adaptability and inexactness of biomarker information is particularly beneficial when robustness (w.r.t. OAR violations) is of high importance. If minor OAR violations are allowed, NOM-FH is a good performing alternative, which can outperform ARO. NOM-FH and ARO thus yield a trade-off between higher performance and OAR violations. Both the difference in performance and the magnitude of OAR violations of NOM-FH are highly influenced by the data quality (i.e., accuracy of parameter estimates).

The current setting can be extended in several ways. In practice the tumor and OAR α/β values would have to be estimated from actual biomarkers (e.g., imaging, blood-based biomarkers, genotyping), which can be incorporated in the model. Furthermore, the approach can be extended to heterogeneous tumor response (different α/β ratios for different tumor subvolumes), or time-dependent response parameters. Other RT applications may also benefit from ARO, such as re-optimization to account for organ motion or setup errors, optimization using the MR-linac or combining RT with chemotherapy.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11081-021-09709-w>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ajdari A, Ghaté A (2016) Robust fractionation in cancer radiotherapy, preprint [arXiv:2108.03209](https://arxiv.org/abs/2108.03209)
- Ajdari A, Ghaté A, Kim M (2018) Adaptive treatment-length optimization in spatiobiologically integrated radiotherapy. *Phys Med Biol* 63(7):075009
- Ajdari A, Niyazi M, Nicolay N, Thieke C, Jeraj R, Bortfeld T (2019) Towards optimal stopping in radiation therapy. *Radiother Oncol* 134:96–100
- Barker CA, Powell SN (2010) Enhancing radiotherapy through a greater understanding of homologous recombination. *Semin Radiat Oncol* 20(4):267–273
- Baumann M, Krause M, Overgaard J, Debus J, Bentzen SM, Daartz J, Bortfeld T (2016) Radiation oncology in the era of precision medicine. *Nat Rev Cancer* 16(4):234–249
- Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math Program* 99:351–376
- Bindra RS, Goglia AG, Jasin M, Powell SN (2013) Development of an assay to measure mutagenic non-homologous end-joining repair activity in mammalian cells. *Nucleic Acids Res* 41(11):e115
- Böck M, Eriksson K, Forsgren A, Hårdemark B (2017) Toward robust adaptive radiation therapy strategies. *Med Phys* 44(6):2054–2065
- Böck M, Eriksson K, Forsgren A (2019) On the interplay between robustness and dynamic planning for adaptive radiation therapy. *Biomed Phys Eng Express* 5(4):045004
- Bortfeld T, Ramakrishnan J, Tsitsiklis JN, Unkelbach J (2015) Optimization of radiation therapy fractionation schedules in the presence of tumor repopulation. *INFORMS J Comput* 27(4):788–803
- Chan TCY, Mišić VV (2013) Adaptive and robust radiation therapy optimization for lung cancer. *Eur J Oper Res* 231:745–756
- Cox JD (1986) Presidential address: fractionation: a paradigm for clinical research in radiation oncology. *Int J Radiat Oncol Biol Phys* 13:1271–1281
- Dabadghao S, Roy A (2020) Optimal interventions for adaptive robust optimization under time-dependent uncertainty with application to radiotherapy, available at SSRN
- De Ruiter FJCT, Brekelmans RCM, den Hertog D (2016) The impact of the existence of multiple adjustable robust solutions. *Math Program* 160:531–545
- De Ruiter FJCT, Ben-Tal A, Brekelmans RCM, den Hertog D (2017) Robust optimization of uncertain multistage inventory systems with inexact data in decision rules. *Comput Manag Sci* 14(1):45–77
- Ehrgott M, Güler Ç, Hamacher HW, Shao L (2008) Mathematical optimization in intensity modulated radiation therapy. *4OR-Q J Oper Res* 6:199–262
- Fowler JF (1989) The linear-quadratic formula and progress in fractionated radiotherapy. *Brit J Radiol* 62(740):679–694
- Fowler JF (2010) 21 years of biologically effective dose. *Brit J Radiol* 83:554–568
- Ghaté A (2014) Dynamic optimization in radiotherapy. *Inf Tutor Oper Res* 60–74. <https://doi.org/10.1287/educ.1110.0088>
- Hall EJ, Giaccia AJ (2012) *Radiobiology for the radiologist*. Lippincott Williams & Wilkins, Philadelphia
- Iancu DA, Trichakis N (2014) Pareto efficiency in robust optimization. *Manag Sci* 60(1):130–147

- Iancu DA, Trichakis N, Yoon DY (2021) Monitoring with limited information. *Manag Sci* 67(7):4233–4251
- Kehwar TS (2005) Analytical approach to estimate normal tissue complication probability using best fit of normal tissue tolerance doses into the NTCP equation of the linear quadratic model. *J Canc Res Ther* 1(3):168–179
- Kim M, Ghate A, Phillips MH (2012) A stochastic control formalism for dynamic biologically conformal radiation therapy. *Eur J Oper Res* 219:541–556
- Lim GJ, Kardar L, Ebrahimi S, Cao W (2020) A risk-based modeling approach for radiation therapy treatment planning under tumor shrinkage uncertainty. *Eur J Oper Res* 280(1):266–278
- Long T (2015) Optimization problems in radiation therapy treatment planning. PhD thesis, University of Michigan, chapter 5: Adaptive treatment planning for lung cancer
- Mahmoudzadeh H (2015) Robust optimization methods for breast cancer radiation therapy. PhD thesis, University of Toronto, Chapter 4: Pareto robust optimization in breast cancer RT
- Mar PA, Chan TCY (2015) Adaptive and robust radiation therapy in the presence of drift. *Phys Med Biol* 60(9):3599–3615
- Mizuta M, Takao S, Date H, Kishimoto N, Sutherland KL, Onimaru R, Shirato H (2012) A mathematical study to select fractionation regimen based on physical dose distribution and the linear-quadratic model. *Int J Radiat Oncol Biol Phys* 84(3):829–833
- Nohadani O, Roy A (2017) Robust optimization with time-dependent uncertainty in radiation therapy. *IIESE Trans Healthc Syst Eng* 7(2):81–92
- Perkó Z, Bortfeld T, Hong T, Wolfgang J, Unkelbach J (2018) Derivation of mean dose tolerances for new fractionation schemes and treatment modalities. *Phys Med Biol* 63(3):035038
- Saberian F, Ghate A, Kim M (2016) Optimal fractionation in radiotherapy with multiple normal tissues. *Math Med Biol* 33(2):211–252
- Saberian F, Ghate A, Kim M (2016) A theoretical stochastic control framework for adapting radiotherapy to hypoxia. *Phys Med Biol* 61(19):7136–7161
- Saberian F, Ghate A, Kim M (2017) Spatiotemporally optimal fractionation in radiotherapy. *INFORMS J Comput* 29(3):422–437
- Saka B, Rardin RL, Langer MP (2014) Biologically guided intensity modulated radiation therapy planning optimization with fraction-size dose constraints. *J Oper Res Soc* 65:557–571
- Santiago A, Barczyk S, Jelen U, Engenhardt-Cabillic R, Wittig A (2016) Challenges in radiobiological modeling: can we decide between LQ and LQ-L models based on reviewed clinical NSCLC treatment outcome data. *Radiat Oncol* 11:67
- Shepard DM, Ferris MC, Olivera GH, Mackie TR (1999) Optimizing the delivery of radiation therapy to cancer patients. *SIAM Rev* 41(4):721–744
- Somaiah N, Rothkamm K, Yarnold J (2015) Where do we look for markers of radiotherapy fraction size sensitivity? *Clin Oncol* 27:570–578
- Van Leeuwen CM, Oei A, Crezee J, Bel A, Franken NAP, Stalpers LJA, Kok HP (2018) The alpha and beta of tumours: a review of parameters of the linear-quadratic model, derived from clinical radiotherapy studies. *Radiation Oncol* 13:96
- Withers HR (1985) Biological basis for altered fractionation schemes. *Cancer* 55:2086–2095
- Yamkoğlu I, Gorissen BL, den Hertog D (2019) A survey of adjustable robust optimization. *Eur J Oper Res* 277:799–813

Authors and Affiliations

S. C. M. ten Eikelder¹ · A. Ajdari² · T. Bortfeld² · D. den Hertog³

✉ S. C. M. ten Eikelder
stefanteneikelder@gmail.com

A. Ajdari
aajdari@mgh.harvard.edu

T. Bortfeld
tbortfeld@mgh.harvard.edu

D. den Hertog
d.denhartog@uva.nl

¹ Department of Econometrics and Operations Research, Tilburg University, Tilburg, The Netherlands

² Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, USA

³ Department of Operations Management, University of Amsterdam, Amsterdam, The Netherlands