



HAL
open science

Toward semi-automatic biologically effective dose treatment plan optimisation for Gamma Knife radiosurgery

Thomas Klinge, Hugues Talbot, Ian Paddick, Sébastien Ourselin, Jamie R McClelland, Marc Modat

► To cite this version:

Thomas Klinge, Hugues Talbot, Ian Paddick, Sébastien Ourselin, Jamie R McClelland, et al.. Toward semi-automatic biologically effective dose treatment plan optimisation for Gamma Knife radiosurgery. *Physics in Medicine and Biology*, 2022, 67 (21), pp.215001. 10.1088/1361-6560/ac8965 . hal-03943763

HAL Id: hal-03943763

<https://hal.science/hal-03943763>

Submitted on 16 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License

PAPER • OPEN ACCESS

Toward semi-automatic biologically effective dose treatment plan optimisation for Gamma Knife radiosurgery

To cite this article: Thomas Klinge *et al* 2022 *Phys. Med. Biol.* **67** 215001

View the [article online](#) for updates and enhancements.

You may also like

- [Assessing small-volume spinal cord dose for repeat spinal stereotactic body radiotherapy treatments](#)
Lijun Ma, Neil Kirby, Renee Korol *et al.*
- [A feasibility study of spatiotemporally integrated radiotherapy using the LQ model](#)
M Kim and M H Phillips
- [Optimization of spatiotemporally fractionated radiotherapy treatments with bounds on the achievable benefit](#)
Melissa R Gaddy, Sercan Yildiz, Jan Unkelbach *et al.*

VERIQA
RT MonteCarlo 3D
Plan selected. Plan verified.
In less than 3 minutes.

Automated. Independent. Web-Based.

PTW THE DOSIMETRY COMPANY

Explore the benefits of streamlined patient QA



PAPER

OPEN ACCESS

RECEIVED

9 November 2021

REVISED

26 July 2022

ACCEPTED FOR PUBLICATION

12 August 2022

PUBLISHED

17 October 2022

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Toward semi-automatic biologically effective dose treatment plan optimisation for Gamma Knife radiosurgery

Thomas Klinge^{1,2,3,*}, Hugues Talbot⁴, Ian Paddick⁵, Sébastien Ourselin³,
Jamie R McClelland^{1,2} and Marc Modat³

¹ Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS), Dept. Medical Physics and Biomedical Engineering, University College London, London, United Kingdom

² Centre for Medical Image Computing, Dept. Medical Physics and Biomedical Engineering, University College London, London, United Kingdom

³ School of Biomedical Engineering & Imaging Sciences, King's College London, London, United Kingdom

⁴ CentraleSupélec, Université Paris-Saclay, Inria, Gif-sur-Yvette, France

⁵ Queen Square Gamma Knife Centre, National Hospital for Neurology and Neurosurgery, London, United Kingdom

* Author to whom correspondence should be addressed.

E-mail: thomas.klinge.17@ucl.ac.uk

Keywords: Gamma Knife, BED, treatment planning, inverse planning, discrete non-convex optimisation

Abstract

Objective. Dose-rate effects in Gamma Knife radiosurgery treatments can lead to varying biologically effective dose (BED) levels for the same physical dose. The non-convex BED model depends on the delivery sequence and creates a non-trivial treatment planning problem. We investigate the feasibility of employing inverse planning methods to generate treatment plans exhibiting desirable BED characteristics using the per iso-centre beam-on times and delivery sequence. *Approach.* We implement two dedicated optimisation algorithms. One approach relies on mixed-integer linear programming (MILP) using a purposely developed convex underestimator for the BED to mitigate local minima issues at the cost of computational complexity. The second approach (local optimisation) is faster and potentially usable in a clinical setting but more prone to local minima issues. It sequentially executes the beam-on time (quasi-Newton method) and sequence optimisation (local search algorithm). We investigate the trade-off between time to convergence and solution quality by evaluating the resulting treatment plans' objective function values and clinical parameters. We also study the treatment time dependence of the initial and optimised plans using BED₉₅ (BED delivered to 95% of the target volume) values. *Main results.* When optimising the beam-on times and delivery sequence, the local optimisation approach converges several orders of magnitude faster than the MILP approach (minutes versus hours–days) while typically reaching within 1.2% (0.02–2.08%) of the final objective function value. The quality parameters of the resulting treatment plans show no meaningful difference between the local and MILP optimisation approaches. The presented optimisation approaches remove the treatment time dependence observed in the original treatment plans, and the chosen objectives successfully promote more conformal treatments. *Significance.* We demonstrate the feasibility of using an inverse planning approach within a reasonable time frame to ensure BED-based objectives are achieved across varying treatment times and highlight the prospect of further improvements in treatment plan quality.

1. Introduction

In stereotactic radiosurgery (SRS) treatments with the Gamma Knife (GK), an array of collimated Cobalt-60 sources is used to precisely deliver therapeutic radiation to an intracranial target volume (TV). Since its introduction over 50 years ago, the GK has become a standard delivery method for SRS procedures (Podgorsak *et al* 1989, Schulder and Patil 2008) and the GK unit has undergone significant updates. So far, there have been

five commercially available models that introduced significant changes to the treatment delivery (e.g. geometry of the radiation unit, patient positioning system, collimator selection).

Given that treatments are reported in terms of total physical dose, the significant changes in the time-domain of GK treatment delivery are currently not taken into account under the assumption that GK treatments are single fraction acute exposures. However, it has been shown that the typical time frame of GK SRS treatments would allow for repair of sublethal radiation damage, an effect that is known to be exposure time dependent (Hallgren *et al* 2019).

There have been a number of studies that demonstrated how the biologically effective dose (BED) across different patients treated with the GK varies with the treatment time despite being evaluated for the same physical dose levels (Hopewell *et al* 2012, 2013, Millar *et al* 2015). The model used for these publications was initially derived for a generalised fractionated protracted irradiation under consideration of incomplete repair processes (Millar and Canney 1993) and subsequently, a version of this model using two repair-rates was used in an iso-effect fit to extract the model parameters (Pop *et al* 2000). More recent treatment outcome studies demonstrated an improved correlation with BED compared to the physical dose for single iso-centre trigeminal neuralgia (Tuleasca *et al* 2019), and multi iso-centre acromegaly (Graffeo *et al* 2020) and pituitary adenoma (Graffeo *et al* 2021) SRS treatments (using a simplified version of the BED for retrospective analysis (Jones and Hopewell 2018)). In addition, it has recently been demonstrated how the BED will change with the sequence in which the iso-centres are delivered and also due to unscheduled interruptions in treatment delivery (Klinge *et al* 2021). This is due to the fact that the BED model tracks all changes in the in-patient dose-rate distribution throughout the entire treatment delivery, including beam-off periods. Any attempt at BED-based treatment planning will thus have to consider the exact delivery sequence in addition to the conventional treatment planning parameters.

These studies raise the need to investigate the value of BED-based treatment planning, especially given the large potential for variations in treatment time, number of iso-centres and possible collimator combinations enabled with the modern GK Perfexion (PFX) and Icon⁶. While there are inverse planning tools available for physical dose GK treatment planning (Sjölund *et al* 2019), they typically rely on formulating a convex treatment planning problem that can be solved efficiently (Levivier *et al* 2018). Using a BED model that incorporates incomplete repair intervals adds a layer of complexity to the problem since the induced radiation damage of an individual iso-centre is linked to the properties of the other iso-centres and can not be treated independently. The nonlinear and non-convex nature of the BED model makes it inherently a hard problem to solve to optimality.

The goal of this study is to explore the value and feasibility of using inverse planning to create BED-based GK treatment plans via adjustments to the beam-on times and delivery sequence (or shot order) using both local and global optimisation techniques. Starting from a plan manually generated using the conventional approach relying only on physical dose, we optimise the original treatment plan in terms of BED by changing the sequence and exposure times only (fixed iso-centre locations and collimator settings). Local optimisation approaches are used to quickly solve for the most beneficial beam-on times (gradient-based (Byrd *et al* 1995)) and delivery sequence (local search (Johnson and McGeoch 1997)). While they are fast, the non-convexity of the BED means that there is also the risk of getting 'stuck' in local minima, which in turn can result in less optimal treatment plans. To tackle the problem of non-convexity, a mixed-integer programming (MIP) approach that can simultaneously solve the discrete delivery sequence and continuous per iso-centre beam-on times was developed. This is realised using a convex hull of the BED model to create a mixed-integer linear programming (MILP) problem⁷. While this approach does consider the entire solution space, the global optimum of such a non-convex function can only be guaranteed with global solvers and (possibly) infinite time. This approach is a trade-off between considering the vast solution space and converging within non-infinite time.

Both optimisation strategies are applied to a cohort of vestibular schwannoma treatments, 14 cases in total, using the original physical dose plans as the starting point. The performances of these approaches are evaluated in terms of their final objective function values, the quality of the optimised treatment plans, and the required time to convergence to determine the clinical feasibility and quality of the individual approaches.

⁶ For the cohort in this study, treatment times were 18.2–75.3 min using 3–17 iso-centres.

⁷ The interested reader is referred to Nocedal and Wright (2006) for a general introduction to optimisation problems (e.g. convexity, continuous optimisation, complexity), to Chen *et al* (2009) for an overview of MIP problems, and to Burer and Letchford (2012) for a more specific review of non-convex MINLP programming (e.g. convex hull). The derivation of the MILP approach using the convex hull of the BED is described in the appendix C.

2. Method

The proposed semi-automatic creation of BED-optimised treatment plans consists of two steps: iso-centre setup and optimisation. In the initial setup step, the location and collimator settings of the iso-centres are defined. This could be done by conventional (manual) physical dose treatment planning, an initial ‘filling’ step or any physical dose inverse planning module. After this initial phase, we optimise the per iso-centre beam-on times and the sequence of delivery.

2.1. Treatment planning problem

2.1.1. BED model

The BED model used in the current study was originally developed by Millar and Canney (1993) as an extension of the linear-quadratic model (Fowler 1989) and refined by Pop *et al* (2000) to include two repair-rates and determine the α/β ratio (Fowler 1989) for central nervous system tissue. The per voxel BED can be determined as follows:

$$\text{BED} = D_T + \frac{1}{\beta} \left[\frac{\Phi(\Xi, \mu_1) + c \cdot \Phi(\Xi, \mu_2)}{1 + c} \right] \sum_{n=1}^N d_n^2 \quad (1)$$

$$\Phi(\Xi, \mu) = \frac{\frac{2}{\mu} \sum_{j=1}^N \left[d_j^2 \frac{\left[\delta t_j - \frac{1}{\mu} (1 - e^{-\mu \delta t_j}) \right]}{\delta t_j^2} - \frac{1}{\mu} \sum_{k=1}^{j-1} d_k d_j \frac{e^{-\mu(t_j - t_k)} (e^{\mu \delta t_k} - 1) (e^{-\mu \delta t_j} - 1)}{\delta t_k \delta t_j} \right]}{\sum_{n=1}^N d_n^2}}, \quad (2)$$

with the number of iso-centres N , the total dose D_T , per iso-centre dose d , per iso-centre beam-on time δt , start time of the iso-centres t , repair-rate μ , partition coefficient c ⁸, and treatment protocol Ξ describing the dose-rate time dependence. The model parameters as determined by Pop *et al* are shown in table 4 in the appendix. Expressed in terms of the per iso-centre dose-rate \dot{d} , the BED equations become:

$$\text{BED} = \sum_{j=1}^N \dot{d}_j \delta t_j + \frac{1}{\beta} \left[\frac{\Psi(\Xi, \mu_1) + c \cdot \Psi(\Xi, \mu_2)}{1 + c} \right] \quad (3)$$

$$\Psi(\Xi, \mu) = \frac{2}{\mu} \sum_{j=1}^N \left[\dot{d}_j^2 \left[\delta t_j - \frac{1}{\mu} (1 - e^{-\mu \delta t_j}) \right] - \frac{1}{\mu} \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j e^{-\mu(t_j - t_k)} (e^{\mu \delta t_k} - 1) (e^{-\mu \delta t_j} - 1) \right]. \quad (4)$$

The model determines the BED-based on the dose-rate profile throughout the treatment. Each individual iso-centre delivers the dose at a constant dose-rate for a period of time. Practically, this means the iso-centres become individual fractions that are defined by a start time, beam-on time and dose-rate. This formulation implicitly includes periods of beam-off time in-between the delivery of iso-centres where the patient is repositioned and a new set of collimators can be selected. For every iso-centre, all preceding iso-centres are taken into account to determine the residual sub-lethal radiation damage (see nested sum in equation (4)). Consequently, changing the order of delivery will change the BED formulation in two ways: firstly, the nested sum now has to be evaluated over the new sequence, and secondly the starting times of the individual iso-centres need to be updated. Even with only the beam-on times as a variable, the BED (δt) constitutes a non-convex function (i.e. Hessian matrix not positive semi-definite for all δt).

2.1.2. Problem definition

After the iso-centre definition (location and shape), the individual per iso-centre dose-rate distributions are fixed, leaving both the per iso-centre beam-on times and the sequence of delivery as free variables. The treatment planning problem using a BED-based objective function $f(\text{BED})$ can then be described as follows:

$$\begin{aligned} & \arg \min_{\delta t_j, seq} f(\text{BED}) \\ & \text{subject to: } \delta t_j \geq 0, \quad \delta t_j \in \mathbb{R}, \quad j \in \{1 \dots N\} \\ & \quad \quad \quad seq \in \mathfrak{S}_N \end{aligned} \quad (5)$$

The δt_j are the beam-on times of the individual iso-centres j and seq is their delivery sequence which is constrained to be a member of the symmetric group \mathfrak{S}_N on the finite set $\{1, \dots, N\}$ of N iso-centres. This group is comprised of all $N!$ possible permutations of the delivery sequence. Due to the characteristics of the BED model and the discrete nature of the sequence, this constitutes a non-convex mixed-integer nonlinear programming problem (MINLP).

⁸The partition coefficient determines the relative contributions from the two repair-rates.

2.1.3. Objective function

All optimisation approaches in this study (MILP, gradient-based methods, combinatorial optimisation) use the same objective in order to enable comparisons across methods. The chosen objective function is comprised of the weighted sum of two objectives: the mean BED under-exposure inside the target volume (TV) and the mean BED over-exposure in the normal tissue (NT) around the target (Rim)

$$f(\mathbf{BED}) = w_{TV} \sum_{v \in TV} \frac{\max(0, \text{BED}_{ref} - \text{BED}_v)}{N_{TV}} + w_{Rim} \sum_{v \in Rim} \frac{\max(0, \text{BED}_v - \text{BED}_{thres})}{N_{Rim}}. \quad (6)$$

Depending on whether a given voxel is part of the TV or the Rim, its BED_v will be compared to the prescription value BED_{ref} or the upper permitted threshold BED_{thres} and the result scaled according to the size of the volume of interest (VOI) (N_{TV}, N_{Rim}). The weighting factors w can be used to focus the optimisation more on target coverage or selectivity.

2.1.4. Solving the treatment planning problem

While a brute force approach could theoretically be employed to identify the most beneficial delivery sequence of a given treatment plan, the fact that the number of possible sequences increases with the factorial of the iso-centre number makes this approach impractical for all but the simplest treatments. The first naive way of tackling this problem is to explore efficient and established local approaches to individually optimise the beam-on time (gradient-based) and the delivery sequence (local search). See section 2.2 for a description of the algorithms used in this study.

However, due to the non-convex nature of the BED, there is the possibility for the optimiser to get stuck in a local minimum and never reach the globally optimal value. Ideally, one would like to simultaneously explore all possible combinations of beam-on times and delivery sequences to avoid reaching convergence at a sub-optimal solution. While there are global optimisation techniques, the nonlinear non-convex nature of the BED together with the mixed-integer variables and the large space of possible permutations of the delivery sequence makes solving this problem to optimality prohibitively expensive in terms of computing power and time.

To tackle both problems mentioned above, a ‘convex mixed-integer underestimator’ of the ‘full BED’ was developed (see section 2.3). This reduces the problem to a mixed-integer linear programming problem (MILP). While this convex MILP approach can still not guarantee global optimality, it does consider the entire search space during optimisation and allows to simultaneously optimise the beam-on time and delivery sequence.

2.2. Local optimisation approaches

Two established algorithms (see 2.2.1 and 2.2.2) are used to optimise the beam-on time and delivery sequence independently. To achieve the best possible plan with these local approaches, they are executed in an alternating fashion until no further improvement of the objective function is observed.

2.2.1. Beam-on time optimisation

To find the most beneficial set of beam-on times, given the defined objectives and constraints, the limited-memory Broyden–Fletcher–Goldfarb–Shanno bound-constrained (L-BFGS-B) algorithm (Byrd *et al* 1995) for the optimisation of nonlinear problems is used. The L-BFGS-B algorithm utilises the gradient and an approximation of the Hessian (2nd derivative) to guide the direction of the optimisation. For this study the implementation of the `scipy.optimize` (Virtanen *et al* 2020) library was used with the BED-based objective function as an input. The gradient is approximated numerically during the optimisation.

2.2.2. Delivery sequence

In order to optimise the sequence of delivery, the treatment planning problem is expressed as a travelling salesperson problem (TSP). Instead of the travelling distance, the BED-based objective function is evaluated to determine the quality of a given sequence. To solve this TSP, a common local search algorithm, the 2-opt approach (Croes 1958, Johnson and McGeoch 1997), is adapted to the BED treatment planning scenario. Generally, the 2-opt approach starts with an initial solution, creates a new connection between 2 nodes on the route and solves the order in which all the other nodes will now be visited to create a new, potentially improved solution. Since in our case, there is no distance measure to identify the potential for improved routing (e.g. longest distances, cross-over), the nodes (iso-centres) to be connected (delivered one after the other) are determined iteratively. Starting from the first iso-centre, the objective function is evaluated for which iso-centre should be delivered next (3, ..., N) until an improvement is found. If no improvement is found, the next iso-centre in the current sequence is chosen as a candidate to be connected to the others. If an improvement was found, the algorithm is started from the beginning. The algorithm is stopped when no more improvement is found for an entire iteration over all iso-centres. Since a treatment delivery is not a closed loop, this approach

never changes the starting iso-centre. To mitigate this issue, the 2-opt algorithm is executed once for every iso-centre in the starting position. The initial solution is always created from the current best solution.

2.3. Convex mixed-integer underestimator approach (Convex MILP)

In general, MINLP problems are difficult to solve since even their continuous relaxation constitutes a non-convex problem that requires a global optimisation approach to solve. To mitigate this issue, we propose the use of a convex relaxation of the full BED model that can be solved to optimality with available MILP solvers. The ‘relaxed BED’ can then be used as an underestimator for the full problem.

Specifically, the ‘relaxed BED’ will be defined by the bounds on the beam-on times δt_j of the individual iso-centres j . When the bounds are initially loose, a large range of beam-on times is permitted for all iso-centres and optimising it returns a lower bound of the feasible objective function value for the ‘full BED’. By iteratively tightening the bounds of the underestimator around its’ current optimum, we can minimise the objective function value (of the ‘full BED’) until the full model and its’ relaxation converge towards each other. For this study, CPLEX 12.10⁹ was used as the MILP solver.

2.3.1. Defining the convex hull of the BED model

The first step in the creation of the ‘relaxed’ BED model is to substitute all nonlinear terms in equation (4) (i.e. all terms including the exponential function) with a new variable. In our case (constant dose-rate), we are left with linear and trilinear terms for the intra and inter iso-centre interactions, respectively.

In order to achieve a convex relaxation, the dual envelope method for general multilinear terms, as described by Costa and Liberti (2012), is applied. The convex hull is defined from the support points created by the combinations of the lower and upper bounds of the variables. Applied to the BED model, these bounds are determined by the values of the substituted nonlinear terms at the lower and upper bounds of the beam-on times δt . The Lagrange multipliers that are used to navigate the convex hull have now become the new decision variables.

To constrain the optimisation to physically deliverable treatment plans, the timing information, i.e. beam-on times and starting times of the individual iso-centres, has to be recovered from the nonlinear terms. This is achieved by applying a piecewise linearisation (PWL) (Lin et al 2013) to each individual substituted term.

In this manner, convex hulls are created for every addend of the nested sum in the full BED model (equation (4)). In total, there are $\frac{N(N-1)}{2}$ convex hulls of the interaction terms for a given delivery sequence.

Due to the limitation of the optimisation variables to beam-on times and sequence, the dual envelope formulation of ‘outer sum’ terms is effectively a linearisation. However, the presented approach could still be applied if the dose-rate were to be included as a variable (due to its’ dependence on iso-centre location and collimator selection). The approach would then create a convex hull equivalent to the McCormick envelopes (bilinear). Using this general approach for multilinear functions allows for great flexibility for future expansion of the model.

2.3.2. Delivery sequence

If the delivery sequence can change, then every iso-centre could feasibly interact with all other iso-centres.

To accommodate all possible delivery sequences, the model is extended to include all $N(N-1)$ interaction terms with a binary variable (h_{jk}) signifying whether the term is active or not, based on the variable start times (t_j , t_k) of the iso-centres

$$\Psi(\Xi, \mu) = \frac{2}{\mu} \sum_{j=1}^N \dot{d}_j^2 \left[\delta t_j - \frac{1}{\mu} (1 - e^{-\mu \delta t_j}) \right] - \frac{2}{\mu^2} \sum_{j=k} h_{jk} \dot{d}_k \dot{d}_j e^{-\mu(t_j-t_k)} (e^{\mu \delta t_k} - 1) (e^{-\mu \delta t_j} - 1) \quad (7)$$

$$h_{jk} = \begin{cases} 1, & \text{if } t_j(\delta t, seq) > t_k(\delta t, seq) \\ 0, & \text{if } t_j(\delta t, seq) \leq t_k(\delta t, seq) \end{cases} \quad (8)$$

Using a set of logical constraints, we can ensure that only the appropriate terms describing a feasible delivery sequence are active.

Since the time in-between the beginning of the delivery of two iso-centres ($t_j - t_k$) depends on the specific sequence, its’ bounds need to accommodate this and do not shrink with the beam-on time bounds anymore. As a result, the number of support points of the PWL of the interaction term needs to be increased to ensure that the ‘relaxed’ BED will be close to the ‘full’ BED when the beam-on time is tightened, leading to a more complex

⁹ ILOG CPLEX Optimization Studio V12.10.0 User’s Manual, International Business Machines Corporation, 2021.

Table 1. Overview of patients treated for vestibular schwannoma with the GK PFX. Listed are the case number, the prescribed dose, the prescription iso-dose, the treatment time T (including beam-off periods), the number of iso-centres N_{iso} , the reference dose-rate on the day of treatment, and the TV.

Case	Pres. Dose [Gy]	Pres. Iso-dose [%]	T [min]	N_{iso}	Reference Dose-rate [Gy min ⁻¹]	TV [cm ³]
1	12	52	44.7	4	1.7	0.2
2	12	51	32.4	9	2.4	3.6
3	12	50	46.2	17	2.1	5.1
4	12	43	62.7	16	1.9	12.1
5	12	44	61.2	11	1.7	2.5
6	12	46	75.3	11	1.7	4.7
7	12	44	48.3	12	1.7	5.9
8	12	50	74.8	17	1.6	1.5
9	12	60	28.1	12	3.4	1.8
10	12	58	18.2	3	3.1	0.1
11	12	50	37.3	12	2.8	3.1
12	12	50	58.1	17	2.1	7.0
13	12	42	51.4	12	1.9	4.6
14	12	46	45.3	11	1.9	1.3

TV: Target Volume

optimisation problem. The number of PWL support points is a parameter of the optimisation problem that allows to control the trade-off between model complexity and underestimator accuracy.

2.4. Workflow

2.4.1. Data import

The cohort for this study consists of 14 cases of vestibular schwannoma treated with the GK PFX. An overview is shown in table 1. All cases were planned with a prescription dose of 12 Gy and a varying prescription iso-dose. Treatment times ranged from 18 to 75 min, assuming a gap of 0.06 min between iso-centre deliveries.

To load the previously created treatment plans into the BED treatment planning framework, a research version of GammaPlan 10.1 is used that allows exporting per-isocentre dose distributions. The exported data includes the dose distributions in a $31 \times 31 \times 31$ voxel grid covering the target volume (TV), a binary mask of the TV, the original iso-centre setup (shape, beam-on times and delivery sequence), and a registration matrix to the imaging study.

With this information, the original dose and BED distributions can be calculated and later compared to the optimised versions.

2.4.2. Problem setup

Since using all voxels (almost 30 000) for the optimisation is impractical and would include regions with very limited dose, the problem is constrained to a region of interest (ROI) where a meaningful dose contribution can occur. To this end, a Rim structure is created, which can also drive the optimisation towards normal tissue (NT) sparing. This is achieved by applying 4 iterations of binary dilation to the TV mask using the `scipy.ndimage.morphology` library from Virtanen *et al* (2020), which effectively grows the TV outwards by 4 voxels. This approach ensures that the number of voxels is approximatively split into a ratio of 60% Rim and 40% TV for all cases of the cohort.

For the present study, both the weights and reference/threshold BED are chosen to be equivalent ($w_{TV} = w_{Rim} = 100$, see equation (6)). A treatment time of 60 min is taken as a reference for the BED, as suggested by Jones and Hopewell (2018). With a prescription dose of 12 Gy, this corresponds to a BED of 53.95 Gy_{2.47}¹⁰. Using the same objective function for all optimisation runs allows for comparing the performance of the different approaches taken to optimise the problem.

2.4.3. Optimisation Runs

We apply the two introduced approaches (local and MILP) to every case in the cohort. For each of the approaches, beam-on time optimisations with and without sequencing (BO-Seq-opt/BO-opt) are executed on a computing cluster using the same setup. Hence, there are four individual optimisation runs per patient. The

¹⁰ Gy_{2.47} is used to signify BED instead of physical dose (Gy). The subscript denotes the α/β -ratio used to determine the BED. α/β -ratio and other BED parameters are shown in table 4 (derived by Pop *et al* 2000).

initial beam-on time is set to 1 min for every case, to avoid accidentally starting from a near-optimal position, e.g. where the original treatment plan was already close to the chosen BED reference.

Due to the large combinatorial complexity for the MILP BO-Seq-opt scenario, a hard limit for the maximum amount of computation time spent in a single iteration is set (deterministic time limit 20e6 ticks corresponds to 20.8–36.8 h¹¹) in order to ensure convergence within a reasonable amount of time. If that time limit is exceeded, the currently best solution will be used to update the model for the next iteration. To ensure a feasible solution is found in every iteration, a pre-solving step is executed (i.e. beam-on time optimisation with a fixed sequence). Further details on the parameters chosen for the different optimisation runs can be seen in table 5 in the appendix.

2.4.4. Plan quality indices

The indices describing the treatment plan quality used in this study are defined using following volume ratios:

$$\text{Coverage: } C = \frac{V(\text{PIV} \cap \text{TV})}{V(\text{TV})}. \quad (9)$$

$$\text{Selectivity: } S = \frac{V(\text{PIV} \cap \text{TV})}{V(\text{PIV})}. \quad (10)$$

$$\text{Paddick Conformity Index: } \text{PCI} = \frac{V(\text{PIV} \cap \text{TV})^2}{V(\text{PIV}) \times V(\text{TV})}. \quad (11)$$

The coverage describes the fraction of the TV that is covered by the prescription iso-dose volume (PIV), while the selectivity describes the fraction of the PIV that is inside the TV. The PCI (Paddick 2000) combines the two proportions to give a measure of conformity. For clinical considerations, it is commonly essential to maintain a particular coverage. Providing the individual values gives the treatment planner additional information (i.e. conformity and how it is ‘split’ between coverage and selectivity). The D_{95} and BED_{95} are defined as the minimum dose and BED values delivered to 95% of the TV.

2.4.5. Analysis

The final objective function values are compared across the entire cohort in order to measure the performance of the individual approaches against each other and determine the benefit of explicitly optimising the sequence over using the beam-on times to compensate for the fixed sequence. The clinical scores for coverage, selectivity and PCI are used to determine the initial plan quality in terms of dose and BED, the benefit of the optimisation methods and how well they correlate to the objective function values. Part of the motivation for the BED optimisation is to not only reduce the intra-patient variability but also the inter-patient variability. To visualise this, the original D_{95} and BED_{95} values in relation to the treatment time are compared to the ones obtained from the optimised treatment plans. In addition (BE)DVHs and dose/BED distributions of a selection of patients are used to qualitatively describe general changes introduced by the optimisations.

3. Results

In this section, the performance of the four different optimisation runs in terms of the minimisation of the objective function (see section 3.1), the resulting quality parameters of the treatment plans (see section 3.2), and examples of the BED-Volume histograms (BEDVHs) and dose/BED distributions (see section 3.3) are presented.

For four of the cases (03, 04, 08, 12) the MILP BED optimisation approach with sequencing did not fully converge, i.e. in the individual iterations of solving the convex underestimator problem no integer-optimal solution was found within the given limit on the computation time¹². Thus, the current best integer solution is returned and used to update the beam-on time. When optimising the sequence of delivery, the complexity of the problem drastically increases with the number of iso-centres which for these cases was especially large (16 and 17 compared to the maximum of 12 for the rest of the cohort. A plan with 12 iso-centres yields $12! = 4.8e8$ possible sequences, while with 17 iso-centres this number increases to $17! = 3.6e14$). Since these cases are not optimised to completion, they will be excluded from the analysis when comparing the individual optimisation approaches to not skew the results for this optimisation method.

¹¹ Deterministic time ticks will be the same for repeated solves even with different loads on the system slowing down the ‘real-world’ optimisation time (seconds/hours).

¹² Optimisation parameters ensure a result within approximately 14 d when running on a cluster using 48 threads.




Figure 1. Objective function values for the original treatment plan and the four optimisation scenarios using the local and MILP approaches with and without sequencing. Adding sequencing to the optimisation leads to clear improvements, where the local optimisation reaches close to the performance of the MILP (excluding cases that exceeded the time limit).

3.1. Performance of the optimisation approaches

The final objective function values are used to assess the performance of the different optimisers. Figure 1 shows an overview of the objective function value for the original treatments compared to the four different optimisation runs. There are large variations for the initial objective function values and all cases in the cohort exhibit significant improvements when using any of the optimisation approaches. Post-optimisation, the objective function values are improved by a factor of approximately 1.25–6.10.

Examining the different optimisation regimens more closely, it becomes evident that the use of the delivery sequence as a variable for the optimisation unlocks the potential for further improvements of the objective function compared to optimising only the beam-on time. Using the local beam-on time optimisation as a reference, the average improvement in objective function value when using sequencing in the optimisation is around 5.86% and 6.95% for the local and MILP optimisations, respectively (see objective function values in table 2).

In comparison, the differences between the local and MILP optimisations are much smaller. When only using the beam-on time optimisation, the optimisation results are similar. For the beam-on time and sequencing runs, the difference is increased but still small (in the order of 1% improvement when using the MILP optimisation). At the same time, the optimisation times increase significantly with the number of iso-centres. For example, the optimisation time for case 09 is considerably increased from 0.1 min for the local beam-on time optimisation to 4.4 min when optimising the sequence as well. With the MILP approach, the optimisation time increases from 3.5 min to over 200 h due to the incredibly large number of possible delivery sequences. The local approach requires only a fraction of that time ($1/2900$) to reach an objective function value within 2.1% of the MILP result.

3.2. Treatment plan quality.

To investigate the quality of the optimised treatment plans, the selectivity, coverage, and PCI are evaluated in terms of the prescribed dose (original plans) and prescribed BED (original and optimised plans). The treatment time relationship is assessed using the D_{95}/BED_{95} values of the treatment plans. Figure 2 gives an overview of the quality parameters for the entire cohort. It can be observed that the original (physical dose) plan favoured coverage (min 95%) over selectivity (min 74%). Evaluating the same original plan in terms of the prescribed BED shows that BED coverage was generally even higher while the selectivity was lower (min 59%). This trade-off at the cost of BED selectivity leads to a lower PCI across the cohort of original treatments.

After optimising these treatment plans, the result shows a slight improvement over the original physical dose PCI. Furthermore, the optimised treatment plans exhibit very similar selectivity and coverage values. Since the generalised objective function used for all cases equally weighs coverage of the TV against sparing of the Rim structure, this behaviour is to be expected.

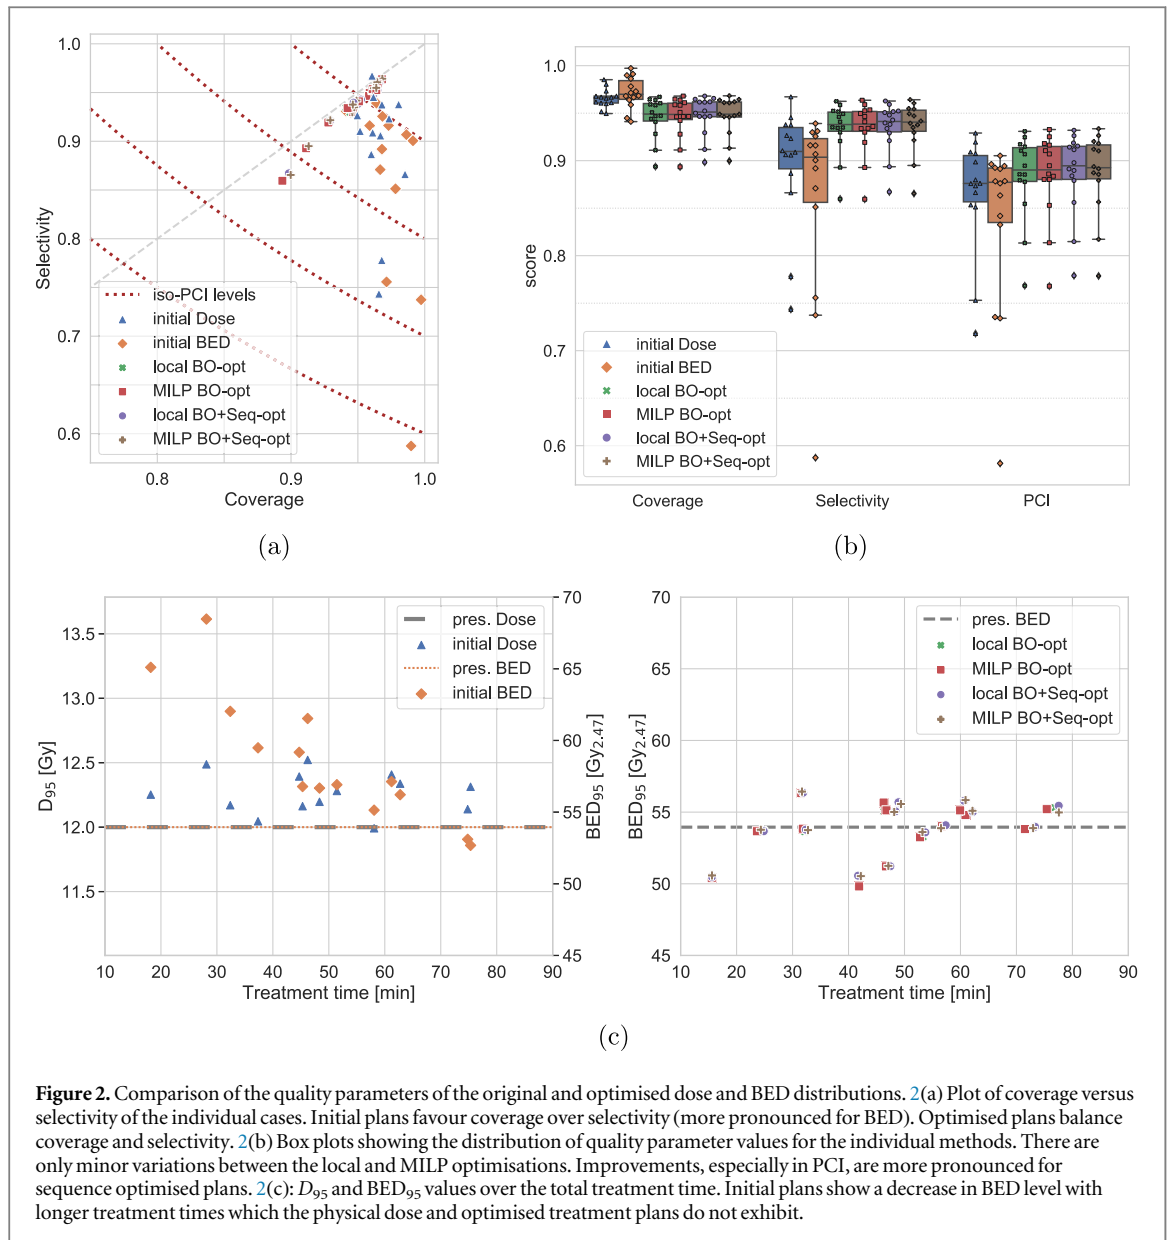
Table 2. Performance Overview: shown are the number of iso-centres N_{iso} , the time T to reach convergence, the final objective function value and the relative difference to the baseline of a local beam-on time optimisation.

Case	Niso	Initial		Beam-on time optimisation					Beam-on time + Sequence optimisation					
		obj.	Δ_{obj} [%]	Local		MILP			local			MILP		
				T [min]	obj.	T [min]	obj.	Δ_{obj} [%]	T [min]	obj.	Δ_{obj} [%]	T [min]	obj.	Δ_{obj} [%]
01	4	243.69	164.00	0.01	92.31	0.14	92.31	0.00	0.02	89.65	-2.88	0.67	88.25	-4.40
02	9	49.40	90.33	0.07	25.95	0.60	25.93	-0.08	0.58	25.25	-2.72	159.25	24.92	-3.98
03	17	58.52	119.17	0.41	26.70	6.64	26.64	-0.21	24.85	25.39	-4.90	17 469.19 ^a	25.32	-5.18
04	16	48.53	31.48	0.45	36.91	3.90	36.90	-0.02	15.67	33.84	-8.33	15 429.26 ^a	34.77	-5.78
05	11	57.10	48.45	0.12	38.46	0.61	38.45	-0.03	1.11	34.41	-10.53	6481.12	34.19	-11.10
06	11	99.02	36.99	0.08	72.28	0.63	72.28	0.00	0.92	64.99	-10.08	3956.95	63.84	-11.67
07	12	45.04	56.54	0.12	28.77	1.25	28.69	-0.27	2.42	26.83	-6.75	10 072.48	26.30	-8.58
08	17	42.63	114.15	0.32	19.91	5.41	19.79	-0.60	21.88	19.04	-4.36	19 043.28 ^a	19.19	-3.58
09	12	260.18	491.65	0.13	43.98	3.46	43.56	-0.94	4.14	43.50	-1.08	12 085.30	42.62	-3.09
10	3	432.85	366.55	0.01	92.78	0.16	92.77	-0.01	0.02	87.04	-6.18	0.58	87.02	-6.20
11	12	107.27	112.95	0.13	50.37	1.60	50.27	-0.21	1.98	49.39	-1.96	10 039.29	48.79	-3.15
12	17	54.81	24.76	0.35	43.93	6.12	43.92	-0.03	15.70	41.57	-5.39	19 896.36 ^a	43.96	0.07
13	12	115.74	53.52	0.15	75.39	1.71	75.38	-0.01	1.40	65.11	-13.64	6650.02	64.80	-14.06
14	11	147.64	57.36	0.10	93.82	1.12	93.81	-0.01	1.12	91.21	-2.78	5567.23	90.77	-3.25
Mean:	11.71	125.89	126.28	0.17	52.97	2.38	52.91	-0.17	6.56	49.80	-5.83	9060.78	49.63	-6.00
w/o ^b	9.70	155.79	147.83	0.09	61.41	1.13	61.35	-0.16	1.37	57.74	-5.86	5501.29	57.15	-6.95

MILP: mixed-integer linear programming.

^a Case did not fully converge within the time limit.

^b Mean value excluding cases that did not converge within the time limit.



To distinguish between the different optimisation scenarios, the distribution of the quality parameters (by optimisation method) across the cohort is shown in figure 2(b). Variations across the individual optimisation approaches are small and the only meaningful differences can be observed between optimisations with and without sequencing. Using the sequence optimisation leads to slight improvements in both selectivity and coverage and thus are more pronounced in the PCI values. On average these cases show a PCI of 0.875 compared to 0.872 and 0.815 for the beam-on times only optimisation and original BED, respectively. The optimised PCI values highlight also an improvement over the original dose PCI of .853 on average. Table 3 presents an overview of the mean and range of quality parameters.

One motivation for the use of BED-based treatment plans is to reduce the inter-patient variability due to the variable timings of the dose deliveries (see Jones and Hopewell 2018). The treatment time dependence of the D_{95} and BED_{95} values in figure 2(c) exhibits a reduction of the BED with increasing treatment time (52.7–68.5 Gy_{2,47}) while there is no time dependence in terms of dose (12–12.5 Gy). After the BED optimisation, the BED_{95} values are distributed around the prescribed BED without any observable treatment time dependence. The range of BED_{95} values across all optimisations is 49.8–56 Gy_{2,47} and is slightly smaller for the optimisations that included the sequencing (beam-on and sequencing: 50.5–56.4 Gy_{2,47}, beam-on: 49.8–56.3 Gy_{2,47}).

3.3. BEDVHs and dose/BED distributions

In this section, we present three example cases, one exhibiting only minor changes from the optimisations (case 06), one with improvements to both coverage and selectivity (case 08), and one that revealed substantially higher BED levels for the original treatment (case 09). BEDVHs of the original and optimised treatment plans for the

Table 3. Overview of the objective function values and quality parameters (Selectivity, Coverage, PCI) for the original treatments and the different optimisation scenarios. Shown are the average and range for all cases that were optimised to convergence.

Optimisation		Objective Mean	Selectivity			Coverage			PCI		
			Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
Orig	Dose	—	0.883	0.744	0.967	0.966	0.952	0.985	0.853	0.718	0.929
	BED	155.79	0.837	0.587	0.930	0.974	0.945	0.997	0.815	0.581	0.896
Beam-on	Local	61.41	0.927	0.860	0.962	0.940	0.894	0.967	0.872	0.768	0.931
	MILP	61.35	0.927	0.859	0.964	0.940	0.894	0.968	0.872	0.768	0.933
Beam-on and sequence	Local	57.74	0.929	0.867	0.963	0.942	0.898	0.968	0.875	0.779	0.932
	MILP	57.15	0.929	0.865	0.964	0.942	0.900	0.968	0.876	0.779	0.934

BED: biologically effective dose, MILP: mixed-integer linear programming, PCI: Paddick conformity index.

selected cases and the corresponding distributions of the physical dose and BED are shown in figure 3. In addition, figure 4 shows a comparison plot for case 09, visualising the difference in the BED distributions between the initial and locally optimised (beam-on time and sequence) treatment plans.

Several cases happened to be close to the chosen BED reference, mostly due to their original treatment time, and as a result, their BEDVHs exhibit only small changes after the optimisation (see case 06 in figure 3(a)). The minimal changes to the BEDVH are reflected in the similarity of all iso-dose and iso-BED lines for both the original and the optimised treatment plans. Changes in the BED distribution after the optimisations are mostly limited to re-distributions in the high dose/BED region inside of the TV.

Case 08 is an example where both the coverage and selectivity could be improved by the BED optimisation. The BED_{95} of the TV after optimisation is increased by approximately 2.3 $Gy_{2.47}$, while the BED_5 of the Rim is reduced by approximately 1.2 $Gy_{2.47}$. These improvements appear to be facilitated by a significant increase of BED levels inside the TV where the peak BED is now increased from originally 181 to 216–227 $Gy_{2.47}$ for the optimised treatment plans. The improvements in conformity observed in the BEDVH are also visible in the BED distributions that show improved adherence to the TV shape, especially on the left side (patient right) of the target. This improved conformity of the optimised treatment plans also coincides with an increase in the volume covered by the 150 $Gy_{2.47}$ iso-BED line.

As mentioned before, most cases exhibit marginal BED levels beyond the chosen prescribed BED. An example is case 09 with an originally short treatment time of 28.1 min. Figure 3(c) shows that the volume covered by the prescribed BED is substantially larger than that covered by the prescribed physical dose. After optimisation, the BED is generally scaled down to the reference level (BED_{95} of 53.7 $Gy_{2.47}$). In addition, the split in the target volume in the chosen slice now also exhibits a split in the prescribed BED iso-lines.

The difference plot in figure 4 shows a large area around the split in the TV where the BED is reduced by more than 40 $Gy_{2.47}$, with a peak reduction of 71.8 $Gy_{2.47}$. In the initial treatment plan, large parts of the healthy tissue in the gap of the TV received more than 100 $Gy_{2.47}$ (peak: 140.8 $Gy_{2.47}$).

In contrast to the initial physical dose treatment plans, the proposed optimisation approaches do not enforce a specific heterogeneity in the distributions. The mean prescription iso-dose level across the cohort originally was 48.8% ($\sigma = 5.23\%$) which translates to 25.9% ($\sigma = 4.2\%$) in terms of the reference BED. After optimising the beam-on time and sequence with the MILP approach, these values change to 50.9% ($\sigma = 8.3\%$) and 28.28% ($\sigma = 7.4\%$), respectively.

4. Discussion

4.1. Optimiser performance

The experiments executed in this study suggest that it is feasible to optimise the per iso-centre beam-on times and the sequence of delivery to create the most beneficial BED distribution. The local approaches do appear to converge to a local minimum, as evident by the fact that the MILP optimisation approach usually converges to a lower objective function value. However, this local minimum appears to be close to the global optimum and no meaningful degradation in plan quality was found as a result of the difference in objective function value. In addition, convergence is reached relatively quickly in the order of seconds/minutes (single-threaded load) compared to hours/days with the MILP approach (multi-threaded).

Using the convex underestimator approach, it is possible to simultaneously consider the beam-on time and order of iso-centre delivery to improve on the local approach. However, the complexity of the MILP models increases dramatically with the number of iso-centres, to the degree that for four of the cases in the cohort (03,

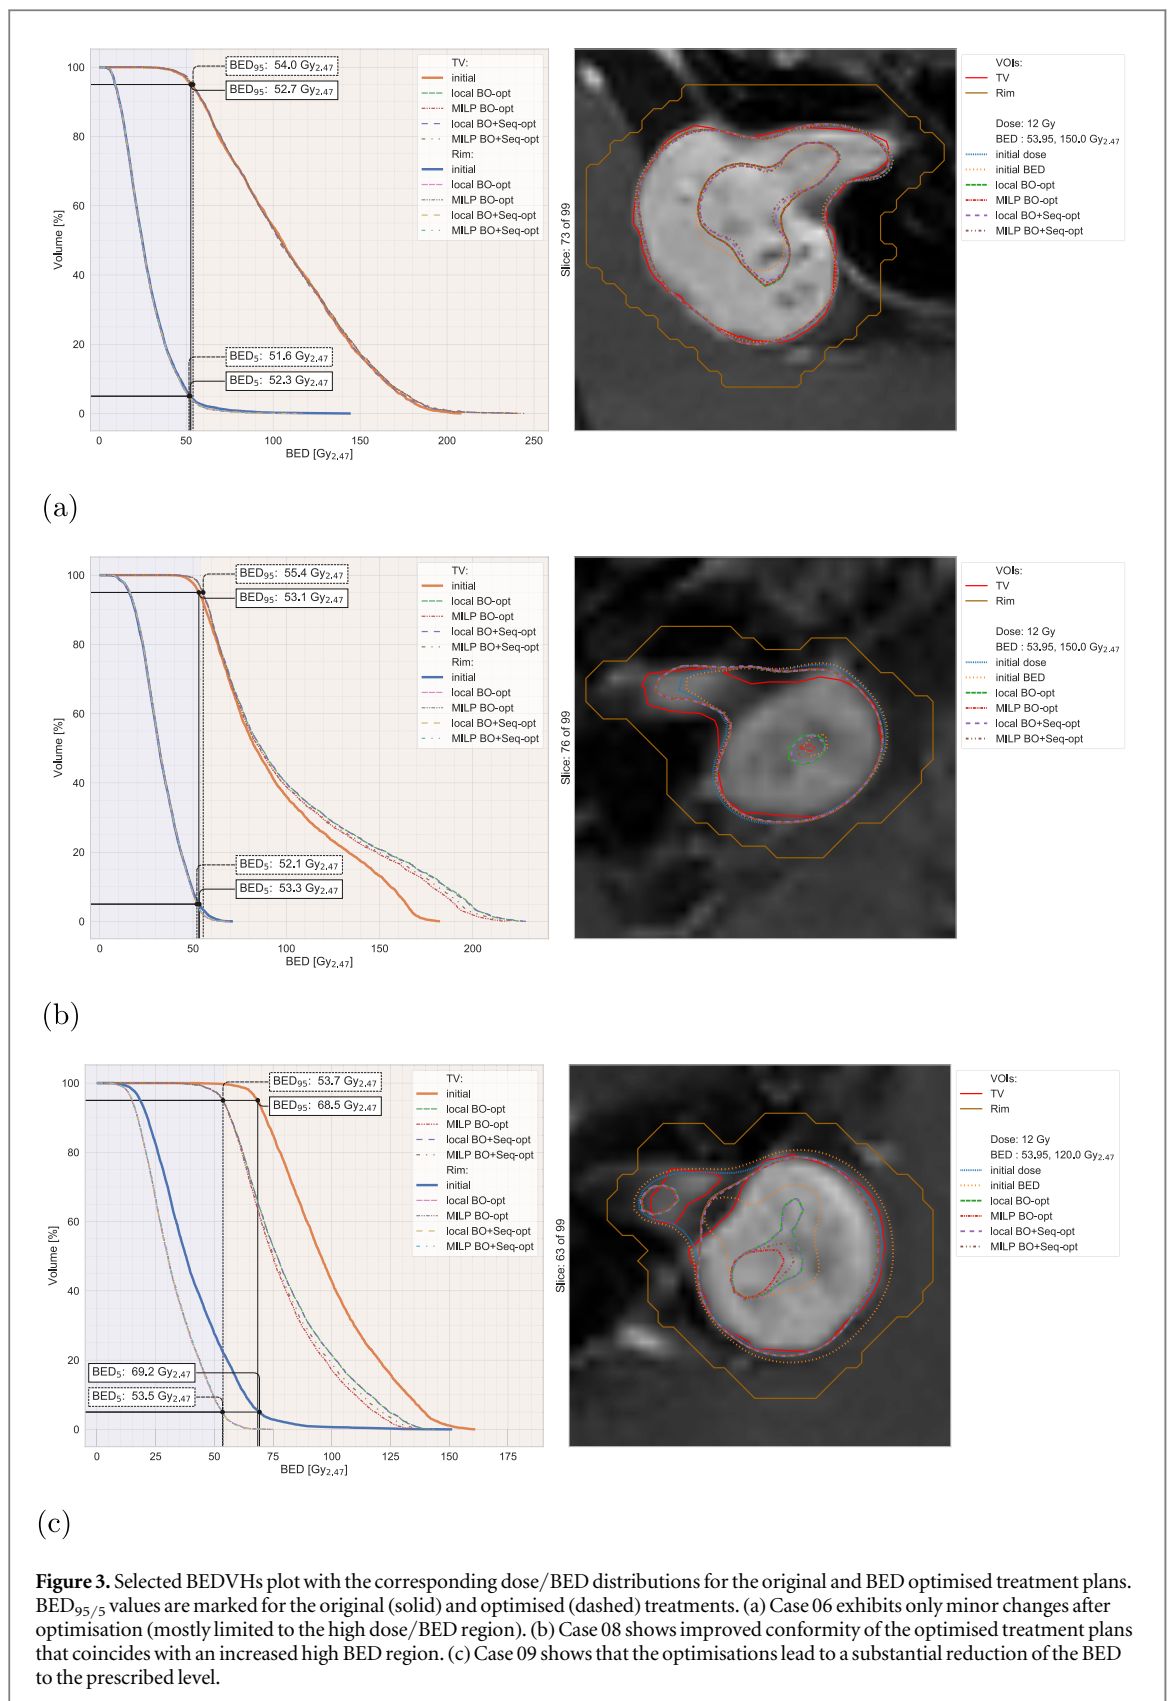



Figure 3. Selected BEDVHs plot with the corresponding dose/BED distributions for the original and BED optimised treatment plans. BED_{95/5} values are marked for the original (solid) and optimised (dashed) treatments. (a) Case 06 exhibits only minor changes after optimisation (mostly limited to the high dose/BED region). (b) Case 08 shows improved conformity of the optimised treatment plans that coincides with an increased high BED region. (c) Case 09 shows that the optimisations lead to a substantial reduction of the BED to the prescribed level.

04, 08, 12) with $N_{iso} \geq 16$ the MILP problem could not be solved to integer-optimality within the defined limit on the computation time. Consequently, the next update of the beam-on time bounds is based on the current best solution which can be sub-optimal and lead to the exclusion of the global optimum from the possible solutions. Thus, the final result after a lengthy optimisation (up to 14 d on a compute cluster) can be worse than the local beam-on time optimisation (see case 08). However, given the results for the cases with fewer iso-centres, this is only an issue of computation time and resources and not a flaw of the approach itself. In addition,



this approach provides a lower bound for the global optimum within the provided limits on the beam-on time whereas it is impossible to estimate the distance to the global optimum with the local approaches.

4.2. Plan quality

Optimising both the beam-on time and sequence can lead to considerably better objective function values, however, this is dependent on the individual cases. In addition, improvements in the objective function value do not necessarily directly translate to increased plan quality.

As mentioned before, most cases exhibit marginal BED levels beyond the chosen prescribed BED. As a result, optimisation typically scales down the BED of the treatment to the appropriate level and increase the selectivity to a similar level of coverage.

The objective function was chosen to be easily compatible across all approaches and create reasonable treatment plans across the entire cohort. It will generally promote both coverage and selectivity to the same degree which results in improved PCI values for the optimised treatment plans. While this relatively simple objective function would be completely minimised (no penalties) for perfect coverage and selectivity, a better objective function value does not directly translate to a better clinical score. In a clinical situation, one would incorporate the planning goals (e.g. minimum 95% coverage, VOI maximum BED thresholds) and quality parameters in the objective function and tailor the weights of the different objectives to the individual patient to achieve the most beneficial treatment plans. The improvements in the overall PCI with this generic objective function indicate that it would be possible to not only reach the same treatment plan quality as for the physical dose plan, but improve on it. It would appear that the added complexity of using a BED model with incomplete repair intervals also provides an additional degree of freedom (in the time domain) that could allow for the creation of even more conformal treatment plans than currently possible.

An example for this is case 08 where both the selectivity and coverage are above the original levels of the physical dose plan. This improvement comes with an increase of the high BED region at the centre of the target. When introducing BED-based treatment plans, one will have to investigate the appropriate prescription BED values and the most beneficial range of BED values inside the TV.

4.3. BED prescription

In this study, the prescription BED was chosen according to Jones and Hopewell (2018) with a reference treatment time of 60 min. For the present cohort, this happens to represent a relatively low BED, requiring the optimisation to generally scale down the overall BED level. This is of course only one proposed way to take the treatment time into account and define a reference BED. If a shorter treatment time was chosen as a basis, the present cohort could have appeared to generally exhibit a lower BED than desired. Nevertheless, the results in this study demonstrate that the optimisation approaches can both scale up or down the overall BED level and optimise the conformity of the given cases. Even without enforcing a certain level of heterogeneity to the BED distribution, the optimised treatment plans exhibit similar prescription iso-dose and iso-BED levels as the clinical treatments. This suggests that the inherent heterogeneity in the delivered GK dose is generally maintained when optimising the BED in the presented scenarios.

Additional studies are needed to determine an appropriate prescription BED that signifies a beneficial trade-off between the therapeutic effect and the incidence of adverse effects. An example would be the publication by

Tuleasca *et al* (2019) that investigated the treatment of trigeminal neuralgia, which found that increasing BED values above a certain threshold increases the risk of complications without increasing the probability of pain control any further. Accurately modelling the tumour control probability (TCP) and normal tissue complication probability (NTCP) can be a valuable tool for determining the therapeutic window of a given treatment regimen.

4.4. Contouring the VOIs

When using inverse planning, the contours that are created determine whether a voxel is considered to be part of the TV, an OAR or the NT. Thus, special care has to be taken to ensure an accurate representation of the treatment planning problem. As can be seen for case 09, where there was a gap in the TV on a single slice (slice thickness: 1.5 mm), the optimiser does not distinguish between voxels of the same class based on their location/surrounding tissue and will adapt the treatment plan accordingly to minimise the objective function. In this case, a lower objective function value is achieved by additional sparing of the NT in this small gap, a characteristic which was not observed in the original treatment plan. This was likely done to ensure coverage of the entire TV. With the inverse planning approach, these high BED values in the NT (regardless of their location) are penalised which leads to the observed BED reductions in that area. This highlights the importance of the drawn contours being consistent with the intention of the treatment planner. If the intention is to accept higher BED values in this gap, an additional contour with a higher BED threshold or lower penalty weight could be added in this location.

In a clinical scenario, where appropriate target coverage is required, the treatment planner would then adapt either the objectives or the contours to ensure the clinical goals are met. In addition, the contouring step will have to include any VOIs in the vicinity of the TV that need to be considered during the treatment plan optimisation. Avoidance of non-delineated critical structures, which a human planner might be aware of during manual planning, is not possible with an inverse planning approach.

4.5. Expansion of the treatment planning framework

Having established the local optimisation methods to be a suitable choice for BED treatment planning, the next step would be to include further variables into this optimisation approach to explore the possibility of not only matching the quality of the original dose plan but to improve upon it. One candidate is the iso-centre location. To use the same L-BFGS-B optimisation framework to simultaneously optimise both the beam-on times and the iso-centre locations, some approximations have to be made. Firstly, the shape of the distribution is assumed to remain unchanged for small displacements of a few millimetres. Secondly, to allow for continuous values of the displacement, the dose-rate matrix is linearly interpolated at the candidate locations.

We integrated the iso-centre location optimisation in our proposed framework and figure 5 shows the quality parameters for the original treatment, the local beam-on time and sequencing approach and the location and beam-on time optimisation. Clear improvements, even beyond the values for the original dose treatment plan can be observed. While the mean coverage now matches the original dose value, the mean selectivity is increased from 89.6% to 95.5%. These preliminary results show the potential for meaningful improvements of the treatment plan quality by taking advantage of the added degree of freedom provided by a BED model.


Additionally, the simultaneous optimisation of sequence and beam-on time could be extended to the selection of the most beneficial collimators for the individual sectors. Different combinations of optimisation variables could then be compared to determine the best trade-off between the increased complexity of the optimisation and the expected improvements of the treatment plans.

Having demonstrated the ability for BED-based inverse planning, the next step is to establish a semi-automatic workflow that allows optimising individual treatment plans in a clinical scenario. This would include more practical objectives to individualise treatment plans and a level of interactivity allowing to explore the feasible trade-offs for a given treatment plan. In general, the investigated semi-automatic approach could be fully automated if an initial 'filling' step was used for the iso-centre definition as available in Leksell GammaPlan.¹³

Given a consensus about BED-based prescriptions is found, such an approach could be translated into the clinical workflow. This would require some investment into the software development of the system, e.g. to adapt the treatment planning workflow and tailor the optimisation algorithms to the available hardware. GammaPlan Lightning is the most recent example of a proposed novel inverse planning approach (see Sjölund *et al* 2019) being integrated into the clinical treatment planning system.

The computationally expensive convex underestimator approach could be used as a benchmark to assess other viable optimisation methods. One could create a number of artificial test cases which can be optimised to a very high degree with the MILP approach. Publishing these test cases with the appropriate results could then allow others to benchmark their solutions without having to run the costly optimisation themselves.

¹³ See Leksell GammaPlan Online Reference Manual. Article No. 102300 Rev. 01, Elekta Instrument AB Stockholm, Sweden, 2011.



4.6. The role of BED in GK SRS

While there is evidence that dose-rate effects like those expressed in the BED formulation influence the treatment outcome in terms of the therapeutic effectiveness and incidence of adverse effect, it is not the only measure influencing treatment success and not all studies find a good correlation between the BED and the investigated treatment endpoints.

For trigeminal neuralgia treatments, a study by Tuleasca *et al* (2019) determined a beneficial BED range to minimise the risk of patients developing hypoesthesia without compromising the long term pain-free incidence. Since these were single iso-centre treatments, the BED could be calculated directly from the physical dose without the need to consider interactions across iso-centres which would require access to per iso-centre dose maps.

For multi iso-centre treatments, several BED approximation methods were developed by Jones and Hopewell (2018) that allow for the determination of the marginal BED value from the prescribed dose and treatment time. These simplified BED models do not capture all the information contained in the delivered 3D distributions but allow for a convenient way to retrospectively estimate marginal BED values. The simplified BED was successfully used to determine treatment outcome correlations in acromegaly (Graffeo *et al* 2020) and pituitary adenoma Graffeo *et al* (2021) SRS treatments. The importance of adhering to the model constraints has been highlighted in a letter to the editor by Hopewell *et al* (2021).

There have also been discussions about the role of the reference dose-rate in the context of outcome correlation. One important distinction to make is that the BED used in the presented study is determined based on the 3D in-patient dose-rate distribution while the reference dose-rate value is based on the calibration measurement of the specific GK unit at the centre of a phantom. Factors like collimator size, sector blocking and patient geometry make it infeasible to infer one from the other. Illustrative examples of this are provided in Paddick *et al* (2019), where similar treatment times and BED values are observed for varying reference dose-rates (1–4 Gy min⁻¹).

Another factor to consider is the α/β -ratio. For this study, the value 2.47 Gy, as determined by Pop *et al* (2000) together with the repair-rates and partition coefficient, is used. Analogous to previous studies Hopewell *et al* (2012, 2013), Millar *et al* (2015), Klinge *et al* (2021), this fixed value of the α/β -ratio is used for the BED calculation. Jones *et al* (2020) report minor effects on equivalent doses for radiation myelopathy in the central nervous system when varying the α/β -ratio (1.5–3.0 Gy). Using the commonly accepted value of 2 Gy resulted in a <1% deviation of the equivalent single doses for treatment times of up to 5 h. If one were able to extract additional and precise patient-specific information about all involved tissues before the treatment, the likely tissue response could be better quantified. This information could include α/β -ratios for all involved tissues and additional or various repair-rates. The presented optimisation approaches allow for the assignment of per-voxel tissue parameters and any number of desired repair-rates. The treatment plan could then be tailored to the specific targeted tissues on a per-patient basis. In practice however, this information is not currently available for treatment planning.

There is certainly a need to further investigate all factors that can influence treatment outcomes to provide treatment planners with the tools required to ensure the best possible treatment. One of these tools is the BED. Previous studies have used the BED model to retrospectively investigate variations within the same prescription dose level (Hopewell *et al* 2012, 2013, Millar *et al* 2015, Klinge *et al* 2021) and investigate outcome correlations

(discussed above). In contrast, the presented work investigates how these findings could be used to further improve SRS treatments in the future. This requires the development of new optimisation methods to enable the use of the more complex BED model developed by Millar and colleagues for inverse planning. There currently exist no BED treatment planning solutions. This study is an important step towards possible future treatment planning systems.

5. Conclusion

This study demonstrates multiple strategies that could be employed to make BED planning feasible in clinical settings. BED treatment planning could significantly reduce the inter-patient variability and has the potential to improve the outcome of GK radiosurgery treatments.

For our cohort, the local optimisation approaches are sufficient to reach the same level of treatment plan quality as the more complex MILP approaches. Furthermore, it is feasible to optimise both the beam-on times and the delivery sequence together to improve on treatment plans created with only the beam-on times as a variable. The MILP approach, while less prone to local minima issues, is too computationally expensive to be used in its current form in a clinical workflow without further significant improvements in optimisation speed. Nevertheless, it can be helpful to provide a lower bound on the achievable objective function value and could be used to benchmark new optimisation approaches.

With the described optimisation methods it is possible to (at least) reach the same treatment plan quality in terms of BED as for the original physical dose treatment plans.

Acknowledgements

This work is supported by the EPSRC-funded UCL Centre for Doctoral Training in Medical Imaging [EP/L016478/1], the Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS) [NS/A000050/1], the Wellcome/EPSRC Centre for Medical Engineering [WT 203 148/Z/16/Z], the NIHR BRC based at Guys and St Thomas Trust and the CRUK ARTNET Network Accelerator Award [A21993]. Ethical statement

This research study was conducted retrospectively using non-identifiable human subject data. Applicable law and standards of ethics have been respected in accordance with the principles embodied in the Declaration of Helsinki and in accordance with local statutory requirements. Conflict of interest

Jamie R McClelland reports grants from Elekta Instrument AB, but outside the scope of the submitted work. Ian Paddick reports personal fees from Elekta Instrument AB, but outside the scope of the submitted work. Other authors have nothing to disclose. The views expressed represent those of the authors and not those of Elekta Instrument AB.

Appendix A. BED model

Millar and Canney (1993) derived a model to determine the BED of a fractionated, protracted treatment with a variable dose-rate and multiple damage and repair processes. In their work, they extend upon the principle assumptions of the linear-quadratic (LQ) model:

1. The fractional cell death rate is proportional to the rate of induction of lethal damage (*extended to include the effects due to a polynomial representation in the dose-rate*)
2. There are two components of cellular damage: direct lethal damage and sublethal damage
3. The rate of induction of sublethal damage is proportional to the dose-rate (*extended to include the effects due to a polynomial representation in the dose-rate*)
4. Sublethal damage is potentially repairable and is governed by a single exponential repair rate process (*modified to include effects due to multiple types of damages, each of which may be repaired by multiple repair mechanisms*)
5. Certain types of sublethal damage can be converted to lethal damage by further irradiation. Only convertible sublethal damage will be considered.
6. Only a small proportion of induced sublethal damage is converted to lethal damage
7. The rate of production of lethal damage derived from sublethal damage is proportional to both the dose-rate and the amount of sublethal damage (*extended to include the effects due to a polynomial representation in the dose-rate*)

Table 4. BED model parameters for α/β -ratio, repair-rates $\mu_{1,2}$ and partition coefficient c used in this study. The partition coefficient determines the relative contributions from the two repair-rates. A value of 1 represents equal contributions from slow and fast repair. These values were determined by Pop *et al* (2000).

Parameter:	α/β [Gy]	μ_1 [$\frac{1}{\text{min}}$]	μ_2 [$\frac{1}{\text{min}}$]	c
Value:	2.47	0.0608	0.0053	0.98

These assumptions allow them to describe the time-dependent fractional rate of induced cell death for an arbitrary fractionated protracted treatment. Integration over the duration of the treatment then allows the determination of the biological effect of the entire treatment. The addition of further assumptions that more closely describe a clinical treatment with the GK (or any radiosurgical treatment) allows the simplification of the model. These assumptions are a constant dose-rate during each fraction (i.e. iso-centre), a varying dose-rate level between individual fractions, a single repair rate process, and first-order polynomials for the induction of lethal damage, sublethal damage and translation from sublethal to lethal damage.

This simplification of the model can be formulated as the general LQ model with an additional correction function describing the dose delivery. Pop *et al* (2000) derived a bi-exponential version of this BED model from animal studies and this BED model has since been well established and used by numerous studies over the past decades. As such we refer the reader to the previous work cited in the manuscript for a thorough analysis of the underlying principles.

Appendix B. BED optimisation settings

Table 5. Settings for the different BED optimisation techniques employed in this study.

Scenario	Solver	Settings	Workflow
Original	—	Beam-on times from dose calculation	Calculate BED with original order of iso-centres as sequence
Local	L-BFGS-B	Approx grad: True, bounds: (0, None), m : 50	Provide BED objective function and initial beam-on time vector
BO-opt		Epsilon: 1 s, maxiter: 250	Optimise to convergence
Convex MILP	CPLEX MIP	Threads: 48, mip tolerance mipgap: 0.002,	Build underestimator for beam-on of 1 min +/-
			10 min - > optimise - >
BO-opt		mip_heuristicfreq: -1, mip_limits_cutpasses: -1	
		Convergence criteria: min bounds: 0.01 min,	Update bounds to dt_{opt} +/- 60% of last interval - > optimise
		rel. diff 'relaxed' versus full: $2x < 1e-3$,	until convergence
		individual objective function rel. diff: $4x < 1e-3$	
Local	L-BFGS-B	Approx grad: True, bounds: (0, None), m : 50,	First execute beam-on optimisation, then optimise sequence
BO+Seq-opt		Epsilon: 1 s, maxiter: 250	
	2-opt	Convergence: obj func rel. diff: $3x < 1e-3$	Alternate to convergence
Convex MILP	CPLEX MIP	Threads: 48, mip tolerance mipgap: 0.002,	Start with beam-on MILP
BO+Seq-opt		mip_heuristicfreq: -1, mip_limits_cutpasses: -1,	Until beam-on time bounds < +/- 1 min
		det time limit: 20e6, mip_limits_repairtries: 300	Switch to model with sequence optimisation
		Convergence criteria: min bounds: 0.01 min,	Update bounds to dt_{opt} +/- 60% of last interval
		rel. diff 'relaxed' versus full: $2x < 1e-3$,	Optimise until convergence
		individual objective function rel. diff: $7x < 1e-3$	

BED: biologically effective dose, L-BFGS-B: limited-memory Broyden–Fletcher–Goldfarb–Shanno bound-constrained algorithm, MILP: mixed-integer linear programming.

Appendix C. Creation of the MILP problem

MINLP problems are generally challenging to solve and can require immense computing resources (Burer and Letchford 2012, Köppe 2012). A common solution approach is to find a reformulation of the problem that has beneficial properties like convexity (Burer and Letchford 2012). This study introduces a method that creates a convex hull of the problem to solve it iteratively.

This section covers the steps required to create this convex relaxation (MILP) of the general MINLP treatment planning problem. Firstly, short general descriptions of the necessary methods are presented, and secondly, their application to the BED optimisation problem is described.

Notation Throughout the derivation of the MILP formulation, the following conventions are used to describe different objects.

The application of a convex envelope is described by the introduction of a \hat{w} . So \hat{w} refers to the convex envelope of w .

Vector and matrix variables are expressed with bold lettering and their individual scalar components are referenced with subscripts (e.g. all beam-on times δt with the j th beam-on time δt_j). For the BED, the subscript index identifying the current voxel is omitted when not specifically required. This is done to aid readability where the BED is completely analogous across all voxels, i.e. the only difference between the voxels v is their dose-rate parameter \hat{d}_{jv} for iso-centre j .

C.1. Convex envelopes

The convex relaxation of the optimisation problem is realised using convex envelopes for multilinear functions $w(x)$

$$w(x) = x_1 \cdot \dots \cdot x_k, \quad \forall j \leq k: x_j \in [x_j^L, x_j^U], \quad k \in \mathbb{N} \quad (12)$$

with known lower x_j^L and upper x_j^U bounds of all variables. The multilinear function $w(x)$ is replaced with a new variable \hat{w} and a set of constraints defining the convex relaxation.

C.1.1. McCormic envelopes. If the function is bilinear ($k = 2$), two linear constraints define a lower convex and upper concave envelope. These explicit constraints are known as the McCormick inequalities (McCormick 1976):

$$\text{convex:} \quad \hat{w} \geq x_1^L x_2 + x_2^L x_1 - x_1^L x_2^L \quad (13)$$

$$\hat{w} \geq x_1^U x_2 + x_2^U x_1 - x_1^U x_2^U \quad (14)$$

$$\text{concave:} \quad \hat{w} \leq x_1^L x_2 + x_2^U x_1 - x_1^L x_2^U \quad (15)$$

$$\hat{w} \leq x_1^U x_2 + x_2^L x_1 - x_1^U x_2^L. \quad (16)$$

These inequalities restrict the values of (x_1, x_2, \hat{w}) to be within the polyhedron defined by the vertex set of the possible combinations of lower and upper bounds of the variables $P_W = \{(x_1^L, x_2^L, x_1^L x_2^L), (x_1^U, x_2^U, x_1^U x_2^U), (x_1^U, x_2^L, x_1^U x_2^L), (x_1^L, x_2^U, x_1^L x_2^U)\}$.

C.1.2. Dual envelopes. In general, convex envelopes of multilinear terms are vertex polyhedral (Rikun 1997, Tardella 2008). To take advantage of this property, Costa and Liberti (2012) propose to express the points inside the envelope as the convex combination of its extreme points $P_W = \{p_1, \dots, p_{2^k}\} \subseteq \mathbb{R}^{k+1}$. Their approach, named dual envelopes, allows them to define the convex hull of a general multilinear term for an arbitrary value of $k \in \mathbb{N}$.

They introduce a vector $\lambda \in \mathbb{R}_{\geq 0}^{2^k}$ of nonnegative Lagrange multipliers that defines a point such that:

$$x = \sum_{i \geq 2^k} \lambda_i p_i \quad \wedge \quad \sum_{i \geq 2^k} \lambda_i = 1. \quad (17)$$

To create the envelopes, the p_i 's need to be expressed as a function of the bounds on the variables x^L, x^U . Costa and Liberti define two parameter sequences, d_{ij} and $b_j(d_{ij})$. The value of $d_{ij} \in \{0, 1\}$ describes whether the j th component of p_i is a lower or upper bound, while $b_j(d_{ij})$ returns the value of the specific bound:

$$\forall i \leq 2^k, j \leq k \quad d_{ij} = \left(\left\lfloor \frac{i-1}{2^{k-j}} \right\rfloor \bmod 2 \right) \quad (18)$$

$$\forall j \leq k \quad b_j(0) = x_j^L \quad \wedge \quad b_j(1) = x_j^U. \quad (19)$$

The k -linear term $w(x) = x_1 \cdot \dots \cdot x_k$ can then be relaxed with the introduction of 2^k constrained variables λ_i and $k+1$ new constraints:

$$\hat{w} = \sum_{i \leq 2^k} \lambda_i \prod_{j \leq k} b_j(d_{ij}) \quad (20)$$

$$\forall j \leq k \quad x_j = \sum_{i \leq 2^k} \lambda_i b_j(d_{ij}) \quad (21)$$

$$\sum_{i \leq 2^k} \lambda_i = 1 \quad \wedge \quad \forall i \leq 2^k, \lambda_i \geq 0. \quad (22)$$

This relaxation is executed individually for every term in the nested sum of the BED formulation (see equation (4)). In this way, the convex hull of a problem containing k -linear terms can be constructed, allowing the relaxed problem to be optimised with the Lagrange multipliers λ_i as the decision variables.

C.2. Big M Method

When formulating the problem, there are situations where the product of two variables needs to be determined. An example is the use of the interaction terms in equation (7), where a binary variable is introduced to determine the active terms. To solve the model with the proposed methods, the product of the two involved variables needs to be linearised (Glover 1975). For the multiplication of a bounded continuous variable $M \in [0, M^U]$ and binary variable $x \in \{0, 1\}$, a new variable $z = M \cdot x$ and a set of constraints can be introduced:

$$z \leq M^U \cdot x \quad (23)$$

$$z \leq M \quad (24)$$

$$z \geq M - (1 - x) \cdot M^U \quad (25)$$

$$z \geq 0. \quad (26)$$

These linear inequalities ensure that z is within the expected bounds $[0, M^U]$, equal to zero if $x = 0$, and equal to M if $x = 1$. This approach is commonly referred to as the big M method. Similarly, if the continuous variable is bounded to a non-positive interval $[M^L, 0]$, the variable $z = M \cdot x$ will require the following constraints:

$$z \geq M^L \cdot x \quad (27)$$

$$z \geq M \quad (28)$$

$$z \leq M - (1 - x) \cdot M^L \quad (29)$$

$$z \leq 0. \quad (30)$$

For two binary variables $x, y \in \{0, 1\}$, we can define $z = x \cdot y$ together with the following three inequalities:

$$z \leq x \quad (31)$$

$$z \leq y \quad (32)$$

$$z \geq x + y - 1. \quad (33)$$

These linearisations will be used to facilitate the sequence optimisation that requires the introduction of additional binary variables (see appendix C.4.6).

C.3. Piece-wise linearisation

To constrain the optimisation to physically deliverable treatment plans, the timing information, i.e. beam-on times and starting times of the individual iso-centres, has to be recovered from the dual envelope formulation without re-introducing nonlinearity to the optimisation. This is achieved by applying a piecewise linearisation (PWL) (Lin *et al* 2013) to each individual substituted nonlinear term.

For a single variable function $f(x): \mathbb{R} \rightarrow \mathbb{R}$ with $x \in [a_0, a_m]$, a set of $m + 1$ monotonically increasing support points $\{a_l \mid l \in \{0, \dots, m\}, a_0 < a_1 < \dots < a_m\}$ can be defined. An approximate PWL of $L(f(x))$ over the interval $[a_0, a_m]$ can then be achieved using the two surrounding support points $a_l \leq x \leq a_{l+1}$ and their individual weights ν_l :

$$L(f(x)) = \sum_{l=0}^m a_l \nu_l \quad (34)$$

$$x = \sum_{l=0}^m a_l \nu_l. \quad (35)$$

This requires the use of the following constraints:

$$\nu_0 \leq \gamma_0, \quad \nu_m \leq \gamma_{m-1} \quad (36)$$

$$\nu_l \leq \gamma_{l-1} + \gamma_l, \quad \text{for } l = 1, \dots, m - 1 \quad (37)$$

$$\sum_{l=0}^{m-1} \gamma_l = 1 \quad (38)$$

$$\sum_{l=0}^m \nu_l = 1 \quad (39)$$

$$\gamma_l \in \{0, 1\}, \quad \nu_l \geq 0, \quad \text{for } l = 0, \dots, m - 1. \quad (40)$$

The binary variable γ_l determines the active interval for any value of $x \in [a_0, a_m]$ and thus which two adjacent ν_l take a nonzero value.

The number of new (binary) variables and linear constraints introduced with this approach scales linearly with the number of support points m . A higher number of support points increases both the accuracy of the PWL and the computational cost to solve the problem.

C.4. Convex relaxation of the MINLP problem

C.4.1. Reformulating the BED. To implement the optimisation of the BED, we need to express it in terms of the actual decision variables and parameters used by the solver. Specifically, we define a new variable Δt that describes the difference in starting time of two iso-centres and relate this back to the individual combination of beam-on times δt and gaps g between iso-centre deliveries. In addition, we substitute the nonlinear terms and combine constants that are the same across multiple summands into a single parameter. *Original:*

$$\text{BED} = \sum_{j=1}^N \dot{d}_j \delta t_j + \frac{1}{\alpha/\beta} \left[\frac{\Psi(\Xi, \mu_1) + c \cdot \Psi(\Xi, \mu_2)}{1 + c} \right] \quad (41)$$

$$\begin{aligned} \Psi(\Xi, \mu) = & \frac{2}{\mu} \sum_{j=1}^N \left[\dot{d}_j^2 \left[\delta t_j - \frac{1}{\mu} (1 - e^{-\mu \delta t_j}) \right] \right. \\ & \left. - \frac{1}{\mu} \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j e^{-\mu(t_j - t_k)} (e^{\mu \delta t_k} - 1) (e^{-\mu \delta t_j} - 1) \right]. \end{aligned} \quad (42)$$

Reformulated:

$$\begin{aligned} \text{BED} = & \sum_{j=1}^N \dot{d}_j \delta t_j + A_1 \sum_{j=1}^N \dot{d}_j^2 w_{A_1}(\delta t_j) + A_2 \sum_{j=1}^N \dot{d}_j^2 w_{A_2}(\delta t_j) \\ & - B_1 \sum_{j=1}^N \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j w_{B_1}(\Delta t_{jk}, \delta t_k, \delta t_j) - B_2 \sum_{j=1}^N \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j w_{B_2}(\Delta t_{jk}, \delta t_k, \delta t_j) \end{aligned} \quad (43)$$

$$A_i = \frac{1}{1 + c} \cdot \frac{2}{\alpha/\beta \cdot \mu_i}, \quad i \in \{1, 2\} \quad (44)$$

$$B_i = \frac{1}{1 + c} \cdot \frac{2}{\alpha/\beta \cdot \mu_i^2}, \quad i \in \{1, 2\} \quad (45)$$

$$w_{A_i} = \delta t - \frac{1}{\mu_i} (1 - e^{-\mu_i \delta t}), \quad i \in \{1, 2\} \quad (46)$$

$$w_{B_i, jk} = e^{-\mu_i \Delta t_{jk}} (e^{\mu_i \delta t_k} - 1) (e^{-\mu_i \delta t_j} - 1), \quad i \in \{1, 2\} \quad (47)$$

$$\Delta t_{jk} = t_j - t_k = \sum_{l=k}^{j-1} (\delta t_l + g_l). \quad (48)$$

C.4.2. Application of the dual envelopes. The dual envelope approach described in appendix C.1.2 is applied to the nonlinear terms $w_{A_{1,2}}$, $w_{B_{1,2}}$. For simplicity and readability, we will limit the formal description to a single term of w_A and w_B . The implementation for the individual repair-rates μ_1, μ_2 , iso-centres, and possible interaction terms is completely analogous. For the interaction terms, the individual beam-on times are designated as δt_c (the *current* iso-centre) and δt_p (a *previously* delivered iso-centre).

The nonlinear function w_A is replaced with the convex envelope \hat{w}_A constructed from the lower and upper bounds of the beam-on time $\underline{\delta t}$, $\overline{\delta t}$.

$$\hat{w}_A = \sum_{i \leq 2^k} \lambda_i \prod_{j \leq k} b_j(d_{ij}) \quad (49)$$

$$\forall i \leq 2^k, j \leq k \quad d_{ij} = \left(\left\lfloor \frac{i-1}{2^{k-j}} \right\rfloor \bmod 2 \right) \quad (50)$$

$$\forall j \leq k \quad b_j(0) = x_j^L = w_A(\underline{\delta t}) \quad \wedge \quad b_j(1) = x_j^U = w_A(\overline{\delta t}) \quad (51)$$

$$\sum_{i \leq 2^k} \lambda_i = 1 \quad \wedge \quad \forall i \leq 2^k, \lambda_i \geq 0. \quad (52)$$

Practically, since $k = 1$ for our application (single variable, $w_A = \prod_{j \leq k} x_j$), this boils down to a substitution of the nonlinear function w_A where the Lagrange multipliers λ_i act as a scaling factor within the bounded interval $[w_A(\underline{\delta t}), w_A(\overline{\delta t})]$. If, for example, we would add the dose-rate d as a decision variable, then this could easily be formulated as a bilinear function ($w_A = x_1 \cdot x_2$) and the dual envelope would act as described in appendix C.1.2.

Similarly, the interaction term between the iso-centres w_B is a nonlinear function of three variables that require a convex relaxation to create a linear problem. First, we substitute the individual terms in the product to

express w_B as a trilinear function:

$$w_B = e^{-\mu\Delta t_{cp}}(e^{\mu\delta t_p} - 1)(e^{-\mu\delta t_c} - 1) = x_1 \cdot x_2 \cdot x_3 \quad (53)$$

$$x_1 = e^{-\mu\Delta t_{cp}}, \quad x_2 = e^{\mu\delta t_p} - 1, \quad x_3 = e^{-\mu\delta t_c} - 1. \quad (54)$$

The bounds on the individual x_j are determined from the lower and upper bounds of the individual variables $\underline{\Delta t}$, $\overline{\Delta t}$, $\underline{\delta t}$, $\overline{\delta t}$.

$$x_1^L = e^{-\mu\overline{\Delta t_{cp}}}, \quad x_2^L = e^{\mu\underline{\delta t_p}} - 1, \quad x_3^L = e^{-\mu\overline{\delta t_c}} - 1 \quad (55)$$

$$x_1^U = e^{-\mu\underline{\Delta t_{cp}}}, \quad x_2^U = e^{\mu\overline{\delta t_p}} - 1, \quad x_3^U = e^{-\mu\underline{\delta t_c}} - 1. \quad (56)$$

With this information, the dual envelop \hat{w}_B of the trilinear interaction term ($k = 3$) is defined as follows:

$$\hat{w}_B = \sum_{i \leq 2^k} \lambda_i \prod_{j \leq k} b_j(d_{ij}) \quad (57)$$

$$\forall j \leq k \quad x_j = \sum_{i \leq 2^k} \lambda_i b_j(d_{ij}) \quad (58)$$

$$\forall i \leq 2^k, j \leq k \quad d_{ij} = \left(\left\lfloor \frac{i-1}{2^{k-j}} \right\rfloor \bmod 2 \right) \quad (59)$$

$$\forall j \leq k \quad b_j(0) = x_j^L \quad \wedge \quad b_j(1) = x_j^U \quad (60)$$

$$\sum_{i \leq 2^k} \lambda_i = 1 \quad \wedge \quad \forall i \leq 2^k, \lambda_i \geq 0. \quad (61)$$

Now the convex relaxation of the BED can be expressed with the beam-on times δt and Lagrange multipliers λ as variables:

$$\begin{aligned} \hat{B}ED = & \sum_{j=1}^N \dot{d}_j \delta t_j + A_1 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_{1j}} + A_2 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_{2j}} \\ & - B_1 \sum_{j=1}^N \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j \hat{w}_{B_{1jk}} - B_2 \sum_{j=1}^N \sum_{k=1}^{j-1} \dot{d}_k \dot{d}_j \hat{w}_{B_{2jk}}. \end{aligned} \quad (62)$$

C.4.3. Application of PWLs. Up until now, all terms in the convex BED formulation are completely independent of each other. However, in practice, each individual iso-centre will contribute to multiple of these terms. For example, iso-centre j will interact with all previously delivered iso-centres $k < j$ in one way and all following iso-centres $k > j$ in another way (see x_1, x_2, x_3 in equation (54)). Thus, we must enforce consistency for all the timing information across the individual envelopes. Since the individual terms substituted in the dual envelope approach are bijective over the domain $\mathbb{R}_{\geq 0}$ of the timing variables (δt and Δt) and their corresponding codomains, they can be inverted. Thus, one can retrieve the timing information from the state of the variables of the dual envelopes. To not re-introduce nonlinearities into our system, we apply PWLs to these nonlinear terms to determine the values of the timing variables δt and Δt .

By enforcing an equality constraint between the PWL of the nonlinear terms L and the corresponding values from the dual envelope (x_j, \hat{w}), the timing variables can be extracted from the current state of the PWL.

For \hat{w}_A the relationship is as follows:

$$L\left(\delta t - \frac{1}{\mu}(1 - e^{-\mu\delta t})\right) = \sum_{l=0}^m \left(a_l - \frac{1}{\mu}(1 - e^{-\mu a_l})\right) \nu_l = \sum_{i \leq 2^k} \lambda_i b_1(d_{i1}) \quad (63)$$

$$\Rightarrow \delta t = \sum_{l=0}^m a_l \nu_l. \quad (64)$$

For \hat{w}_B we need to create a PWL for each of the three terms (x_1, x_2, x_3) and enforce the same equality constraints:

$$L(x_j) = \sum_{l=0}^m x_j(a_l) \nu_l = \sum_{i \leq 2^k} \lambda_i b_j(d_{ij}), \quad j \in \{1, 2, 3\}, k = 3 \quad (65)$$

$$j = 1: L(e^{-\mu\Delta t_{cp}}) = \sum_{l=0}^m (e^{-\mu a_l}) \nu_l = \sum_{i \leq 2^k} \lambda_i b_1(d_{i1}) \Rightarrow \Delta t_{cp} = \sum_{l=0}^m a_l \nu_l \quad (66)$$

$$j = 2: L(e^{\mu\delta t_p} - 1) = \sum_{l=0}^m (e^{\mu a_l} - 1) \nu_l = \sum_{i \leq 2^k} \lambda_i b_2(d_{i2}) \Rightarrow \delta t_p = \sum_{l=0}^m a_l \nu_l \quad (67)$$

$$j = 3: L(e^{-\mu\delta t_c} - 1) = \sum_{l=0}^m (e^{-\mu a_l} - 1) \nu_l = \sum_{i \leq 2^k} \lambda_i b_3(d_{i3}) \Rightarrow \delta t_c = \sum_{l=0}^m a_l \nu_l. \quad (68)$$

As mentioned above, these equations are exemplary for individual envelopes. The introduced constraints need to be enforced across all terms in the BED formulation. For each of the two repair-rates there are N_{iso} \hat{w}_A 's and $N_{ia} = \frac{N_{iso}(N_{iso}-1)}{2}$ \hat{w}_B 's. With the constraints on every individual term, we can ensure that the timing of the iso-centres is consistent across all envelopes even after switching to the Lagrange multipliers as the variables.

C.4.4. Objective function. To implement the objective function (see equation (6)), a new cost variable is introduced:

$$f(\hat{\mathbf{B}}\hat{\mathbf{E}}\mathbf{D}) = w_{TV} \sum_{v \in TV} \frac{cost_v}{N_{TV}} + w_{Rim} \sum_{v \in Rim} \frac{cost_v}{N_{Rim}}. \quad (69)$$

Using inequality constraints for every voxel v , it can be ensured that only BED values beyond the defined lower BED_{ref} and upper bounds BED_{thres} contribute to the penalty value (i.e. no negative penalty is incurred for voxels within their defined bounds)

$$cost_v \geq \begin{cases} BED_{ref} - \hat{B}ED_v & \forall v \in TV \\ \hat{B}ED_v - BED_{thres} & \forall v \in Rim \end{cases}. \quad (70)$$

C.4.5. Beam-on time MILP problem. With the approaches introduced above, we can now define the beam-on time optimisation as a MILP problem. The parameters controlling the number of support points used for the PWLs (m_A and $m_{B,1}, m_{B,2}, m_{B,3}$) allow for a trade-off between the accuracy of the approximations and the incurred computational cost. An appropriate choice of this parameter becomes crucial when dealing with the sequence optimisation (see appendix C.5.1). For the beam-on time MILP problem, two support points were used for every envelope.

$$\begin{aligned} & \text{minimise } f(\hat{\mathbf{B}}\hat{\mathbf{E}}\mathbf{D}) \\ & \text{w.r.t. } \delta \mathbf{t} \in \mathbb{R}_{\geq 0}^{N_{iso}} \quad \Delta \mathbf{t} \in \mathbb{R}_{\geq 0}^{N_{ia}} \\ & \quad \hat{\mathbf{w}}^A \in \mathbb{R}_{\geq 0}^{2 \times N_{iso}} \quad \lambda^A \in \mathbb{R}_{\geq 0}^{2 \times N_{iso} \times 2} \\ & \quad \mathbf{L}(\mathbf{w}^A) \in \mathbb{R}_{\geq 0}^{2 \times N_{iso}} \quad \nu^A \in \mathbb{R}_{\geq 0}^{2 \times N_{iso} \times m_A} \quad \mathbf{y}^A \in \{0, 1\}^{2 \times N_{iso} \times m_A - 1} \\ & \quad \hat{\mathbf{w}}^B \in \mathbb{R}_{\leq 0}^{2 \times N_{ia}} \quad \lambda^B \in \mathbb{R}_{\geq 0}^{2 \times N_{ia} \times 8} \\ & \quad \mathbf{x}_1^B \in \mathbb{R}_{\geq 0}^{2 \times N_{ia}} \quad \mathbf{x}_2^B \in \mathbb{R}_{\geq 0}^{2 \times N_{iso}} \quad \mathbf{x}_3^B \in \mathbb{R}_{\leq 0}^{2 \times N_{iso}} \\ & \quad \mathbf{L}(\mathbf{x}_1^B) \in \mathbb{R}_{\geq 0}^{2 \times N_{ia}} \quad \nu_1^B \in \mathbb{R}_{\geq 0}^{2 \times N_{ia} \times m_{B,1}} \quad \mathbf{y}_1^B \in \{0, 1\}^{2 \times N_{ia} \times m_{B,1} - 1} \\ & \quad \mathbf{L}(\mathbf{x}_2^B) \in \mathbb{R}_{\geq 0}^{2 \times N_{iso}} \quad \nu_2^B \in \mathbb{R}_{\geq 0}^{2 \times N_{iso} \times m_{B,2}} \quad \mathbf{y}_2^B \in \{0, 1\}^{2 \times N_{iso} \times m_{B,2} - 1} \\ & \quad \mathbf{L}(\mathbf{x}_3^B) \in \mathbb{R}_{\leq 0}^{2 \times N_{iso}} \quad \nu_3^B \in \mathbb{R}_{\geq 0}^{2 \times N_{iso} \times m_{B,3}} \quad \mathbf{y}_3^B \in \{0, 1\}^{2 \times N_{iso} \times m_{B,3} - 1} \\ & \quad cost_i \in \mathbb{R}_{\geq 0}^{N_{vox}} \\ & \text{s.t. (49) - (52) (57) - (61) (dual env. constr.)} \\ & \quad (34) - (40) (63) - (68) \quad \text{(PWL constr.)} \\ & \quad (62) \quad (69) - (70) \quad \text{(cost constr.)} \end{aligned} \quad (71)$$

C.4.6. Adding sequence optimisation. To add the ability to change the delivery sequence, we introduce a square matrix variable $\mathbf{h} \in \{0, 1\}^{N \times N}$ controlling the active terms

$$\begin{aligned} \hat{\mathbf{B}}\hat{\mathbf{E}}\mathbf{D} = & \sum_{j=1}^N \dot{d}_j \delta t_j + A_1 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_{1,j}} + A_2 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_{2,j}} \\ & - B_1 \sum_{j \neq k} h_{jk} \dot{d}_k \dot{d}_j \hat{w}_{B_{1,jk}} - B_2 \sum_{j \neq k} h_{jk} \dot{d}_k \dot{d}_j \hat{w}_{B_{2,jk}} \end{aligned} \quad (72)$$

and enforce a constraint that limits the active terms to the physically possible ones. That is: there are always $\frac{N(N-1)}{2}$ active terms and only one active interaction term for any combination of two iso-centres j, k .

$$\forall j, k \in \{1, \dots, N\}, j \neq k \quad h_{jk} + h_{kj} = 1. \quad (73)$$

Now the solver would be able to activate and deactivate all the individual interaction terms. However, when changing the sequence of delivery there is more changing than only the active interaction terms. The time between two iso-centres depends on the beam-on times of all the iso-centres delivered in-between them. Thus,

we need to be able to relate the active terms, to the delivery sequence. To that end we introduce the matrix $\mathbf{s} \in \{0, 1\}^{N \times N}$, which describes which iso-centre is delivered at which position in the order of delivery. The column describes the position in the sequence and the non-zero row index determines the iso-centre used in that position

$$\mathbf{seq} = [2 \ 1 \ 3] \Leftrightarrow \mathbf{s} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \Leftrightarrow \mathbf{h} = \begin{bmatrix} - & 1 & 0 \\ 0 & - & 0 \\ 1 & 1 & - \end{bmatrix}. \quad (74)$$

Again, we need to ensure that \mathbf{s} is constrained to a feasible treatment, i.e. exactly one iso-centre delivered at every position (column) and each iso-centre (row index) must be used once

$$\forall i \in \{1, \dots, N\}, \quad \sum_{j=1}^N s_{ij} = 1 \quad \wedge \quad \sum_{j=1}^N s_{ij} = 1. \quad (75)$$

To enforce the consistency between the variables for iso-centre order \mathbf{s} and the active terms \mathbf{h} , we introduce a comparison function $comp_{jk}$ between two iso-centres j and k :

$$comp_{jk} = \sum_{i=1}^{N-1} (N - i)(s_{ij} - s_{ik}). \quad (76)$$

$comp_{jk} \in \{-(N-1), \dots, -1, 1, \dots, N-1\}$ will be positive if iso-centre j is delivered before iso-centre k (i.e. $t_j < t_k$), and negative if iso-centre j is delivered after iso-centre k (i.e. $t_j > t_k$). Thus, the binary interaction term control variable h_{jk} needs to be one for all combinations of j and k where $comp_{jk} < 0$. This can be enforced with the following inequality constraint:

$$\forall j, k \in \{1, \dots, N\}, j \neq k \quad h_{jk} \cdot comp_{jk} + h_{kj} \cdot comp_{kj} \leq -1. \quad (77)$$

Now that the active terms \mathbf{h} directly represent the current sequence, the starting time of every iso-centre j can be defined as follows:

$$t_j = \sum_{k=1, k \neq j}^N h_{jk}(\delta t_k + g_k). \quad (78)$$

We can now use this starting time definition to constrain the value of Δt_{jk} to the appropriate value based on the delivery sequence:

$$h_{jk} \Delta t_{jk} = h_{jk}(t_j - t_k). \quad (79)$$

The use of the binary variable h_{jk} is required here to ensure compatibility with the dual envelopes and PWLs which only allow values of Δt_{jk} within the predefined bounds $[\underline{\Delta t}_{jk}, \overline{\Delta t}_{jk}]$ where $\underline{\Delta t}_{jk} > 0$.

At this point, the problem fully is described and the convex relaxation of the BED could be optimised. However, with the additional binary variables introduced above, the problem is no longer linear. There are multiplications of the binary variable \mathbf{h} with the interaction terms \hat{w}_B , sequence \mathbf{s} , beam-on times δt , starting times t , and time in-between iso-centre deliveries Δt . Every one of these nonlinearities is resolved using the big M method described in appendix C.2. The required inequalities are completely analogous to that description and are omitted here for clarity:

$$\forall j, k \in \{1, \dots, N\}, j \neq k:$$

$$\forall i \in \{1, 2\}: \quad h_{jk} \hat{w}_{B,i,jk} \Rightarrow z_{jk}^{B_i} \quad (80)$$

$$\forall i \in \{1, \dots, N-1\}, l = j \vee l = k: \quad h_{jk} s_{il} \Rightarrow z_{jk,il}^S \quad (81)$$

$$h_{jk} \delta t_k \Rightarrow z_{jk}^{\delta t} \quad (82)$$

$$h_{jk} t_j \Rightarrow z_{jk}^{T^+} \quad (83)$$

$$h_{jk} t_k \Rightarrow z_{jk}^{T^-} \quad (84)$$

$$h_{jk} \Delta t_{jk} \Rightarrow z_{jk}^{\Delta t}. \quad (85)$$

Now the convex relaxation of the BED for the simultaneous beam-on time and sequence optimisation can be defined as follows:

$$\begin{aligned} \widehat{\text{BED}} = & \sum_{j=1}^N \dot{d}_j \delta t_j + A_1 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_1,j} + A_2 \sum_{j=1}^N \dot{d}_j^2 \hat{w}_{A_2,j} \\ & - B_1 \sum_{j \neq k} \dot{d}_k \dot{d}_j \hat{z}_{jk}^{B_1} - B_2 \sum_{j \neq k} \dot{d}_k \dot{d}_j \hat{z}_{jk}^{B_2} \end{aligned} \quad (86)$$

C.5. Beam-on time and sequence MILP problem

With the variables and constraints introduced above, we can now formulate a MILP problem for the beam-on time and sequence optimisation.

$$\begin{aligned} & \text{minimise } f(\widehat{\text{BED}}) \\ \text{w.r.t. } & \mathbf{s} \in \{0, 1\}^{N_{\text{iso}} \times N_{\text{iso}}} \quad \mathbf{h} \in \{0, 1\}^{N_{\text{iso}} \times N_{\text{iso}}} \quad \mathbf{z}^S \in \{0, 1\}^{N_{\text{iso}} \times N_{\text{iso}} \times N_{\text{iso}} - 1 \times N_{\text{iso}}} \\ & \delta \mathbf{t} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}}} \quad \mathbf{t} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}}} \quad \Delta \mathbf{t} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}} \times N_{\text{iso}}} \\ & \mathbf{z}^{\delta t} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}} \times N_{\text{iso}}} \quad \mathbf{z}^{T^+} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}} \times N_{\text{iso}}} \quad \mathbf{z}^{T^-} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}} \times N_{\text{iso}}} \\ & \mathbf{z}^{\Delta t} \in \mathbb{R}_{\geq 0}^{N_{\text{iso}} \times N_{\text{iso}}} \quad \hat{\mathbf{z}}^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}}} \\ & \hat{\mathbf{w}}^A \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}}} \quad \lambda^A \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times 2} \\ & \mathbf{L}(w^A) \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}}} \quad \nu^A \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times m_A} \quad \mathbf{y}^A \in \{0, 1\}^{2 \times N_{\text{iso}} \times m_A - 1} \\ & \hat{\mathbf{w}}^B \in \mathbb{R}_{\leq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}}} \quad \lambda^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}} \times 8} \\ & \mathbf{x}_1^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}}} \quad \mathbf{x}_2^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}}} \quad \mathbf{x}_3^B \in \mathbb{R}_{\leq 0}^{2 \times N_{\text{iso}}} \\ & \mathbf{L}(\mathbf{x}_1^B) \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}}} \quad \nu_1^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times N_{\text{iso}} \times m_{B,1}} \quad \mathbf{y}_1^B \in \{0, 1\}^{2 \times N_{\text{iso}} \times N_{\text{iso}} \times m_{B,1} - 1} \\ & \mathbf{L}(\mathbf{x}_2^B) \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}}} \quad \nu_2^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times m_{B,2}} \quad \mathbf{y}_2^B \in \{0, 1\}^{2 \times N_{\text{iso}} \times m_{B,2} - 1} \\ & \mathbf{L}(\mathbf{x}_3^B) \in \mathbb{R}_{\leq 0}^{2 \times N_{\text{iso}}} \quad \nu_3^B \in \mathbb{R}_{\geq 0}^{2 \times N_{\text{iso}} \times m_{B,3}} \quad \mathbf{y}_3^B \in \{0, 1\}^{2 \times N_{\text{iso}} \times m_{B,3} - 1} \\ & \text{cost}_i \in \mathbb{R}_{\geq 0}^{N_{\text{iso}}} \\ & \text{s.t. } \quad (73) \quad (75) - (79) \quad (\text{seq. constr.}) \\ & \quad (23) - (33) (80) - (85) \quad (\text{big } M \text{ constr.}) \\ & \quad (49) - (52) (57) - (61) \quad (\text{dual env. constr.}) \\ & \quad (34) - (40) (63) - (68) \quad (\text{PWL constr.}) \\ & \quad (69) - (70) \quad (86) \quad (\text{cost constr.}) \end{aligned} \quad (87)$$

Note: All variables that were specific to the $N_{ia} = \frac{N_{\text{iso}}(N_{\text{iso}} - 1)}{2}$ interaction terms in equation (71) are now expressed with $N_{\text{iso}} \times N_{\text{iso}}$ individual components to cover all possible interactions. Their diagonal elements (i.e. ‘self-interaction’ of an iso-centre) do not contribute to the optimisation problem.

C.5.1. *PWL considerations for the sequence optimisation.* As described above, to solve the treatment planning problem the convex relaxation of the BED is built from a set of bounds on the iso-centre timings (δt , Δt). For the beam-on time optimisation, the bounds of Δt follow directly from the bounds on the individual beam-on times. As the bounds on beam-on times tighten, so do the bounds on the possible Δt . Thus, the convex relaxation of the BED converges to the value of the full nonlinear version even if we only use the minimum of two support points in the PWLs. This substantially reduces the computational cost of solving the model.

When the order of iso-centre delivery can change, then the range of possible values for Δt will not shrink with the bounds on the beam-on time. Instead, its lower limit is determined by the shortest possible beam-on time and the upper limit is defined by the $(N - 1)$ longest possible beam-on times. In order to ensure accurate values for Δt (and thus the BED) are determined from the PWL of the corresponding nonlinear term, the number of support points needs to be increased. Preliminary tests determined that using 10 support points was sufficient to accurately determine the BED from the convex relaxation.

ORCID iDs

Thomas Klinge  <https://orcid.org/0000-0002-8559-0841>
 Hugues Talbot  <https://orcid.org/0000-0002-2179-3498>
 Ian Paddick  <https://orcid.org/0000-0002-3820-8654>
 Sébastien Ourselin  <https://orcid.org/0000-0002-5694-5340>
 Jamie R McClelland  <https://orcid.org/0000-0002-4922-0093>

References

- Burer S and Letchford A N 2012 Non-convex mixed-integer nonlinear programming: a survey *Surveys Oper. Res. Manage. Sci.* **17** 97–106
- Byrd R H, Lu P, Nocedal J and Zhu C 1995 A limited memory algorithm for bound constrained optimization *SIAM J. Sci. Comput.* **16** 1190–208
- Chen D-S, Batson R G and Dang Y 2009 *Applied Integer Programming* (Hoboken: John Wiley & Sons, Ltd) (<https://doi.org/10.1002/9781118166000>)
- Costa A and Liberti L 2012 Relaxations of multilinear convex envelopes: dual is better than primal *Experimental Algorithms (Lecture Notes in Computer Science)* ed R Klasing et al vol 7276 (Berlin, Heidelberg: Springer Berlin Heidelberg) (https://doi.org/10.1007/978-3-642-30850-5_9)
- Croes G A 1958 A method for solving traveling-salesman problems *Oper. Res.* **6** 791–812
- Fowler J F 1989 The linear-quadratic formula and progress in fractionated radiotherapy *Br. J. Radiol.* **62** 679–94
- Glover F 1975 Improved linear integer programming formulations of nonlinear integer problems *Manage. Sci.* **22** 455–60
- Graffeo C S, Perry A, Link M J, Brown P D, Young W F and Pollock B E 2021 Biological effective dose as a predictor of hypopituitarism after single-fraction pituitary adenoma radiosurgery: dosimetric analysis and cohort study of patients treated using contemporary techniques *Neurosurgery* **88** E330–5
- Graffeo C S, Donegan D, Erickson D, Brown P D, Perry A, Link M J, Young W F and Pollock B E 2020 The impact of insulin-like growth factor index and biologically effective dose on outcomes after stereotactic radiosurgery for acromegaly: cohort study *Neurosurgery* **87** 538–46
- Hallgren S, Hill M A, Thompson J M, Elliott A, Paddick I, Jones B and Hopewell J W 2019 Effects of variations in overall treatment time on the clonogenic survival of V79-4 cells: implications for radiosurgery *J. Radiosurg. SBRT* **6** 1–9 PMID: 30775069
- Hopewell J W, Millar W T, Lindquist C, Nordström H, Lidberg P and Gårding J 2013 Application of the concept of biologically effective dose (BED) to patients with vestibular schwannomas treated by radiosurgery *J. Radiosurg. SBRT* **2** 257–71
- Hopewell J W, Paddick I, Jones B and Klinge T 2021 Letter to the editor. Biologically effective dose and the treatment of AVMs *J. Neurosurg.* **134** 2007–8
- Hopewell J W, Millar W T and Lindquist C 2012 Radiobiological principles: their application to gamma knife therapy *Prog. Neurolog. Surg.* **25** 39–54
- Johnson D S and McGeoch L A 1997 The traveling salesman problem: a case study in local optimization *Local Search in Combinatorial Optimization* ed E H L Aarts and J K Lenstra (London: John Wiley and Sons) 215–310
- Jones B, Klinge T and Hopewell J W 2020 The influence of the α/β ratio on treatment time iso-effect relationships in the central nervous system *Int. J. Radiat. Biol.* **96** 903–9
- Jones B and Hopewell J W 2018 Modelling the influence of treatment time on the biological effectiveness of single radiosurgery treatments: derivation of ‘protective’ dose modification factors *Br. J. Radiol.* **92** 20180111
- Klinge T, Modat M, McClelland J R, Dimitriadis A, Paddick I, Hopewell J W, Walton L, Rowe J, Kitchen N and Ourselin S 2021 The impact of unscheduled gaps and iso-centre sequencing on the biologically effective dose in gamma knife radiosurgery *J. Radiosurg. SBRT* **7** 213–21 PMID: PMC8055240
- Köppe M 2012 On the complexity of nonlinear mixed-integer optimization *Mixed Integer Nonlinear Programming* ed J Lee and S Leyffer (New York, NY: Springer) 533–57
- Levivier M, Carrillo R E, Charrier R, Martin A and Thiran J-P 2018 A real-time optimal inverse planning for gamma knife radiosurgery by convex optimization: description of the system and first dosimetry data *J. Neurosurg.* **129** 111–7
- Lin M-H, Carlsson J G, Ge D, Shi J and Tsai J-F 2013 A review of piecewise linearization methods *Math. Problems Eng.* **2013** 1–8
- McCormick G P 1976 Computability of global solutions to factorable nonconvex programs: Part I. — Convex underestimating problems *Math. Program.* **10** 147–75
- Millar W T and Canney P A 1993 Derivation and application of equations describing the effects of fractionated protracted irradiation, based on multiple and incomplete repair processes. Part I. Derivation of equations *Int. J. Radiat. Biol.* **64** 275–91
- Millar W T, Hopewell J W, Paddick I, Lindquist C, Nordström H, Lidberg P and Gårding J 2015 The role of the concept of biologically effective dose (BED) in treatment planning in radiosurgery *Phys. Med. : PM : Int. J. Devoted Appl. Phys. Med. Biol. : Official J. Italian Assoc. Biomed. Phys. (AIFB)* **31** 627–33
- Nocedal J and Wright S J 2006 *Numerical Optimization (Springer Series in Operations Research and Financial Engineering)* 2nd edn (New York: Springer) (<https://doi.org/10.1007/978-0-387-40065-5>)
- Paddick I, Hopewell J W, Klinge T, Graffeo C S, Pollock B E and Sneed P K 2020 Letter: treatment outcomes and dose rate effects following gamma knife stereotactic radiosurgery for vestibular schwannomas *Neurosurgery* **86** E407–9
- Paddick I 2000 A simple scoring ratio to index the conformity of radiosurgical treatment plans *J. Neurosurg.* **93** 219–22
- Podgorsak E B, Bruce Pace G, Olivier A, Pla M and Souhami L 1989 Radiosurgery with high energy photon beams: a comparison among techniques *Int. J. Radiat. Oncol. *Biol. *Phys.* **16** 857–65
- Pop L A M, Millar W T, van der Plas M and van der Kogel A J 2000 Radiation tolerance of rat spinal cord to pulsed dose rate (PDR-) brachytherapy: the impact of differences in temporal dose distribution *Radiother. Oncol.* **55** 301–15
- Rikun A D 1997 A convex envelope formula for multilinear functions *J. Global Optim.* **10** 425–37
- Schulder M and Patil V 2008 The history of stereotactic radiosurgery *Principles and Practice of Stereotactic Radiosurgery* ed L S Chin and W F Regine (New York, NY: Springer) 3–7
- Sjölund J, Riad S, Hennix M and Nordström H 2019 A linear programming approach to inverse planning in gamma knife radiosurgery *Med. Phys.* **46** 1533–44
- Tardella F 2008 Existence and sum decomposition of vertex polyhedral convex envelopes *Optim. Lett.* **2** 363–75
- Tuleasca C, Paddick I, Hopewell J W, Jones B, Millar W T, Hamdi H, Porcheron D, Levivier M and Régis J 2019 Establishment of a therapeutic ratio for gamma knife radiosurgery of trigeminal neuralgia: the critical importance of biologically effective dose versus physical dose *World Neurosurg.* **134** E204–13
- Virtanen P et al 2020 SciPy 1.0: fundamental algorithms for scientific computing in python *Nat. Methods* **17** 261–72