




ELSEVIER

Contents lists available at ScienceDirect

Physics and Imaging in Radiation Oncology

journal homepage: www.elsevier.com/locate/phro

Original Research Article

Utilisation of Pareto navigation techniques to calibrate a fully automated radiotherapy treatment planning solution

Philip A. Wheeler^{a,*}, Michael Chu^a, Rosemary Holmes^a, Maeve Smyth^a, Rhydian Maggs^a, Emiliano Spezi^b, John Staffurth^c, David G. Lewis^a, Anthony E. Millin^a

^a Velindre Cancer Centre, Medical Physics, Cardiff, United Kingdom

^b Cardiff University, School of Engineering, Cardiff, United Kingdom

^c Cardiff University, School of Medicine, Cardiff, United Kingdom

ARTICLE INFO

Keywords:

Intensity modulated radiotherapy

VMAT

Treatment planning

Automation

Pareto navigation

Prostate cancer

Multicriteria optimization

ABSTRACT

Background and purpose: Current automated radiotherapy planning solutions do not allow for the intuitive exploration of different treatment options during protocol calibration. This work introduces an automated planning solution, which aims to address this problem through incorporating Pareto navigation techniques into the calibration process.

Materials and methods: For each tumour site a set of planning goals is defined. Utilising Pareto navigation techniques an operator calibrates the solution through intuitively exploring different treatment options: selecting the optimum balancing of competing planning goals for the given site. Once calibrated, fully automated plan generation is possible, with specific algorithms implemented to ensure trade-off balancing of new patients is consistent with that during calibration. Using the proposed methodology the system was calibrated for prostate and seminal vesicle treatments. The resultant solution was validated through quantitatively comparing the dose distribution of automatically generated plans (VMAT_{Auto}) against the previous clinical plan, for ten randomly selected patients.

Results: VMAT_{Auto} yielded statistically significant improvements in: PTV conformity indices, high dose bladder metrics, mean bowel dose, and the majority of rectum dose metrics. Of particular note was the reduction in mean rectum dose (median 25.1 Gy vs. 27.5 Gy), rectum V_{24.3Gy} (median 41.1% vs. 46.4%), and improvement in the conformity index for the primary PTV (median 0.86 vs. 0.79). Dosimetric improvements were not at the cost of other dose metrics.

Conclusions: An automated planning methodology with a Pareto navigation based calibration has been developed, which enables the complex balancing of competing trade-offs to be intuitively incorporated into automated protocols.

1. Introduction

Inverse radiotherapy planning is a time consuming, iterative process where optimal plan quality is not guaranteed [1]. A solution to this problem is automated planning (AP) where high quality treatment plans are generated fully autonomously [2–11]. AP has been implemented using a range of methodologies, which can be categorised within the following three broad domains: knowledge based planning (KBP), sequential ϵ -constraint optimisations (ϵ c) and protocol based automatic iterative optimisations (PB-AIO).

KBP utilises information from previously treated patients to inform the optimisation of future patients. The most common methods use

machine learning algorithms, trained on databases of historical treatment plans, to predict the achievable dose distribution [12,13] or dose volume histograms [14,15] for new patients. This information is utilised during the inverse optimisation process to generate a plan whose dose distribution best matches that predicted.

ϵ c generates plans according to a list of prioritised clinical goals, which are minimised in strict sequential order under the condition that lower priority goals must not compromise higher priority goals. Through the appropriate selection and ordering of goals, a single prioritised list can generate desirably balanced plans for individual patients within a given treatment site [2,8].

Finally, PB-AIO techniques load an initial set of objectives (either

* Corresponding author.

E-mail address: Philip.Wheeler@wales.nhs.uk (P.A. Wheeler).

<https://doi.org/10.1016/j.phro.2019.04.005>

Received 11 January 2019; Received in revised form 10 April 2019; Accepted 14 April 2019

2405-6316/ © 2019 The Authors. Published by Elsevier B.V. on behalf of European Society of Radiotherapy & Oncology. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

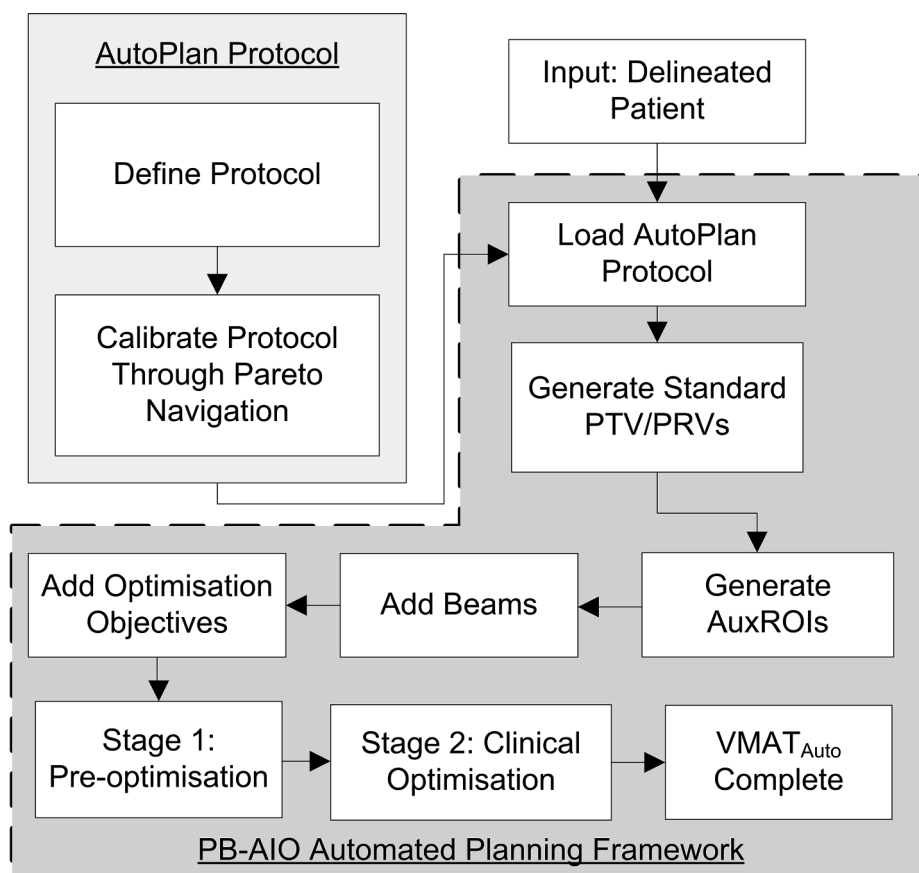


Fig. 1. Flowchart depicting the workflow of the proposed solution, with all items within the PB-AIO framework (as represented by the dashed area) fully automated. For each tumour site a calibrated AutoPlan protocol is required.

hard coded or derived from a site specific template) into the planning system's native optimiser. During the optimisation process, automated algorithms iteratively adjust objectives based on information from the optimised dose distribution to tailor the plan to the desired clinical aims [4,16,17]. Specific examples include: regularly updating objective positions such that a constant distance below the corresponding DVH line is maintained [18] and modulating objective weights such that the function's objective value (OV) tends towards a target OV during the course of the optimisation [4,9]. Through implementing these methods of dynamic objective adjustment it has been shown that a single set of initial objectives can yield plans with minimised organ at risk (OAR) doses and consistently balanced trade-offs across all patients within a given treatment site [4,9,18]. In this manuscript, we define objectives whose weight and position are modified in this specific manner as 'dynamic objectives'.

A key challenge in all three approaches is adequately and intuitively capturing the oncologist's experience and decision making during the calibration process, such that automated plans are congruent with clinical preference. KBP is dependent on the optimality of large datasets of previous clinical plans, which is not guaranteed, and both ϵ c and PB-AIO rely on trial and error to develop and refine automated protocols [19]. In this paper a novel solution to this problem is proposed through integrating Pareto navigation techniques directly into the calibration process.

Pareto navigation enables operators to intuitively explore differing treatment plan options such that an informed choice can be made on the optimal balancing of competing trade-offs [20,21]. Navigation is performed on a pre-calculated set of Pareto optimal plans, which aim to sample clinically relevant parts of the Pareto surface. In this regard, a treatment plan is considered Pareto optimal when improvement of a given trade-off is only possible at the detriment of another, with the

Pareto surface being represented by an infinite set of such plans. On an individual patient basis, intuitively exploring the Pareto surface through Pareto navigation has shown to reduce the need for trial and error, and yield plans more congruent with oncologists' clinical preferences [22]. It is expected that incorporating Pareto navigation into the automated planning calibration process will yield similar benefits at the patient cohort level.

The purpose of this work is to present the methodology of EdgeVcc (Experience Driven plan Generation Engine by Velindre Cancer Centre): a PB-AIO based automated planning solution, designed to be applicable across a range of radiotherapy treatment sites and uniquely calibrated using Pareto navigation principles. The first section of this paper provides a detailed description of the proposed methodology alongside its associated algorithms. The second section presents an example of its application to the tumour site of prostate and seminal vesicles (PSV).

2. Methods and materials

2.1. Patient dataset

20 patients previously treated at Velindre Cancer Centre between July and December 2015 were randomly selected into a calibration ($n = 10$) and validation ($n = 10$) cohort. Patients were planned on computed tomography scans of 3 mm slice thickness with prostate, seminal vesicles, rectum, bladder and bowel delineated. Prostate + seminal vesicles were expanded 10 mm isotropically to form the planning target volume PTV48 and prostate expanded by 5 mm (6 mm craniocaudally) to form PTV60, with the PTV's suffix denoting its prescribed dose in Gy. All study patients were previously treated with VMAT on Elekta Agility linear accelerators with treatment plans generated manually in Oncentra Masterplan (v4.3, Elekta Ltd, Crawley)

Table 1
Summary of objective weight scaling factors.

Scaling Factor	Description
F_V^i	Scales objective weight according the volume of its corresponding ROI.
F_T^i	Scales objective weight according to the objective's target dose level (D_T^i). This removes an unwanted dependency of RayStation's objective functions on D_T^i .
F_C^i	A hardcoded constant utilised to reduce the weight of PTV sub-volume objectives to avoid skin boosting and reduce conflicts within the PTV/OAR overlap region.
F_N^i	F_N^i enables $w_{initial}^i$ to be modified for an individual planning goal. The purpose is to bring $w_{initial}^i$ closer to the anticipated final weight for dynamic objectives.

using a single 6MV 360° arc, simultaneous integrated boost technique. Treatments were prescribed for 20 fractions, and manually planned to local clinical goals (Supplementary table 1) using a class solution based methodology.

2.2. Automated system

2.2.1. System overview

The proposed solution (Fig. 1) was developed in the treatment planning system RayStation (Raysearch Laboratories, Stockholm) using custom python scripts. For each tumour site a set of planning goals is defined within an 'AutoPlan protocol' (Section 2.2.2). On a selected calibration patient(s) Pareto navigation techniques are utilised to derive a set of planning goal weighting factors, which correspond to a clinically desirable point on the Pareto surface for the given patient (Section 2.2.4). The weighting factors are stored within the AutoPlan protocol, which then forms the input for automated planning of new patients. Automated plan generation is based on a PB-AIO framework (Section 2.2.3) that utilises 'dynamic objectives' to ensure OAR doses are minimised and trade-off balancing for new patients is consistent with that selected during protocol calibration.

2.2.2. AutoPlan protocol

The treatment modality, beam arrangement, standard PTV and planning volume at risk (PRV) margins, and planning goals are defined within the AutoPlan protocol. Planning goals guide the optimisation process and are stratified into three priority levels: primary normal tissue goals (P_1), target goals (P_2) and trade-off goals (P_3). The optimisation methodology aims to meet goals in order of their priority, with compromise to target goals permissible by P_1 but not P_3 . Trade-off goals are assigned a group number, which determines the order in which they are explored during the calibration process. Goals of the same parameter type and clinical relevance (e.g. low dose rectum objectives) are grouped to reduce degrees of freedom during calibration. The planning goals for PSV are presented in the Supplementary table 2.

Planning goals are designed to be clinically intuitive, with no specification of weighting factors required. Weighting factors are instead derived through two distinct processes. For P_1 and P_2 , the clinical preference across all tumour sites when balancing conflicting goals is explicitly defined: target coverage is compromised to maintain normal tissue goals. Conflicting goals are therefore explicitly handled through region of interest (ROI) retraction algorithms, enabling weights to be defined by simple hard coded algorithms. In contrast, conflicting P_3 trade-offs require careful balancing for each tumour site; a complex process requiring specialist clinical judgment. Weights are therefore derived through utilisation of Pareto navigation techniques (Section 2.2.4).

2.2.3. PB-AIO framework

The following PB-AIO framework is used to generate both the final automated plans and those utilised during the calibration process.

2.2.3.1. Auxiliary optimisation volumes. Following PTV and PRV creation, a standard set of auxiliary optimisation ROIs (AuxROIs) are generated according to the algorithms detailed in the Supplementary

file S1. AuxROIs have two purposes. For conformity related planning goals AuxROIs enable a higher level of geometric specificity. For target goals, in line with ICRU 83 [23], AuxROIs subdivide each PTV into three sub-volumes to avoid conflicting planning aims: PTV_{SV-1} is retracted from the skin surface and proximal primary OARs, PTV_{SV-2} consists of areas of PTV within the skin surface or extending into air, and PTV_{SV-3} is the PTV volume not covered by PTV_{SV-1} or PTV_{SV-2}, which represents parts of PTV proximal to primary OARs. It is through this subdivision that P_2 goals are compromised for P_1 goals and IMRT flash is secured for superficial PTVs.

2.2.3.2. Initial optimisation objectives. Following treatment beam/arc definition an initial set of optimisation objectives are loaded into RayStation's native optimiser. Optimisation objectives are derived from the defined planning goals according to the algorithms specified in Supplementary file S2, with the initial weight, $w_{initial}^i$, for the i^{th} objective defined by:

$$w_{initial}^i = w_{nom}^i F_V^i F_T^i F_C^i F_N^i \quad (1)$$

where w_{nom}^i is the nominal weight of the planning goal from which the objective is derived, and F_V^i , F_T^i , F_C^i and F_N^i are optimisation objective specific scaling factors. Each scaling factor is summarised in Table 1, with full definitions provided in the Supplementary file S3. For P_1 and P_2 goals, w_{nom}^i is an empirically derived hardcoded value (Supplementary table 3), intended to be common across all treatment sites. For P_3 goals, w_{nom}^i is generated through the Pareto navigation calibration process.

2.2.3.3. Plan optimisation. The employed optimisation algorithms (Fig. 2) consist of two stages: a pre-optimisation, which sets initial P_3 objective target values (T_{p3}^i) and a main optimisation, which generates the final clinical plan. During the optimisation process, P_3 goals are implemented as 'dynamic objectives' with the aim of minimising OAR doses and keeping trade-off balancing across patients consistent. In the following description of the methodology, we define Δ^i by the equation:

$$\Delta^i = \frac{D_{p3}^i - T_{p3}^i}{x^i} \quad (2)$$

where D_{p3}^i is the current value of the planning goal's corresponding dose parameter (c.f. Supplementary table 2) and x^i equals D_{presc} for dose objectives and V_{ROI}^i for volume objectives, where V_{ROI}^i is the volume of the i^{th} objective's corresponding ROI.

For stage one, a fluence-based optimisation, which allows beam intensity to be modulated with minimal physical limits, is performed. Following the optimisation, if Δ^i does not lie within the range [0.15–0.5] (or [0.0–0.5] if $T_{p3}^i = 0$) for each dynamic objective, T_{p3}^i is updated according to equation (3) (with the variable δ set to 0.35) and the optimisation rerun.

$$T_{p3}^i = D_{p3}^i - x^i \delta \quad (3)$$

The process is repeated until Δ^i lies within the specified bounds across all dynamic objectives. Bounding Δ^i in this manner ensures P_3 optimisation objectives are a significant but not dominant component of the composite objective function, resulting in P_3 goals being minimised without significantly compromising P_1 or P_2 goals. The resulting dose distribution, which is generated within 1–2 minutes, provides an

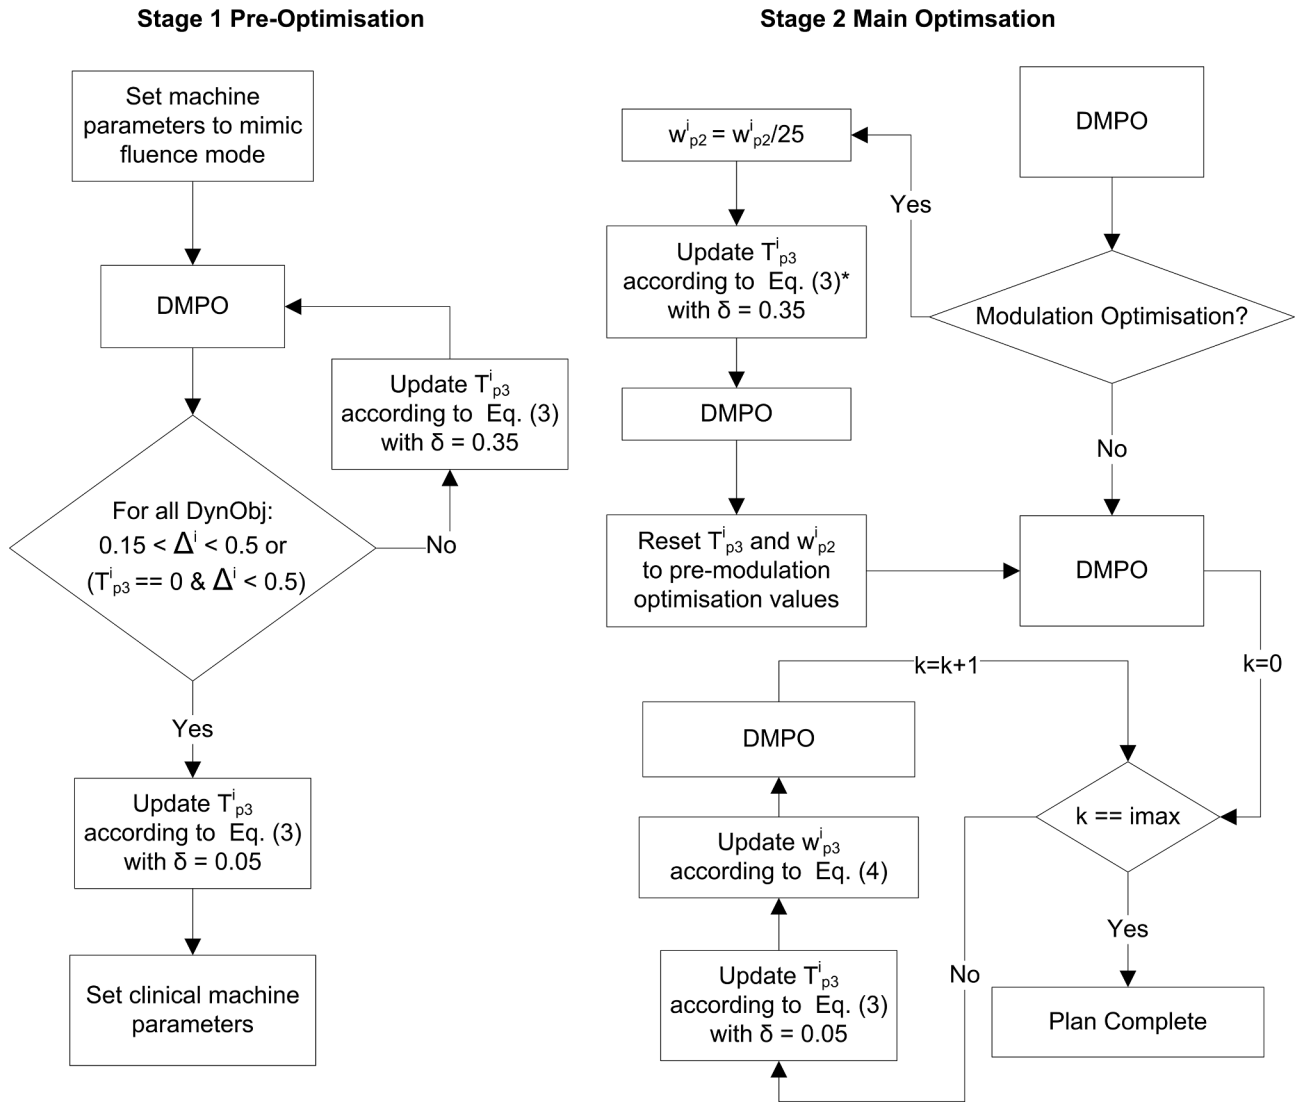


Fig. 2. Flowchart of the stage 1 and stage 2 optimisation algorithms, where: $w_{p2}^i, T_{p3}^i, w_{p3}^i, \Delta^i$ and δ are defined in the main manuscript, DynObj is an abbreviation for dynamic objectives and DMPO indicates a direct machine parameter optimisation. *During the modulation optimisation, for Eq. (3), D_{p3}^i is calculated from the final stage one pre-optimisation distribution.

approximate prediction of the final, fully optimised, clinical solution. Based on this distribution, T_{p3}^i for the main optimisation is set according to Eq. (3) with $\delta = 0.05$. Fluence-based VMAT optimisations are not possible in RayStation, therefore stage one treatment arcs are approximated through 15 equi-spaced static IMRT fields.

For stage two, a preliminary direct machine parameter optimisation (DMPO), where optimisations are bound by the machine’s physical limits, is executed to generate an initial set of segments. An optional modulation optimisation is performed where P_2 objective weights (w_{p2}^i) are reduced by a factor of 25 and, using the stage one pre-optimisation distribution as the reference dose, T_{p3}^i set according to Eq. (3) with $\delta = 0.35$. This prioritises the minimisation of P_3 objectives during the initial phases of the plan generation process and results in a reduction of OAR doses at the lower dose levels in the final clinical plan. This however is at the expense of increased modulation and MU. After the modulation optimisation, objective positions and weights are reverted to their original values. Finally, in the main stage of the plan generation process, multiple DMPOs are performed with T_{p3}^i and w_{p3}^i adjusted after each round. Using the stage two dose distribution as the reference, T_{p3}^i is updated according to Eq. (3) (with $\delta = 0.05$) and w_{p3}^i according to:

$$w_{p3}^i = \frac{OV_c^i w_c^i}{OV_c^i} \tag{4}$$

where w_c^i, OV_c^i and OV_i^i are the current weight, current OV and target OV of the i^{th} objective respectively, with OV_i^i derived from w_{norm}^i according to the Supplementary file S4. A function’s OV is defined as the product of its weight and function, therefore this iterative weight adjustment ensures that across the multiple DMPOs, OV_c^i tends towards OV_i^i .

2.2.4. AutoPlan protocol calibration

A flowchart of the calibration process, used to generate w_{norm}^i for each P_3 planning goal, is provided in the Supplementary file S5. A calibration patient data set, consisting of typically 10–20 delineated patients, is defined. From this dataset a single ‘navigation patient’ is selected and the Pareto navigation process started.

Initially all P_3 nominal weights are set to zero. For the first P_3 group, multiple plans (nominally five) are generated, each with a different value of w_{norm}^i applied to the group. The operator uses a slider to navigate through convex combinations of the differently weighted plans, with the navigated dose distribution and associated DVH updated in real time to inform the decision-making (Fig. 3). The operator selects

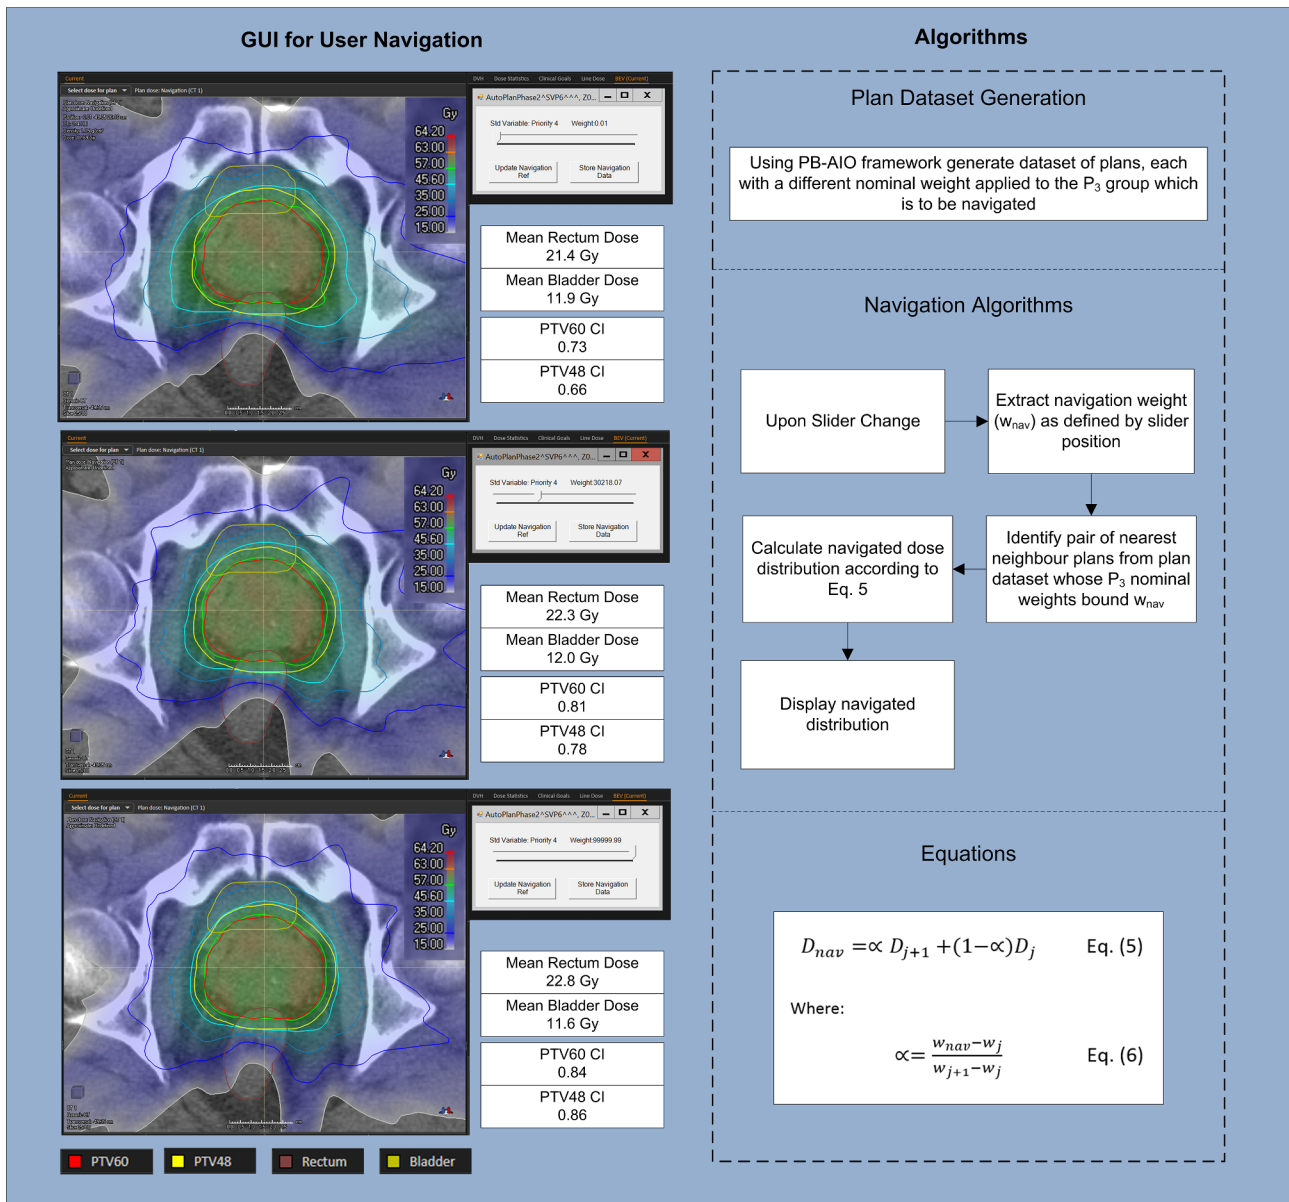


Fig. 3. (LHS) Screenshots demonstrating using the slider GUI to navigate through different weighted options for the PSV conformality goal (P_3 group 4). The displayed DVH metrics, which are not part of the calibration GUI, demonstrate the trade-off between the Paddick’s conformity index (CI) for both PTVs and organ at risk mean doses. Isodose legend is enlarged for clarity. (RHS) Algorithms associated with the navigation module where: w_j and w_{j+1} are the nominal weights of the nearest neighbour plans j and $j + 1$ respectively, whose weights bound the navigation weight, w_{nav} ; D_j and D_{j+1} correspond to the dose distribution of plan j and $j + 1$ respectively; and D_{nav} represents the estimated navigated dose distribution.

what they consider to be the optimum group weighting and the navigated weight is stored in the AutoPlan protocol. The process is then repeated for the next group using the updated protocol. Once all groups are navigated a final ‘rebalancing’ navigation is performed on a set of plans with differing factors (range 0.25–1.25) applied globally to all P_3 nominal weights. This process allows the ratio of P_3 weights, and P_1 and P_2 weights to be explored, ensuring a solution can be selected where higher priority goals are not compromised. Once this first calibration round is complete the solution is tested across all calibration patients, with amendments to planning goals or additional navigations (on selected P_3 groups) performed as required to refine the solution.

Generating one set of navigation plans for a P_3 group takes 1–3 h, depending on plan complexity, and each group must be optimised and navigated sequentially. Navigations are initially performed on a single patient, however where there are large inter-patient anatomical variations, repeat navigations over population outliers may be required to

ensure the solution is robust across the whole patient cohort. When navigating over multiple patients the operator decides whether the P_3 group weighting is based a particular patient, or averaged over multiple patients. The calibration process can be considered equivalent to navigating the Pareto surface one dimension (or P_3 group) at a time, with operators using clinical experience and expertise to balance competing trade-offs.

2.3. Application to prostate cancer

All 10 calibration patients alongside their previous clinically approved treatment plan (VMAT_{Clinical}) were available during the AutoPlan protocol calibration for PSV. Once successfully calibrated a final automated plan (VMAT_{Auto}) was generated for all study patients using identical arc configurations to VMAT_{Clinical}. To assess the efficacy of the calibrated automated solution, plan quality was quantitatively

compared to VMAT_{Clinical} using the local clinical goals, alongside D98%, D2% and Paddick's Conformity Index [24] for each target volume. The statistical significance of any differences was assessed using two-sided Wilcoxon matched-paired signed-rank tests.

3. Results

3.1. AutoPlan protocol calibration

Protocol calibration was performed by one physicist with a radiation oncologist providing clinical input on trade-off prioritisation prior to calibration. 15 individual trade-off navigations were required to calibrate the AutoPlan protocol. All navigations were performed on a single patient, with planning goals manually modified twice after reviewing results across all patients in the calibration dataset. Key planning goal updates included the addition of bowel and low dose bladder planning goals. The final nominal weights are presented in [Supplementary table 3](#).

3.2. Comparison with VMAT_{Clinical}

A summary of the quantitative comparison of VMAT_{Auto} and VMAT_{Clinical} for the validation cohort is presented in [Table 2](#). In comparison with the previous clinical plans, VMAT_{Auto} yielded statistically significant ($p < 0.05$) improvements in: PTV conformity indices, high dose bladder metrics, mean bowel dose, and the majority of rectum dose metrics. Of particular note was the reduction in mean rectum dose (median 25.1 Gy vs. 27.5 Gy), rectum V_{24.3Gy} (median 41.1% vs 46.4%), and improvements in CI_{PTV60} (median 0.86 vs. 0.79) and CI_{PTV48} (median 0.84 vs. 0.77). Dosimetric improvements were not at the expense of other dose metrics, with observed detriments either statistically or clinically insignificant. In terms of modulation, VMAT_{Auto} yielded plans with a median MU 10% higher than VMAT_{Clinical}.

Extending the comparison across all 20 study patients yielded similar results ([Supplementary tables 4 and 5](#)), with all treatment plans meeting the locally defined mandatory goals for clinical acceptability.

Table 2
Dosimetric comparison of VMAT_{Auto} and VMAT_{Clinical} for the validation patient cohort.

Metric	VMAT _{Auto}		VMAT _{Clinical}		p value	
	Median	Range	Median	Range		
PTV60	D98% (Gy)	57.8	57.7–58.0	57.9	57.6–58.3	0.17
	D2% (Gy)	61.7	61.6–61.7	61.7	61.3–62.2	0.33
	CI	0.86	0.85–0.87	0.79	0.76–0.82	0.01
PTV48	D98% (Gy)	46.8	46.5–47.4	47.0	46.6–47.7	0.06
	D2% (Gy)	58.8	58.5–59.2	59.4	58.8–59.8	0.01
	CI	0.84	0.82–0.87	0.77	0.75–0.79	0.01
Rectum	V24.3 Gy (%)	41.1	27.3–63.6	46.4	30.6–66.1	0.01
	V40.5 Gy (%)	24.2	15.3–39.6	24.6	16.4–41.2	0.06
	V52.7 Gy (%)	11.0	4.7–16.5	12.8	5.6–18.9	0.01
	V60.8 Gy (%)	0.1	0.0–0.5	0.0	0.0–0.3	0.09
	DMean (Gy)	25.1	17.5–32.3	27.5	20.5–34.0	0.01
Bladder	V40.5 Gy (%)	15.3	8.8–31.5	15.4	8.9–31.0	0.39
	V52.7 Gy (%)	7.8	3.0–16.8	8.0	3.3–17.9	0.04
	V56.8 Gy (%)	5.3	2.1–11.7	5.7	2.3–12.8	0.03
	DMean (Gy)	18.7	13.1–30.7	19.1	13.9–31.2	0.51
Bowel	V36.5 Gy (cm ³)	0.0	0.0–0.7	0.0	0.0–0.9	0.27
	V44.6 Gy (cm ³)	0.0	0.0–0.0	0.0	0.0–0.1	0.32
	DMean (Gy)	6.4	3.4–11.2	7.5	3.9–14.1	0.01
External	D1.8 cm ³ (Gy)	61.6	61.5–61.8	61.7	61.2–62.4	0.33
Plan MU	MU	600	582–653	546	496–627	0.01

Statistical significance: results where $p \leq 0.05$ are presented in bold.
CI: Paddick's Conformity Index for the specified PTV.

4. Discussion

To the authors' knowledge this paper presents the first automated planning solution that directly incorporates Pareto navigation techniques into the calibration process. Compared to clinical practice ([Table 2](#)), VMAT_{Auto} consistently yielded plans with improved conformity and reduced organ at risk doses, with no clinically relevant compromise to other dose metrics. As VMAT_{Clinical} and VMAT_{Auto} are generated in differing planning systems (Oncontra vs RayStation), these results are not intended to form a robust assessment as to their relative efficacy, this is the subject of future work. Instead they provide sound evidence that directly calibrating automated solutions through Pareto navigation is feasible and yield plans of improved dosimetric quality compared to current clinical practice.

Utilisation of Pareto navigation within the calibration process was observed to have two main benefits. Firstly, exploring differently weighted options via a sliding interface and live dose distribution allowed trade-off options to be explored in a visually intuitive manner. Secondly, an automated solution was derived in a time efficient manner with minimal trial and error. These advantages have been demonstrated on a per-patient basis by a number of studies [[22,25,26](#)] and this work indicates that similar benefits can be realised by applying this technique at a patient cohort level.

A potential weakness of Pareto navigation is that there can be clinically relevant discrepancies between the navigated dose distribution and that of the final deliverable plan [[27](#)]. To minimise these discrepancies the following approaches were adopted. Firstly, the Pareto surface was sampled one dimension at a time to reduce interpolation errors during navigation (whilst maintaining a reasonable computational cost). Secondly, the Pareto dataset was populated with deliverable plans, ensuring navigations were performed on clinically achievable solutions. By utilising these approaches, discrepancies throughout the calibration process were of negligible clinical significance ([Supplementary table 6](#)). This navigation methodology does however have a potential weakness, in that by limiting the navigation to one P₃ group at a time a full exploration of the Pareto surface is not performed. Whilst this was not considered a problem for the relatively simple site of PSV cancer, for more complex sites, navigation of multiple P₃ groups in parallel may be required to derive the most clinically desirable solution.

An interesting finding from this study was that Pareto navigation across a single patient appears sufficient for a successful calibration. Whilst modifications to planning goals were required after reviewing results across all calibration patients, these adjustments were due to deficiencies in the original set of planning goals, which were highlighted by differing patient geometries (e.g. proximity of bowel to PTV demonstrating requirement for bowel planning goals), rather than an inappropriate calibration. The ability to calibrate automated solutions against small patient cohorts (5–10 subjects) is in-line with examples in the literature for ec [[8](#)] and PB-AIO [[10](#)] solutions, and should enable a more efficient automation of novel techniques or protocols than KBP solutions, which require the manual generation of large patient datasets for each change in clinical practice.

The implemented calibration methodology requires algorithms that balance trade-offs consistently across differing patients. Dynamically adapting trade-off objective positions and weights during the optimisation was hypothesised to fulfil this function. By implementing 'dynamic objectives' a single calibrated AutoPlan protocol was found to yield appropriately balanced, clinically acceptable plans across all 20 study patients. In terms of robustness to patient geometry, even when PTV/OAR overlap differed considerable from the navigation patient, trade-off balancing was observed to be appropriate ([Fig. 4](#)).

In summary, a novel automated planning solution has been developed, which for the first time directly incorporates Pareto navigation into the calibration process. The solution has been successfully calibrated for the site of PSV, yielding clinically acceptable, appropriately

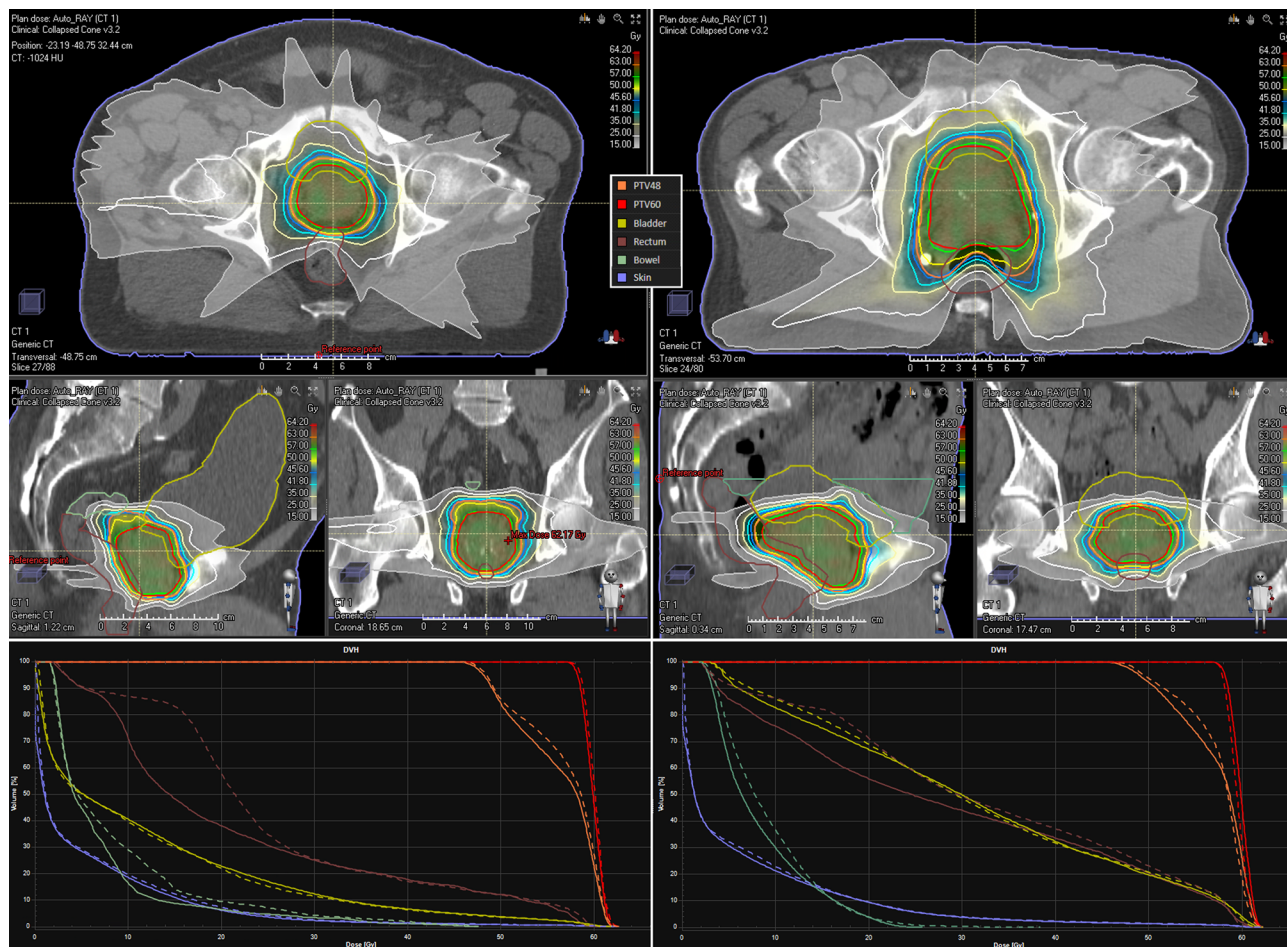


Fig. 4. DVH and dose distributions for the navigation patient (LHS) and patient 7 in the validation cohort (RHS), demonstrating the robustness of the automated solution to different anatomy. For both patients the DVH results for VMAT_{Clinical} are provided for reference (dashed line). For patient 7 the overlap of rectum with PTV60 and PTV48 was 9% and 24% respectively (c.f. 4% and 13% respectively for the navigation patient), and for bladder 8% and 19% respectively (c.f. 1% and 4% for the navigation patient). The PTV/OAR overlaps for patient 7 were all greater than the 89th percentile when considering the overlaps of all 20 study patients.

balanced treatment plans, fully autonomously.

Acknowledgements

This research is funded by Health and Care Research Wales and sponsored by Velindre NHS Trust. The authors would like to thank Mrs Sue Campbell and Dr Jim Fitzgibbon who represented the views of the patients and public throughout the research.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.phro.2019.04.005>.

References

- [1] Nelms BE, Robinson G, Markham J, Velasco K, Boyd S, Narayan S, et al. Variation in external beam treatment plan quality: an inter-institutional study of planners and planning systems. *Pract Radiat Oncol* 2012;2:296–305. <https://doi.org/10.1016/j.prro.2011.11.012>.
- [2] Voet PWJ, Dirx MLP, Breedveld S, Al-Mamgani A, Incrocci L, Heijmen BJM. Fully automated volumetric modulated arc therapy plan generation for prostate cancer patients. *Int J Radiat Oncol Biol Phys* 2014;88:1175–9. <https://doi.org/10.1016/j.ijrobp.2013.12.046>.
- [3] Wu B, McCutt T, Zahurak M, Simari P, Pang D, Taylor R, et al. Fully automated simultaneous integrated boosted-intensity modulated radiation therapy treatment planning is feasible for head-and-neck cancer: a prospective clinical study. *Int J Radiat Oncol Biol Phys* 2012;84:e647–53. <https://doi.org/10.1016/j.ijrobp.2012.06.047>.
- [4] Song Y, Wang Q, Jiang X, Liu S, Zhang Y, Bai S. Fully automatic volumetric modulated arc therapy plan generation for rectal cancer. *Radiother Oncol* 2016;119:531–6. <https://doi.org/10.1016/j.radonc.2016.04.010>.
- [5] Hansen CR, Bertelsen A, Hazell I, Zukauskaite R, Gyldenkerne N, Johansen J, et al. Automatic treatment planning improves the clinical quality of head and neck cancer treatment plans. *Clin Transl Radiat Oncol* 2016;1:1–7. <https://doi.org/10.1016/j.ctro.2016.08.001>.
- [6] Purdie TG, Dinniwell RE, Fyles A, Sharpe MB. Automation and intensity modulated radiation therapy for individualized high-quality tangent breast treatment plans. *Int J Radiat Oncol Biol Phys* 2014;90:688–95. <https://doi.org/10.1016/j.ijrobp.2014.06.056>.
- [7] Sharfo AWM, Stieler F, Kupfer O, Heijmen BJM, Dirx MLP, Breedveld S, et al. Automated VMAT planning for postoperative adjuvant treatment of advanced gastric cancer. *Radiat Oncol* 2018;13:74. <https://doi.org/10.1186/s13014-018-1032-z>.
- [8] Buschmann M, Sharfo AWM, Penninkhof J, Seppenwoolde Y, Goldner G, Georg D, et al. Automated volumetric modulated arc therapy planning for whole pelvic prostate radiotherapy. *Strahlentherapie Und Onkol* 2017;333–42. <https://doi.org/10.1007/s00066-017-1246-2>.
- [9] Quan EM, Chang JY, Liao Z, Xia T, Yuan Z, Liu H, et al. Automated volumetric modulated arc therapy treatment planning for stage III lung cancer: How does it compare with intensity-modulated radiotherapy? *Int J Radiat Oncol Biol Phys* 2012;84:e69–76. <https://doi.org/10.1016/j.ijrobp.2012.02.017>.
- [10] Vanderstraeten B, Goddeeris B, Vandecasteele K, van Eijkeren M, De Wagter C, Lievens Y. Automated instead of manual treatment planning? A plan comparison based on dose-volume statistics and clinical preference. *Int J Radiat Oncol* 2018. <https://doi.org/10.1016/j.ijrobp.2018.05.063>.
- [11] Buergy D, Sharfo AWM, Heijmen BJM, Voet PWJ, Breedveld S, Wenz F, et al. Fully automated treatment planning of spinal metastases – a comparison to manual planning of Volumetric Modulated Arc Therapy for conventionally fractionated

- irradiation. *Radiat Oncol* 2017;12:1–7. <https://doi.org/10.1186/s13014-017-0767-2>.
- [12] McIntosh C, Purdie TG. Voxel-based dose prediction with multi-patient atlas selection for automated radiotherapy treatment planning. *Phys Med Biol* 2017;62:415–31. <https://doi.org/10.1088/1361-6560/62/2/415>.
- [13] Fan J, Wang J, Chen Z, Hu C, Zhang Z, Hu W. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Med Phys* 2018. <https://doi.org/10.1002/mp.13271>.
- [14] Yuan L, Ge Y, Lee WR, Yin FF, Kirkpatrick JP, Wu QJ. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans. *Med Phys* 2012;39:6868. <https://doi.org/10.1118/1.4757927>.
- [15] Appenzoller LM, Michalski JM, Thorstad WL, Mutic S, Moore KL. Predicting dose-volume histograms for organs-at-risk in IMRT planning. *Med Phys* 2012;39:7446–61. <https://doi.org/10.1118/1.4761864>.
- [16] Zhang X, Li X, Quan EM, Pan X, Li Y. A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Phys Med Biol* 2011;56:3873–93. <https://doi.org/10.1088/0031-9155/56/13/009>.
- [17] Winkel D, Bol GH, van Asselen B, Hes J, Scholten V, Kerkmeijer LGW, et al. Development and clinical introduction of automated radiotherapy treatment planning for prostate cancer. *Phys Med Biol* 2016;61:8587–95. <https://doi.org/10.1088/1361-6560/61/24/8587>.
- [18] Tol JP, Dahele M, Peltola J, Nord J, Slotman BJ, Verbakel WFAR. Automatic interactive optimization for volumetric modulated arc therapy planning. *Radiat Oncol* 2015;10:75. <https://doi.org/10.1186/s13014-015-0388-6>.
- [19] Hussein M, Heijmen BJM, Verellen D, Nisbet A. Automation in intensity modulated radiotherapy treatment planning—a review of recent innovations. *Br J Radiol* 2018;91:20180270. <https://doi.org/10.1259/bjr.20180270>.
- [20] Craft D, Halabi T, Shih HA, Bortfeld T. An approach for practical multiobjective IMRT treatment planning. *Int J Radiat Oncol Biol Phys* 2007;69:1600–7. <https://doi.org/10.1016/j.ijrobp.2007.08.019>.
- [21] Thieke C, Küfer KH, Monz M, Scherrer A, Alonso F, Oelfke U, et al. A new concept for interactive radiotherapy planning with multicriteria optimization: first clinical evaluation. *Radiother Oncol* 2007;85:292–8. <https://doi.org/10.1016/j.radonc.2007.06.020>.
- [22] Craft DL, Hong TS, Shih HA, Bortfeld TR. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *Int J Radiat Oncol Biol Phys* 2012;82:83–90. <https://doi.org/10.1016/j.ijrobp.2010.12.007>.
- [23] International Commission on Radiation Units and Measurements. Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy IMRT (Report 83). *J ICRU* 2010;10. doi: 10.1093/jicru/10.1.Report83.
- [24] Paddick I. A simple scoring ratio to index the conformity of radiosurgical treatment plans. *J Neurosurg* 2000;93:219–22.
- [25] Müller BS, Shih HA, Efstathiou JA, Bortfeld T, Craft D. Multicriteria plan optimization in the hands of physicians: a pilot study in prostate cancer and brain tumors. *Radiat Oncol* 2017;12:1–11. <https://doi.org/10.1186/s13014-017-0903-z>.
- [26] Kierkels RG, Visser R, Bijl HP, Langendijk JA, van't Veld AA, Steenbakkers RJ, et al. Multicriteria optimization enables less experienced planners to efficiently produce high quality treatment plans in head and neck cancer radiotherapy. *Radiat Oncol* 2015;10. <https://doi.org/10.1186/s13014-015-0385-9>.
- [27] Kyrouti A, Petersson K, Ghandour S, Pachoud M, Matzinger O, Ozsahin M, et al. Discrepancies between selected Pareto optimal plans and final deliverable plans in radiotherapy multi-criteria optimization. *Radiother Oncol* 2016;120:346–8. <https://doi.org/10.1016/j.radonc.2016.05.018>.