



Keeping your best options open with AI-based treatment planning in prostate and cervix brachytherapy

Leah R.M. Dickhoff^{1,a}, Renzo J. Scholman^{2,3,a,*}, Danique L.J. Barten⁴, Ellen M. Kerkhof¹,
Jelmen J. Roorda⁴, Laura A. Velema¹, Lukas J.A. Stalpers^{4,5,6}, Bradley R. Pieters^{4,5},
Peter A.N. Bosman^{2,3}, Tanja Alderliesten¹

¹Department of Radiation Oncology, Leiden University Medical Center, Leiden, The Netherlands

²Evolutionary Intelligence Group, Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

³Faculty Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Delft, The Netherlands

⁴Department of Radiation Oncology, Amsterdam University Medical Centers (location University of Amsterdam), Amsterdam, The Netherlands

⁵Cancer Center Amsterdam, Cancer Treatment and Quality of Life, Amsterdam, The Netherlands

⁶Cancer Center Amsterdam, Cancer Biology and Immunology, Amsterdam, The Netherlands

ABSTRACT

PURPOSE: Without a clear definition of an optimal treatment plan, no optimization model can be perfect. Therefore, instead of automatically finding a single “optimal” plan, finding multiple, yet different near-optimal plans, can be an insightful approach to support radiation oncologists in finding the plan they are looking for.

METHODS AND MATERIALS: BRIGHT is a flexible AI-based optimization method for brachytherapy treatment planning that has already been shown capable of finding high-quality plans that trade-off target volume coverage and healthy tissue sparing. We leverage the flexibility of BRIGHT to find plans with similar dose-volume criteria, yet different dose distributions. We further describe extensions that facilitate fast plan adaptation should planning aims need to be adjusted, and straightforwardly allow incorporating hospital-specific aims besides standard protocols.

RESULTS: Results are obtained for prostate ($n=12$) and cervix brachytherapy ($n=36$). We demonstrate the possible differences in dose distribution for optimized plans with equal dose-volume criteria. We furthermore demonstrate that adding hospital-specific aims enables adhering to hospital-specific practice while still being able to automatically create cervix plans that more often satisfy the EMBRACE-II protocol than clinical practice. Finally, we illustrate the feasibility of fast plan adaptation.

CONCLUSIONS: Methods such as BRIGHT enable new ways to construct high-quality treatment plans for brachytherapy while offering new insights by making explicit the options one has. In particular, it becomes possible to present to radiation oncologists a manageable set of alternative plans that, from an optimization perspective are equally good, yet differ in terms of coverage-sparing trade-offs and shape of the dose distribution. © 2023 The Authors. Published by Elsevier Inc. on behalf of American Brachytherapy Society. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

Keywords:

Automated treatment planning; Artificial intelligence; Multi-objective optimization; Cervical cancer; Prostate cancer

Received 22 June 2023; received in revised form 26 September 2023; accepted 11 October 2023; Available online xxx

Disclosures: This work was supported by the Dutch Cancer Society [KWF Kankerbestrijding; project number 12183] and by Elekta Brachytherapy, Veenendaal, The Netherlands. The authors report no proprietary or commercial interest in any product mentioned or concept discussed in this article.

* Corresponding author. Evolutionary Intelligence Group, Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, The Netherlands.

E-mail address: renzo.scholman@cw.nl (R.J. Scholman).

^a These authors contributed equally.

1538-4721/\$ - see front matter © 2023 The Authors. Published by Elsevier Inc. on behalf of American Brachytherapy Society. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

<https://doi.org/10.1016/j.brachy.2023.10.005>

Please cite this article as: L.R.M. Dickhoff *et al.*, Keeping your best options open with AI-based treatment planning in prostate and cervix brachytherapy, Brachytherapy, <https://doi.org/10.1016/j.brachy.2023.10.005>

Introduction

In the last decade, most clinically used treatment planning methods in brachytherapy (BT) have been semi-automatic. Often, an inverse planning method - the most prevalent ones being inverse planning simulated annealing (IPSA) (1) and hybrid inverse treatment planning and optimization (HIPO) (2) - is used to generate a single treatment plan. This plan is then manually adjusted, by changing parameters in the inverse planning method, modifying specific dwell times, or dragging isodose lines graphically. With many dwell positions, ranging from roughly 60 up to over 400 for some patients, adjusting treatment plans by hand can be challenging, resulting in a time-consuming process and treatment plans that may still be improved (3–5).

Therefore, automated optimization for brachytherapy can play an important role, in quickly creating high-quality patient-specific treatment plans. However, in order to perform optimization appropriately, a clear definition is needed of what constitutes the ideal treatment plan. Recent studies, such as EMBRACE-I and EMBRACE-II for cervical cancer BT (6), have laid important groundwork, defining limits and desired values for planning criteria in terms of dose-volume indices (DVIs).

Plan quality is typically evaluated using the set of DVIs as given in a clinical protocol together with a visual inspection of the isodose lines. Nonetheless, other patient characteristics such as comorbidity and the radiation oncologists' preferences play a key role in determining what the most suitable treatment plan is for the patient at hand. These aspects are however difficult to quantify and therefore, not straightforward to include in an optimization model used for automated treatment planning.

Consequently, optimization models are inherently flawed with respect to describing the characteristics of a single, optimal treatment plan. Furthermore, most treatment planning methods solely optimize on quantitative metrics, in particular the DVIs, even though multiple different treatment plans (i.e., constituting different 3D dose distributions) can exist that exhibit the same metric values. We believe that it is important to unveil this truth and present radiation oncologists with insight into the possible options one has, even for plans that are optimized for criteria such as DVIs. Upon inspecting these options, it becomes clear what can be achieved for a particular patient, and the plan that is deemed most fitting can then be intuitively selected.

It has to be noted, nevertheless, that methods that generate only a single plan have successfully been developed. These include BiCycle (7,8) and intelligent inverse treatment planning via deep reinforcement learning (9). Both methods have been shown capable of outperforming previously used manual planning. Although it hides the fact that alternative plans are possible, the single-plan output can be preferred by radiation on-

cologists for reasons of efficiency, simplicity, or other preferences.

One recent method that is capable of generating multiple plans is GPU-based multicriteria optimization (gMCO) (10,11). This method however requires certain weights to be defined a priori for each organ (12). The most crucial region of interest to spare will not be similar for all patients, however. Therefore, to be able to obtain good results for all patients it is highly likely that these weights need to be set for each patient individually (or for subgroups of patients). Alternatively, many plans using many weight combinations can be computed first, and a medical physicist can attempt to a posteriori find the weights that lead to a preferred plan, which is how gMCO is proposed to be used in practice. However, the number of control parameters is the number of weights to be set and the number of potential plans to inspect is usually very large. Moreover, as the optimization model concerns mainly DVIs, it is not possible to discern different dose distributions with similar DVI values.

Most methods, including BiCycle and gMCO, change and simplify the problem formulation of optimizing BT treatment plans by making the problem smooth and convex so that gradient-based methods can be used. This means that these methods no longer optimize directly on the DVIs calculated according to the AAPM TG-43 dose calculation protocol (13). Evolutionary algorithms (EAs) are artificial intelligence (AI) based optimization techniques that do not rely on gradients, making it possible to optimize DVIs directly. Moreover, EAs can do so in a multi-objective (MO) fashion, finding multiple plans that trade-off key objectives (14,15).

An MO problem formulation together with a state-of-the-art EA has been introduced as BRIGHT (BRachytherapy via artificially Intelligent GOMEA-Heuristic based Treatment planning) (16). Each clinical goal in a clinical protocol is incorporated as-is in the optimization model in BRIGHT. By leveraging the GPU-parallelizability of BRIGHT, optimization on high-fidelity representations of dose distributions (100,000 dose calculation points) can be performed in within just 3 min (17). This then results in multiple plans that trade-off coverage and sparing, which has been proven to be insightful and can lead to desirable treatment plans (18).

While the first version of BRIGHT was deemed successful, challenges remain in the quest for the ideal treatment planning system that offers both high-quality plans and the most insight into what is possible for each patient. In this work, we highlight three key challenges in automated BT treatment planning and show how BRIGHT can be used to overcome these challenges. First, since different distinct 3D dose distributions exist for treatment plans which have similar DVI values, we demonstrate the need to find and present these options. Second, we outline the importance of being able to incorporate additional optimization criteria. Moreover, these criteria should be easy to

extend and be kept separate from the criteria that are commonly accepted to be important, such as the EMBRACE-II criteria for cervix brachytherapy. Lastly, we point out the importance of being able to quickly re-optimize treatment plans if the overall aims need to be altered (i.e., deviate from the protocol) for a specific patient.

Methods and materials

BRIGHT

As there is an inherent trade-off between irradiating the tumor target enough and sparing the organs at risk (OARs) more, a bi-objective optimization problem formulation that naturally reflects this has been introduced (19,20). This problem formulation is based directly upon the DVIs calculated according to the TG-43 formalism (13), which contributes to making it a nonconvex, non-linear, and nonsmooth optimization problem (14). EAs are especially suitable for the optimization of such problems. However, EAs can be relatively slow. For this reason, BRIGHT uses a modern model-based EA, called the Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA) (21). MO-RV-GOMEA can leverage the vast parallel computing power of modern GPUs (17). The use of MO-RV-GOMEA to solve the bi-objective problem formulation tailored to high-dose-rate (HDR) brachytherapy is called the BRachytherapy via artificially Intelligent GOMEA-Heuristic based Treatment planning (BRIGHT) method. The resulting output is a set of multiple plans varying from high coverage to high sparing, consisting of up to several hundred plans. However, since different radiation oncologists might prefer to be presented with more or less options, the number of presented plans is a user-defined setting. BRIGHT was shown to outperform both IPSA and HIPO, even when parameters in both methods were automatically tuned (22). The development of BRIGHT is done in close collaboration with a team of radiation oncologists, medical physicists, and radiation technologists, from two different treatment centers. Plans resulting from BRIGHT, including those from techniques presented in this paper, are regularly presented to them in feedback loops to ensure clinical acceptability. Following a successful retrospective evaluation study (18), BRIGHT has been clinically introduced at the Amsterdam University Medical Centers (Amsterdam UMC, location Academic Medical Center) for prostate cancer patients (16).

Optimization objectives

The DVIs in a clinical protocol can be divided into sparing and coverage aims, depending on whether the associated DVI should be minimized (“<”) or maximized (“>”). For cervical cancer BT, the aims follow from the EMBRACE-II protocol, except for the A point aims (23).

We do not include these, since this can lead to unnaturally forced dose distributions in that specific point only, instead of the whole parametrial area that point A clinically represents. For prostate cancer, the protocol from the Amsterdam UMC is used. All DVIs can be found in Table 1.

A detailed technical description of the two objective functions can be found in Supplementary Material A. In short, the objectives - called the least coverage index (LCI) and the least sparing index (LSI) - for a plan p are computed as the sum over either every coverage or sparing aim a as:

$$\begin{aligned} \text{LCI}_w(p) &= \sum_{a \in \text{coverage aims}} w_a(\delta(\text{DVI}^a)) \\ \text{LSI}_w(p) &= \sum_{a \in \text{sparing aims}} w_a(\delta(\text{DVI}^a)) \end{aligned} \quad (1)$$

where $\delta(\text{DVI}^a)$ is the difference between the current DVI value and its aim. Originally, LCI and LSI represented only the value of the worst achieved aim, that is, the largest δ (19). The weighted version above was introduced later so that, while the most violated DVI will still be given by far the most attention, other DVIs will still be improved in case the most violated DVI cannot be improved further (17). To this end, exponentially higher weights w_a are attributed to larger $\delta(\text{DVI}^a)$ values. This also implies that optimization is continued even when all planning aims are reached to obtain the best possible trade-off between coverage and sparing even beyond the aims.

Data

Data used for this work includes both cervical cancer as well as prostate cancer patients. All experiments showcased in this paper are done retrospectively.

The cervical cancer dataset includes 36 patients, treated with four fractions of 7 Gy HDR BT each at the Leiden University Medical Center (LUMC). This encompasses a total of 81 treatment plans with mostly MRI-based treatment planning (since most patients are treated with the same treatment plan in 2 consecutive fractions after CT-based verification of the patients’ anatomy relative to the applicator implant). The patients were treated during 2017–2020 after the EMBRACE-II protocol was introduced. Only patients for which all regions of interest as stated in this protocol were delineated and were treated with four fractions, were included.

The clinically used plans were manually optimized, in Oncentra Brachy (Elekta, Veenendaal, The Netherlands), by starting from a so-called library plan in which only applicator dwell positions are used, and which is normalized to a slight adjustment of the A points (closer to the applicator). Then, the plan is modified by adjusting (blocks of) selected dwell times and by manual optimization through drag-and-drop of isodose lines.

The prostate cancer dataset comprises 12 HDR BT patients, which were treated with a single fraction of 15 Gy.

Table 1

HDR BT aims for prostate and cervical cancer. Prostate aims are based on the protocol used by the Amsterdam UMC, given as a percentage of the single planning-aim dose of 15 Gy. Cervix aims are based on the EMBRACE-II protocol, based on a single fraction in percentages of 7 Gy (limits are provided between brackets).

	Volume	Use	Coverage criteria			Sparing criteria			Added aims		
Prostate	Prostate	Target	$V_{100\%}$	> 95%		$D_{90\%}$	> 15 Gy	$V_{150\%}$	< 40%	$V_{200\%}$	< 15%
	Vesicles	Target	V_{11Gy}	> 95%							
	Bladder	OAR						D_{1cm^3}	< 13 Gy	D_{2cm^3}	< 12 Gy
	Rectum	OAR						D_{1cm^3}	< 11 Gy	D_{2cm^3}	< 9.5 Gy
	Urethra	OAR						$D_{30\%}$	< 16.5 Gy	$D_{0.1cm^3}$	< 18 Gy
Cervix	CTV _{HR}	Target	$D_{90\%}$	> 7.8 Gy	(> 7.1 Gy)	$D_{98\%}$	> 5.8 Gy	$D_{90\%}$	< 8.3 Gy	$V_{100\%}$	> 99.9%
	CTV _{IR}	Target	$D_{98\%}$	> 3.5 Gy						$V_{50\%}$	> 99.9%
	GTV _{RES}	Target	$D_{98\%}$	> 8.3 Gy	(> 7.8 Gy)						
	Bladder	OAR						D_{2cm^3}	< 5.5 Gy	(< 6.3 Gy)	
	Rectum	OAR						D_{2cm^3}	< 4.0 Gy	(< 6.3 Gy)	
	Sigmoid	OAR						D_{2cm^3}	< 4.5 Gy	(< 6.3 Gy)	
	Bowel	OAR						D_{2cm^3}	< 4.5 Gy	(< 6.3 Gy)	
	Recto-vaginal point	OAR						D_{point}	< 4.0 Gy	(< 6.3 Gy)	
	Mid-CTV _{IR}	Target								$V_{100\%}$	< 25%
	Mid-normal-tissue	OAR								$V_{100\%}$	< 0.1%
Top-normal-tissue	OAR								$V_{100\%}$	< 0.2%	

Aims: D_v : dose index - minimum dose to the most irradiated subvolume v cm³; V_d : volume index - subvolume which is planned to receive at least dose d Gy; D_{point} : dose at point.

Abbreviations: CTV_{HR}=high risk clinical target volume; CTV_{IR}=intermediate risk clinical target volume; GTV_{RES}=residual gross tumor volume; see Supplementary Material Figure 1 for definitions of Mid-CTV_{IR}, Mid-normal-tissue, Top-normal tissue.

The clinically used plans were optimized using BRIGHT in 2020-2021 at the Amsterdam UMC (16). The plans which were used for clinical treatment are not further utilized in this work.

Adding custom DVIs for the cervix

For cervical cancer BT, treatment plans optimized solely on the planning aims from the EMBRACE-II protocol were deemed not clinically acceptable by radiation oncologists (24). Consequently, additional aims were formulated (20). The necessary regions of interest and corresponding aims were tuned through an iterative feedback loop in close collaboration with radiation oncologists, a BT medical physicist, and a BT technologist. The added aims are shown in Table 1, and definitions of the corresponding regions of interest are visualized in Supplementary Material Figure 1 (20).

A key question now is how to incorporate these aims intuitively, separating adherence to protocol from achieving additional, hospital-specific aims. In BRIGHT, this can be achieved relatively straightforwardly, by adding a third objective. Thus, using a similar form as Eq. (1), for cervical cancer BT, we define as a third objective the least added index (LAI):

$$\text{LAI}_w(p) = \sum_{a \in \text{added aims}} w_a(\delta(\text{DVI}^a)) \quad (2)$$

We present preliminary results from using tri-objective BRIGHT by means of a retrospective numerical comparison with the clinically used plans for the cervical cancer dataset as described in the last subsection. The BRIGHT plans were obtained by running BRIGHT for 3 min.

Radiation oncologists might opt for a higher coverage in one of the fractions due to a favorable applicator/needle implantation. This could be counteracted in a later fraction, as EMBRACE-II describes a total EQD2 dose over all fractions. Therefore, we compare the cumulative dose per patient, over all four fractions combined (where two different implantations at fraction 1 and 3 and possible plan adjustments for fractions 2 and 4 are taken into account). As such, the clinical plans were evaluated using the true planned dose based on (potentially) four different plans. BRIGHT plans have also been optimized whenever a new clinical plan was constructed.

The rationale is that in clinical practice, the radiation oncologist will be able to inspect the set of high-quality plans generated by BRIGHT, each plan representing different trade-offs between the objectives, and thereby enabling making a well-informed decision regarding the preferred treatment plan for the patient at hand. Because we do not perform a clinical validation study in this work, but still want to be able to compare the clinical plans for a patient with BRIGHT plans, we selected a single BRIGHT plan for each fraction of the patient. Specifically, we chose plan p for which the LCI and LSI are most balanced in

the sense that the worst of them is maximized, that is,

$$p = \operatorname{argmax}_{p \in \text{plans}} [\min(\text{LCI}(p), \text{LSI}(p))] \quad (3)$$

Picking BRIGHT plans this way is solely meant to give an indication of whether BRIGHT plans, accumulated over all fractions, could achieve better coverage and sparing compared to the clinical plans. BRIGHT plans are taken as resulting directly from the algorithm, without further manual optimization.

We note that BRIGHT is a stochastic approach that results in slightly different treatment plans after each optimization run (unless the random seed is fixed). In practice, BRIGHT would be run only once. However, in a research setting, it is prudent to run BRIGHT 30 times for each fraction. The median result overall runs is then taken for further analysis as it is a fair expectation of what would be achieved in clinical practice. Note that differences between runs are very small and therefore most likely not clinically significant; this is verified by calculating the interquartile ranges.

A comparison is made by examining the number of times for which all the aims and limits of the EMBRACE-II protocol are reached for a patient, since abiding by this internationally recommended protocol is important. As the added aims are included in the third objective, resulting plans exhibit corresponding desirable properties up to the maximum degree. These aims were however not analyzed further in our evaluation as the aims and limits from the EMBRACE-II protocol are internationally accepted whereas the added aims include local preferences.

Finally, we evaluate whether BRIGHT plans or clinical plans have overall better coverage and sparing for the limits in the EMBRACE-II protocol over all four fractions. Here, better is defined as the LCI and LSI simultaneously being at least 1% higher.

We further perform a two-sided Wilcoxon signed rank test with significance level $\alpha = 0.05$ to investigate whether results differ significantly. This is done separately for levels of sparing (of OARs) and coverage (of targets), on a per-patient basis over all given fractions.

Note that a retrospective observer study for prostate cancer BT has already been conducted (18), which is why we solely present experiments for cervical cancer BT in this section.

Distinct dose distributions with the same DVI values

Treatment plans which have the same DVI values can still be characterized by different dose distributions. These alternatives cannot be found if just a single plan is optimized based solely on DVIs. With BRIGHT, however, it is possible to optimize for multiple sets of treatment plans that each exhibit coverage-sparing trade-offs, while each set is focused on a distinctly different shape of the dose distribution.

To achieve this, we run BRIGHT multiple times and for each run t , we take a plan p_t from the resulting set of solutions. For this, we choose again the most balanced plan as outlined in Eq. (3). We define a fourth objective that maximizes the distance to these plans in terms of dose distribution. We denote the set of all previously selected plans p_0, p_1, \dots, p_{t-1} by \mathbb{P}_t .

For run t , the distance between plan p and set \mathbb{P}_t is defined as the minimal Euclidean distance to any plan in \mathbb{P}_t . Plan distances are calculated in the space of dose calculation points d_i^p , where d_i^p is the i -th dose calculation point for plan p , i.e.:

$$\text{distance}(p, \mathbb{P}_t) = \min_{p_t \in \mathbb{P}_t} \left(\sqrt{\sum_{i=0}^{n-1} (d_i^{p_t} - d_i^p)^2} \right) \quad (4)$$

In the first run of BRIGHT, the normal problem formulation is used. In four subsequent runs, the new fourth objective is also taken into consideration. It is worth noting that multiple runs are only necessary here to show that different dose distributions are possible at all - in future work, we incorporate optimizing for different dose distributions in one single run. We thereby find sets of treatment plans that differ from each other in that the values in their respective dose calculation points are as different as possible.

This extension is retrospectively tested on all 36 cervical cancer patients and 12 prostate cancer patients as described above.

Reoptimization

Occasionally, patients can have a challenging organ-tumor geometry, an unfavorable applicator/needle implantation, or other specific characteristics that require deviating from the standard clinical protocol in a specific way. These desirable deviations can be too large to be represented by one of the trade-off plans that results from using BRIGHT on the standard clinical protocol. It can be the case that one only sees this after optimization is performed. In such situations it is important to be able to actively steer the optimization to these desirable deviating aims and thus to specify how much to deviate from the standard clinical protocol.

Should the need to deviate from the protocol be discovered only after optimization has been performed already, reoptimization is required. Instead of restarting from scratch, BRIGHT can be configured to continue optimization using plans obtained so far as the starting point. This results in a substantial speed-up compared to starting from scratch, requiring only 30 s instead of 3 min to achieve similar outcomes (25).

As seen in clinical practice for BRIGHT, manual adjustments were sometimes needed in the generated plans for prostate cancer (16). Therefore, a technique for reoptimization is desired. We present first results of running

Table 2

Statistics for cervical treatment plans from BRIGHT and clinical plans in terms of reaching planning limits and aims of the EMBRACE-II protocol (percentage of patients indicated for which these are achieved).

	Planning limits	Planning aims
BRIGHT	100.0%	83.6%
Clinic	88.9%	47.2%

the re-optimization once for a patient as a proof of concept. For illustration purposes, it is tested by adjusting the bladder $D_{2\text{cm}^3}$ and $D_{1\text{cm}^3}$ to stricter values on 36 patients for cervical cancer and 12 patients for prostate cancer BT. The effects thereof are displayed via dose distributions.

Results

Adding custom DVIs for the cervix

The BRIGHT results presented in Table 2 indicate that the inclusion of the added aims does not prevent the standard planning aims and limits from being met. Clinical choices often included planning a slightly higher dose at the cost of some sparing due to which not all limits could be achieved. Often, more dose was planned for the CTV_{HR} to ensure that the CTV_{IR} would not be underdosed. Importantly, all OAR limits were always satisfied. Only the limit for the recto-vaginal point was violated for two patients.

Results for achieved LCI and LSI values (medians, interquartile ranges, minima, and maxima) for both BRIGHT and clinical plans, their median differences for each individual DVI, as well as a per-fraction comparison can be found in the Supplementary Material C and D.

For 41.7% of the patients, both better overall coverage and sparing is achieved by BRIGHT compared to the clinical plans. For all other patients, equal levels of coverage and sparing were achieved. The two-sided Wilcoxon signed rank test ($\alpha=0.05$) leads to a p -value of 0.12 for coverage (LCI) and 1.08×10^{-7} for sparing (LSI). Thus, BRIGHT plans are significantly better than the clinically used plans in terms of sparing, and perform equally in terms of coverage. These results further emphasize that BRIGHT can generate plans with at least equally favorable DVIs for the clinical protocol, whilst also shaping the dose distribution according to additional aims. A visualization thereof is provided in Fig. 1, which shows the difference in accumulated LCI and LSI over four fractions between the picked BRIGHT plans and the clinical plans. As can be seen, the patients for which equal performance is found between BRIGHT plans and clinically used plans, are characterized by a high level of sparing and/or coverage, though not necessarily both. This could suggest that manual optimization was a complex procedure. It is however more likely that additional patient-specific information was incorporated into the decision to go for high coverage or high sparing plans. This underlines the need to not only present one automatically optimized plan, but a set of

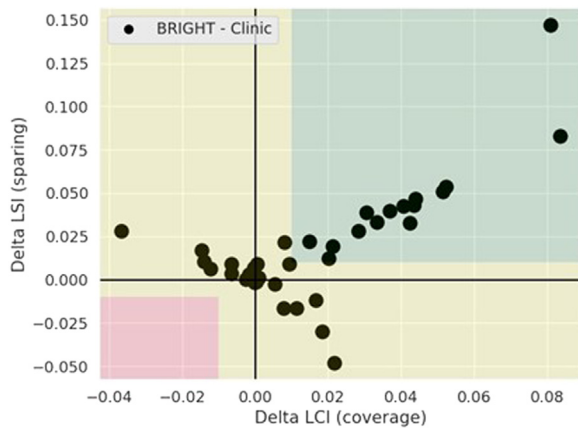


Fig. 1. Comparison of cervical brachytherapy BRIGHT plans and the clinically used plans, for each patient, accumulated over all fractions for that patient. Values for the black dots are calculated by, first, for each fraction, selecting the plan according to 3 from the set of plans generated by BRIGHT. Then, the selected four plans are added together to evaluate the cumulative dose given over four fractions. Subsequently, for both LCI and LSI, per patient, the total value associated with the clinical plans is subtracted from the total value associated with the BRIGHT plans (both calculated using 1). This way, dots in the green quadrant denote that for those patients, better LCI as well as LSI (by at least 1%) was achieved using BRIGHT as compared to the clinic. Dots in the upper-left quadrant indicate that BRIGHT achieved a higher LSI but lower LCI value, whereas dots in the lower-right quadrant imply a higher LCI but lower LSI value. Dots in the red quadrant would mean that the clinic performed better than BRIGHT. (For color figure the reader is referred to the web version of this article.)

plans representing the full spectrum of different trade-offs between coverage and sparing.

Benefits of being presented with options are further emphasized by the fact that in this work, we solely included

cervical cancer patients who were treated with four fractions. Clinically, depending on the target volume, the implant, and the favorability of the position of the OARs, some patients are treated with three fractions to reduce hospitalization time and patient discomfort. However, to achieve the same target dose with three fractions instead of four, plans with much higher coverage are required. Arguably, therefore, it would be favorable to present these high coverage plans, even if some sparing aims would be achieved less. This would aid in identifying those patients for whom a more favorable coverage/sparing ratio can be achieved using three fractions.

It has to be noted, however, that this comparison exclusively considers the objective values and neglects the dose distributions which play an important role in defining the quality of a treatment plan. A full blinded observer study still has to be conducted, but discussions with radiation oncologists that occurred in the process of tuning the added aims, have nonetheless already indicated that BRIGHT plans are clinically acceptable.

Distinct dose distributions with the same DVI values

We show two examples of optimized treatment plans with similar DVI values, yet different dose distributions, for both cervix and prostate BT, respectively in Figs. 2 and 3. For reasons of legibility, results are showcased only on one cervical and one prostate cancer patient, which were found to be representative of all other patients. All aims from the respective protocols were achieved for the visualized treatment plans.

In Fig. 2, two plans for cervical cancer are illustrated. The corresponding DVI values are given in Table 3. As can be seen in the coronal view along the applicator axis,

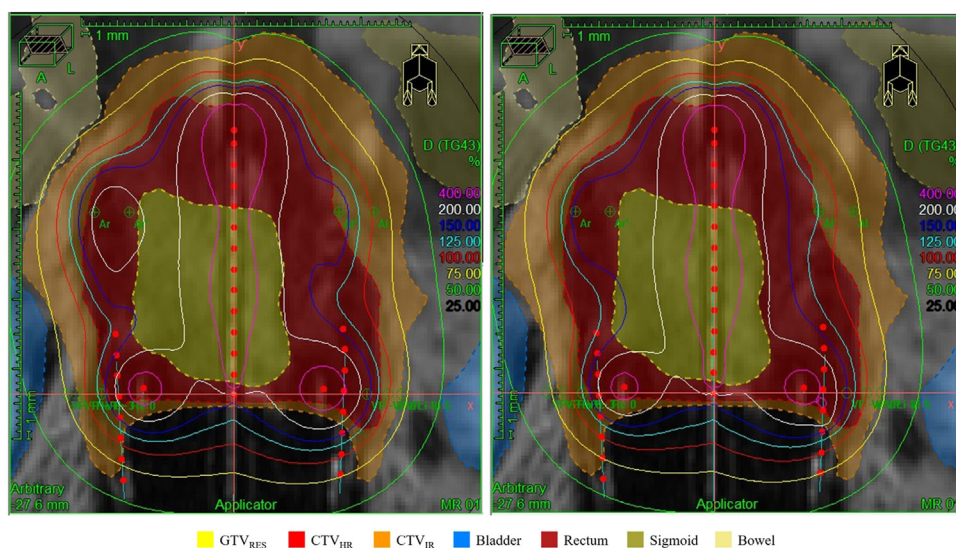


Fig. 2. Distinct dose distributions with similar DVI values for cervix HDR BT. The 100% isodose line equals the prescribed dose of 7 Gy. Visualized with Oncentra Brachy (Elekta, Veenendaal, The Netherlands). (For color figure the reader is referred to the web version of this article.)

Table 3

DVI values for plans with distinct dose distributions from Fig. 2.

DVI	Left plan	Right plan
GTV _{RES} $D_{98\%}$	9.9 Gy	9.9 Gy
CTV _{HR} $D_{90\%}$	8.1 Gy	8.1 Gy
CTV _{HR} $D_{98\%}$	7.0 Gy	7.0 Gy
CTV _{IR} $D_{98\%}$	4.2 Gy	4.2 Gy
Bladder $D_{2\text{cm}^3}$	5.1 Gy	5.1 Gy
Rectum $D_{2\text{cm}^3}$	2.2 Gy	2.2 Gy
Sigmoid $D_{2\text{cm}^3}$	4.2 Gy	4.2 Gy
Bowel $D_{2\text{cm}^3}$	3.6 Gy	3.6 Gy
Recto-vaginal point	3.3 Gy	3.4 Gy

different dwell positions in different needles have been loaded. Despite these differences, the DVI values differ by at most a couple of cGy. For prostate cancer, similar results are found (see Fig. 3 and Table 4). Rather large differences occur in used dwell positions and thus in the 3D dose distribution. Analogous to the cervix case, the differences in DVI values are again negligible.

Reoptimization

The capability of BRIGHT to re-optimize treatment plans after a (small) protocol adaptation is shown for two cases, one for cervix and one for prostate BT. These two cases are found to be representative of results found for other patients.

Figure 4 shows the transverse view from a patient with cervical cancer, where the left dose distribution corresponds to a tri-objective BRIGHT treatment plan obtained using the clinical protocol and the added aims. Then, the protocol was modified such that, for the bladder, instead of the initial $D_{2\text{cm}^3} < 5.5$ Gy, the aim was set to $D_{2\text{cm}^3} < 4.8$ Gy, that is, more sparing for the bladder is needed. The dose distribution on the right displays the reoptimized plan, generated in under 30 s. It is characterized by a shift of the dose distribution, especially visible when focusing

Table 4

DVI values for plans with distinct dose distributions from Fig. 3.

DVI	Left plan	Right plan
Prostate $D_{90\%}$	16.3 Gy	16.3 Gy
Prostate $V_{100\%}$	95.5%	95.4%
Prostate $V_{150\%}$	38.5%	39.5%
Prostate $V_{200\%}$	15.2%	15.7%
Vesicles $V_{11\text{Gy}}$	99.8%	100.0%
Bladder $D_{1\text{cm}^3}$	12.5 Gy	12.5 Gy
Bladder $D_{2\text{cm}^3}$	11.0 Gy	11.0 Gy
Rectum $D_{1\text{cm}^3}$	10.0 Gy	9.8 Gy
Rectum $D_{2\text{cm}^3}$	8.4 Gy	8.4 Gy
Urethra $D_{0.1\text{cm}^3}$	17.6 Gy	17.6 Gy
Urethra $D_{30\%}$	16.0 Gy	16.0 Gy

on the yellow 75% isodose line. The bladder now receives less dose. The plan is assuredly evaluated using the aims from the original protocol, which have still all been met.

Similarly, Fig. 5 shows the dose distributions for a prostate cancer patient. The treatment plan on the left is optimized on the regular prostate protocol with the bladder aims being $D_{1\text{cm}^3} < 13$ Gy and $D_{2\text{cm}^3} < 12$ Gy. For the reoptimized treatment plan on the right, the aims for the bladder have been reduced to $D_{1\text{cm}^3} < 12$ Gy and $D_{2\text{cm}^3} < 10.5$ Gy. The reoptimized treatment plan shows that the dose distribution is now shifted more towards the outside of the bladder, most notably resulting in the cyan 100% isodose line being located closer to the prostate. The corresponding DVI values for the bladder have been reduced by 0.5 Gy whilst still achieving all of the other aims, except for a very slight violation of the prostate $V_{150\%}$ and $V_{200\%}$.

Discussion and conclusion

To this day, treatment plan evaluation consists of inspection of the DVI values and visual inspection of the associated 3D dose distribution. Clinical practice is guided by protocols in which desirable values for the DVIs are given.




Fig. 3. Distinct dose distributions with similar DVI values for prostate HDR BT. The 100% isodose line equals the prescribed dose of 15 Gy. Visualized with Oncentra Brachy (Elekta, Veenendaal, The Netherlands). (For color figure the reader is referred to the web version of this article.)

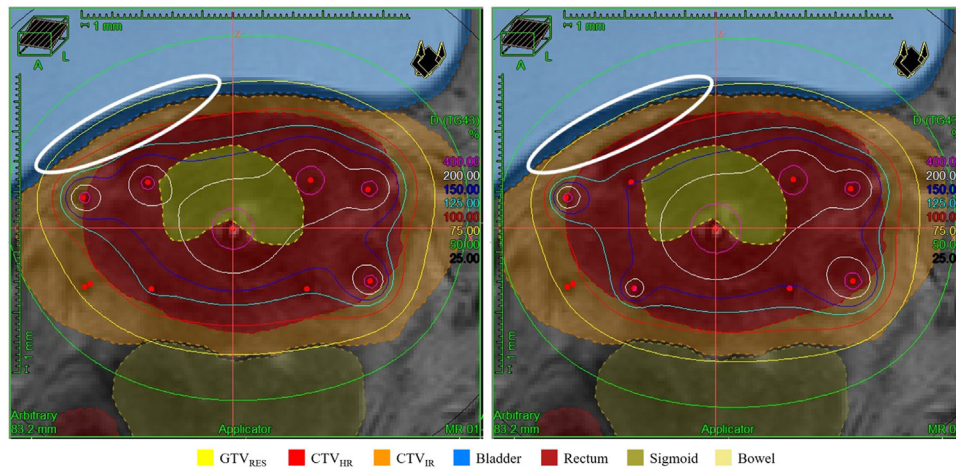


Fig. 4. Left: initial dose distribution of a treatment plan for cervix HDR BT optimized by BRIGHT, resulting in $D_{2\text{cm}^3} = 5.5$ Gy for the bladder. Right: reoptimized treatment plan with a lower aim for the bladder, obtaining $D_{2\text{cm}^3} = 5.2$ Gy. Visualized with Oncentra Brachy (Elekta, Veenendaal, The Netherlands). (For color figure the reader is referred to the web version of this article.)




Fig. 5. Left: initial dose distribution of a by BRIGHT optimized treatment plan for prostate HDR BT optimized with the regular aims, resulting in $D_{1\text{cm}^3} = 12.3$ Gy and $D_{2\text{cm}^3} = 10.8$ Gy for the bladder. Right: reoptimized treatment plan with lower aims for the bladder, leading to $D_{1\text{cm}^3} = 11.7$ Gy and $D_{2\text{cm}^3} = 10.3$ Gy. Visualized with Oncentra Brachy (Elekta, Veenendaal, The Netherlands). (For color figure the reader is referred to the web version of this article.)

Although modern optimizers (of which relevant literature is presented in the introduction as a motivational point) are capable of optimizing plans using the DVIs, when using DVIs only, not all clinical wishes can be optimized. This makes it impossible to calibrate models and algorithms so that they can generate for each individual patient a single preferred treatment plan, making overcoming this a key challenge.

Overcoming this includes developing criteria that can be optimized, to more closely describe plan properties that are desirable or undesirable. For instance, by including the use of contiguous volumes, hotspots can be detected and minimized (26), or specific shapes of the dose distribution can be optimized. As an example, in cervical cancer BT, since the introduction of the A points, favorable dose distributions were historically described to resemble a pear shape (27,28). Incorporating contiguous volumes requires

an AI technique which can handle nonconvex optimization, for which EAs are particularly suited.

Even so, it is likely that for the foreseeable future, multiple different 3D dose distributions exist for plans which have similar values for the criteria that we optimize for. At the same time, the quality of a BT treatment plan is additionally defined by other non-quantifiable characteristics (e.g., comorbidity or the radiation oncologists' preferences) which cannot easily be included in the current optimization criteria. It is therefore important to find all plan alternatives and present these as options to the radiation oncologists so that they can pick the best plan according to (their expertise and) all plan characteristics. In this work, we demonstrate that it is possible to explicitly find plans with similar values for the optimization criteria but different underlying 3D dose distributions for both prostate and cervix BT.

Even if in the (near) future a more detailed definition of what entails a good treatment plan would be known, in clinical practice different patients can have particular differences in anatomy (i.e., OARs position relative to the target), needing distinct implants, and therefore leading to individual problems and needs. Thus, approaches with which it is possible to quickly deviate from a standard protocol are desirable to obtain a tailored high-quality treatment plan for the specific patient at hand.

In this work, we focus on BRIGHT, an AI-based treatment planning method for BT. The optimization algorithm in BRIGHT is sufficiently flexible to support the required functionality outlined above. Bringing this all together and extending this with an intuitive user interface that allows insightful and fast plan selection and comparison, is a near-future vision that has the potential to innovate BT treatment planning while keeping options open to discover what it is we wish to see in optimal treatment plans, rather than become overly focused on single plans that are deemed optimal by DVI-based optimization models.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at [doi:10.1016/j.brachy.2023.10.005](https://doi.org/10.1016/j.brachy.2023.10.005).

References

- Lessard E, Pouliot J. Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Med Phys* 2001;28:773–779. doi:10.1118/1.1368127.
- Karabis A, Giannouli S, Baltas D. HIPO: a hybrid inverse treatment planning optimization algorithm in HDR brachytherapy. *Radiother Oncol* 2005;76:S29. doi:10.1016/S0167-8140(05)81018-7.
- Lessard E, Hsu IC, Pouliot J. Inverse planning for interstitial gynecologic template brachytherapy: truly anatomy-based planning. *Int J Radiat Oncol - Biol - Phys* 2002;54(4):1243–1251. doi:10.1016/S0360-3016(02)03802-6.
- Morton GC, Sankrecha R, Halina P, Loblaw A. A comparison of anatomy-based inverse planning with simulated annealing and graphical optimization for high-dose-rate prostate brachytherapy. *Brachytherapy* 2008;7(1):12–16. doi:10.1016/j.brachy.2007.10.001.
- Jamema SV, Sharma S, Mahantshetty U, et al. Comparison of IPSA with dose-point optimization and manual optimization for interstitial template brachytherapy for gynecologic cancers. *Brachytherapy* 2011;10(4):306–312. doi:10.1016/j.brachy.2010.08.011.
- EMBRACE. Accessed: May 25, 2023. URL: www.embracestudy.dk.
- Bijman R, Rossi L, Breedveld S, et al. Fully automated adaptive multi-criteria planning for EMBRACE II cervical cancer brachytherapy. *Radiother Oncol* 2022;170:S390 ESTRO 2022, 6–10 May 2022, Copenhagen. Onsite in Copenhagen and Online. doi:10.1016/S0167-8140(22)02582-8.
- Breedveld S, Bannan ABA, Aluwini S, et al. Fast automated multi-criteria planning for HDR brachytherapy explored for prostate cancer. *Phys Med Biol* 2019;64(20):205002. doi:10.1088/1361-6560/ab44ff.
- Shen C, Gonzalez Y, Klages P, et al. Intelligent inverse treatment planning via deep reinforcement learning, a proof-of-principle study in high dose-rate brachytherapy for cervical cancer. *Phys Med Biol* 2019;64:S29. doi:10.1088/1361-6560/ab18bf.
- Bélanger C, Cui S, Ma Y, et al. A GPU-based multi-criteria optimization algorithm for HDR brachytherapy. *Phys Med Biol* 2019;64:105005. doi:10.1088/1361-6560/ab1817.
- Bélanger C, Poulin E, Aubin S, et al. Inter-observer evaluation of a GPU-based multicriteria optimization algorithm combined with plan navigation tools for HDR brachytherapy. *Brachytherapy* 2022;21(4):551–560. doi:10.1016/j.brachy.2022.04.003.
- Cui S, Després P, Beaulieu L. A multi-criteria optimization approach for HDR prostate brachytherapy: I. Pareto surface approximation. *Phys Med Biol* 2018;63(20):205004. doi:10.1088/1361-6560/aae24c.
- Rivard MJ, Coursey BM, DeWerd LA, et al. Update of AAPM task group no. 43 report: A revised AAPM protocol for brachytherapy dose calculations. *Med Phys* 2004;31(2):633–674. doi:10.1118/1.1646040.
- Lahanas M, Baltas D, Zamboglou N. A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy. *Phys Med Biol* 2003;48(3):399. doi:10.1088/0031-9155/48/3/309.
- Milickovic NB, Lahanas M, Papagiannopoulou M, et al. Application of multiobjective genetic algorithms in anatomy based dose optimization in brachytherapy and its comparison with deterministic algorithms. In: 2001 Conference Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society. 2001; vol 4, pp. 3919–3922. pages. doi:10.1109/IEMBS.2001.1019698.
- Barten DLJ, Bouter A, van Wieringen N, et al. Artificial intelligence based planning of HDR prostate brachytherapy: first clinical experience. *Radiother Oncol* 2021;161:653–655. doi:10.1016/S0167-8140(21)07100-0.
- Bouter A, Alderliesten T, Pieters BR, et al. GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy. *Med Phys* 2019;46:3776–3787. doi:10.1002/mp.13681.
- Maree SC, Luong NH, Kooreman ES, et al. Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy—a retrospective observer study. *Brachytherapy* 2019;18:396–403. doi:10.1016/J.BRACHY.2018.12.010.
- Luong NH, Bouter A, van der Meer MC, et al. Efficient, effective, and insightful tackling of the high-dose-rate brachytherapy treatment planning problem for prostate cancer using evolutionary multi-objective optimization algorithms. In: Proceedings of the 2017 Genetic and Evolutionary Computation Conference Companion, Berlin, Germany. 2017, pp. 1372–1379. doi:10.1145/3067695.3082491.
- Dickhoff LRM, Kerkhof EM, Deuzeman H, et al. Adaptive objective configuration in bi-objective evolutionary optimization for cervical cancer brachytherapy treatment planning. In: Proceedings of the 2022 Genetic and Evolutionary Computation Conference, GECCO '22, New York, NY, USA. 2022. Association for Computing Machinery, pp. 1173–1181. doi:10.1145/3512290.3528822.
- Bouter A, Luong NH, Witteveen C, et al. The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In: Proceedings of the Genetic and Evolutionary Computation Conference GECCO '17. 2017. New York, NY, USA: Association for Computing Machinery, pp. 537–544. doi:10.1145/3071178.3071274.
- Maree SC, Bosman PAN, Van Wieringen N, et al. Automatic bi-objective parameter tuning for inverse planning of high-dose-rate prostate brachytherapy. *Phys Med Biol* 2020;65. doi:10.1088/1361-6560/ab7362.
- Srivastava A, Datta NR. Brachytherapy in cancer cervix: time to move ahead from point A? *World J Clin Oncol* 2014;5:764–774. doi:10.5306/wjco.v5.i4.764.
- Dickhoff LRM, Kerkhof EM, Pieters BR, et al. Automated optimization for cervix brachytherapy requires more than the EMBRACE-II planning aims. *Radiother Oncol* 2022;170:S388–S390.
- Bakker M. Warm-starting evolutionary plan optimization for high-dose-rate brachytherapy treatment to reduce optimization time Master's thesis. TU Delft Electrical Engineering, Mathematics and Computer Science, 2022.

- [26] Commandeur JL, Bouter A, Dickhoff LRM, et al. Incorporating control of contiguous highdose volumes in automated optimization for prostate BT. *Radiother Oncol* 2022;170:S232–S233.
- [27] Fletcher GH, Wall JA, Bloedorn FG, et al. Direct measurements and isodose calculations in radium therapy of carcinoma of the cervix. *Radiology* 1953;61:885–902. doi:10.1148/61.6.885.
- [28] Tanderup K, Nielsen SK, Nyvang GB, et al. From point A to the sculpted pear: MR image guidance significantly improves tumour dose and sparing of organs at risk in brachytherapy of cervical cancer. *Radiother Oncol* 2010;94:173–180. doi:10.1016/j.radonc.2010.01.001.