



PAPER

Automatic bi-objective parameter tuning for inverse planning of high-dose-rate prostate brachytherapy

S C Maree¹, P A N Bosman², N van Wieringen¹, Y Niatsetski³, B R Pieters¹, A Bel¹ and T Alderliesten¹¹ Department of Radiation Oncology, Amsterdam UMC, University of Amsterdam, Amsterdam, The Netherlands² Life Sciences and Health Group, Centrum Wiskunde & Informatica, Amsterdam, The Netherlands³ Elekta, Veenendaal, The NetherlandsE-mail: s.c.maree@amsterdamumc.nl**Keywords:** brachytherapy, prostate cancer, treatment planning, parameter tuning, multi-objective optimization, IPSA, HIPORECEIVED
5 September 2019REVISED
29 January 2020ACCEPTED FOR PUBLICATION
6 February 2020PUBLISHED
2 April 2020**Abstract**

We present an automatic bi-objective parameter-tuning approach for inverse planning methods for high-dose-rate prostate brachytherapy, which aims to overcome the difficult and time-consuming manual parameter tuning that is currently required to obtain patient-specific high-quality treatment plans. We modelled treatment planning as a bi-objective optimization problem, in which dose-volume-based planning criteria related to target coverage are explicitly separated from organ-sparing criteria. When this model is optimized, a large set of high-quality plans with different trade-offs can be obtained. This set can be visualized as an insightful patient-specific trade-off curve. In our parameter-tuning approach, the parameters of inverse planning methods are automatically tuned, aimed to maximize the two objectives of the bi-objective planning model. By generating trade-off curves for different inverse planning methods, their maximally achievable plan quality can be insightfully compared. Automatic parameter tuning furthermore allows to construct standard parameter sets (class solutions) representing different trade-offs in a principled way, which can be directly used in current clinical practice. In this work, we considered the inverse planning methods IPSA and HIPO. Thirty-nine previously treated prostate cancer patients were included. We compared automatic parameter tuning, random parameter sampling, and the maximally achievable plan quality obtained by directly optimizing the bi-objective planning model with the state-of-the-art optimization software GOMEA. We showed that for each patient, a set of plans with a wide range of trade-offs could be obtained using automatic parameter tuning for both IPSA and HIPO. By tuning HIPO, better trade-offs were obtained than by tuning IPSA. For most patients, automatic tuning of HIPO resulted in plans close to the maximally achievable plan quality obtained by optimizing the bi-objective planning model directly. Automatic parameter tuning was shown to improve plan quality significantly compared to random parameter sampling. Finally, from the automatically-tuned plans, three class solutions were successfully constructed representing different trade-offs.

1. Introduction

High-dose-rate (HDR) brachytherapy is a treatment modality for prostate cancer with excellent therapeutic outcomes (Bachand *et al* 2009, Fatyga *et al* 2009). In HDR prostate brachytherapy, catheters are implanted in the prostate. Within each catheter, a radioactive source can be temporarily stopped at multiple dwell positions to locally irradiate the surrounding tissue. Treatment planning is the process of determining the time this source stops at each dwell position, which is essential for delivering an effective treatment that optimizes the trade-off between tumor control and normal tissue complications. The quality of a treatment plan is assessed by a visual inspection of the 3D dose distribution. In addition, and to quantify certain aspects, treatment planning criteria are often formulated in terms of dose-volume indices on the target volumes and organs at risk (OARs), derived from the dose-volume histogram (Hoskin *et al* 2013, RTOG 2016). The purpose of treatment planning is to obtain a plan that has desirable trade-offs between the planning criteria. Achievable trade-offs are however

patient specific, which makes treatment planning in clinical practice difficult and time consuming (Venselaar *et al* 2012, Maree *et al* 2019).

Inverse treatment planning methods aim to aid in the planning process by reducing planning time and effort (Dinkla *et al* 2015). In these methods, treatment planning is modeled as an optimization problem and optimization software is used to obtain high-quality plans according to the specified model. Two clinically widely-available planning methods are inverse planning simulated annealing (IPSA) (Lessard and Pouliot 2001) and hybrid inverse plan optimization (HIPO) (Lahanas *et al* 2003). In both IPSA and HIPO the treatment planning problem is modeled with a dose-penalty model. This model allows fast optimization, but resulting plans do not necessarily adhere to the planning criteria even when such plans exist (Holm *et al* 2012, Gorissen *et al* 2013). To overcome this, the planner needs to tune the parameters of the dose-penalty model and rerun the planning method, or adapt the plan manually using graphical optimization, either of which is difficult and time consuming (Dinkla *et al* 2015, Maree *et al* 2019).

Instead of formulating the treatment planning problem by using a dose-penalty model, it can also be approached by optimizing the dose-volume indices directly (Morton *et al* 2008, Siau *et al* 2011). One such approach uses a bi-objective planning model (Luong *et al* 2018a, Maree *et al* 2019). In the bi-objective planning model, the planning criteria in the clinical protocol are grouped into two separate objectives; one related to target coverage and one related to OAR sparing. Per objective, the grouped planning criteria are combined in a worst-case manner such that optimization searches for satisfactory plans, i.e. plans that satisfy all planning criteria. The result of optimizing the bi-objective planning model is not a single best plan, but a large set of plans that we call the *trade-off set*, in which each plan has a different high-quality trade-off between the two objectives. Since the bi-objective planning model reduces all planning criteria to only two objectives, the trade-off set is easy to visualize as a *trade-off curve*, which shows maximally achievable plan quality in an insightful way. The trade-off curve can be used in clinical practice to intuitively select the preferable plan from the trade-off set. It was shown in a retrospective observer study that optimizing the bi-objective planning model with the gene-pool optimal mixing evolutionary algorithm (GOMEA) (Bouter *et al* 2017) resulted in plans that were considered clinically acceptable for all patients considered ($n = 18$) (Maree *et al* 2019). Furthermore, these automatically generated plans were preferred over the clinically used plans in 98% of the cases, based on inspection of the 3D dose distribution and the achieved planning criteria (Maree *et al* 2019).

In this work, we developed an automatic bi-objective parameter-tuning approach that automatically tunes the parameters of an inverse planning method, aimed to maximize the two objectives of the bi-objective planning model. By specifying the parameter settings of an inverse planning method, a single treatment plan is generated. When the parameter settings are varied, a different plan can be generated. A large set of different parameter settings therefore results in a large set of different plans. However, not all parameter settings will result in high-quality plans. Our bi-objective parameter-tuning approach therefore automatically searches for multiple parameter settings aimed at obtaining a set of high-quality plans. By this approach, maximally achievable trade-off sets can be generated with any inverse planning method. Corresponding trade-offs insightfully show the entire range of achievable high-quality plans of an inverse planning method. By comparing trade-off curves for different inverse planning methods, their maximally achievable plan quality can be insightfully compared.

We applied our automatic bi-objective parameter-tuning approach to the inverse planning methods IPSA and HIPO. We investigated how automatic parameter tuning differs from the maximally achievable plan quality obtained by directly optimizing dwell times with GOMEA and from the clinically used treatment plans. We furthermore compared automatic bi-objective parameter tuning to random parameter sampling, i.e. a straightforward approach to generate a trade-off set, but without further optimizing the resulting plans (Bélanger *et al* 2019).

Our automatic parameter-tuning approach aims to overcome the need for manual tuning, but this approach is not yet available in clinical practice today. Therefore, we showed that trade-off sets generated by automatic bi-objective parameter tuning can be used to construct standard parameter sets (class solutions) representing different plan trade-offs in a principled way. These class solutions can directly be applied in current clinical practice.

2. Materials and methods

2.1. Patient data and treatment protocol

Between November 2014 and December 2018, 39 prostate cancer patients were treated in our medical center with a single-dose HDR brachytherapy of 13 Gy delivered a week after external beam radiation treatment with a dose schedule of 20×2.2 Gy. Catheter implantation was performed using transrectal ultrasound under general or epidural anesthesia according to a preplan, made in the operating room in Oncentra Prostate (version 4.2.2, Elekta AB, Stockholm, Sweden) based on ultrasound imaging (Pieters *et al* 2006). A transurethral catheter with a bladder balloon was used to enhance visibility of the urethra. After catheter implantation, three orthogonal pelvic T2-weighted turbo spin echo magnetic resonance imaging (MRI) scans (Ingenua 3.0T Philips Healthcare,

Table 1. HDR BT prostate protocol of a single planning-aim dose of 13 Gy. Volume (V) indices are expressed in percent of the planning-aim dose, and have unit volume percentage. Dose (D) indices are expressed in volume percentage or absolute volume (cm^3), and have unit Gy. Equivalent dose (EQD2) computed with an α/β -ratio = 3.

	Volume	Use			Criteria	
Coverage criteria	Prostate	Target	$V_{100\%} > 95\%$	42 Gy EQD _{2,3}	$D_{90\%} > 100\%$	42 Gy EQD _{2,3}
	Vesicles	Target	$V_{80\%} > 95\%$	28 Gy EQD _{2,3}		
Sparing criteria	Prostate	Organ	$V_{150\%} < 50\%$	89 Gy EQD _{2,3}	$V_{200\%} < 20\%$	150 Gy EQD _{2,3}
	Bladder	Organ	$D_{1\text{ cm}^3} < 86\%$	32 Gy EQD _{2,3}	$D_{2\text{ cm}^3} < 74\%$	24 Gy EQD _{2,3}
	Rectum	Organ	$D_{1\text{ cm}^3} < 78\%$	27 Gy EQD _{2,3}	$D_{2\text{ cm}^3} < 74\%$	24 Gy EQD _{2,3}
	Urethra	Organ	$D_{0,1\text{ cm}^3} < 110\%$	50 Gy EQD _{2,3}		

Best, The Netherlands) with an in-plane resolution of 0.6×0.7 mm and 3.0 mm slice thickness with 0.3 mm gap were acquired and used for treatment planning. Imaging was done in supine position, similar to the treatment position. These images were loaded into Oncentra Brachy (version 4.3–4.5, Elekta AB, Stockholm, Sweden) and used for catheter reconstruction, delineation of the volumes of interest (prostate, vesicles, bladder, rectum, and urethra) and further treatment planning. The median prostate volume defined on these images after catheter placement was 31.5 (range: 16.6–74.2) cm^3 .

Treatment planning criteria, as shown in table 1, are based on the dose-volume indices, which are adapted from the GEC-ESTRO HDR Prostate Guidelines (Hoskin *et al* 2013). As starting point for treatment planning, dwell positions were activated inside the target volumes plus a 5 mm margin, excluding dwell positions located within 1 mm of the urethra. A median of 17 (range: 14–20) catheters were implanted with a source-step of 2.5 mm, which resulted in a median of 230 (range: 121–353) active dwell positions.

Clinically used treatment plans of all 39 patients were used for comparison. Clinical treatment planning started with an initial plan constructed with a class solution. This plan was then manually fine-tuned using graphical optimization. Initial plans of patients treated before mid-2015 were constructed using a class solution for IPSA (patients 1–3, 19, 23–29), plans of patients treated after mid-2015 were initialized with a class solution for HIPO (Maree *et al* 2019). Quality checks were done by a medical physicist and final plans were assessed for clinical acceptability by a physician using the planning criteria (table 1), and by a visual inspection of the dose distribution. Clinical plans of four patients (10, 11, 16, and 17), satisfied all planning criteria, additionally, 24 patients satisfied all coverage criteria, but violated at least one sparing criterion. The remaining 11 patients (4, 5, 7, 12, 14, 19, 20, 23, 24, 30, and 31) violated at least one sparing criterion and one coverage criterion.

2.2. Dose-volume calculations

We compared different approaches to generate treatment plans, i.e. a list of dwell times, either obtained using an inverse planning method, or using GOMEA. To compute the dose distribution and dose-volume indices corresponding to a list of dwell times, we used our in-house developed TG-43 (Rivard *et al* 2004) dose engine, which was validated against Oncentra Brachy (van der Meer *et al* 2019). This dose-engine is based on randomly sampled dose-calculation points within each volume of interest. Using more dose-calculation points results in a more accurate estimate of the dose-volume index values, but also requires more computation time and more computer memory. In this work, we chose to use a larger number of dose-calculation points than would be strictly necessary in clinical practice to show the maximally achievable performance of the methods. During optimization of the bi-objective planning model, 10^5 dose-calculation points were used for both the bladder and the prostate, 5×10^4 for the rectum, and 2×10^4 for both the vesicles and urethra, which were heuristically chosen based on the average size of each volume to provide sufficiently accurate computations of the dose-volume indices (van der Meer *et al* 2019). To overcome a bias towards the dose-calculation points used in the optimization process, all obtained plans were afterwards re-evaluated based on newly sampled dose-calculation points; 10^5 for each of the five volumes of interest. This is the same number of dose-calculation points as used in the dose-volume histogram calculations of Oncentra Brachy in clinical practice in our medical center.

2.3. Dose-penalty model setup

The dose-penalty model, used by IPSA and HIPO, can be optimized efficiently, but does not directly relate to the planning criteria (Gorissen *et al* 2013). To set up the dose-penalty model, the planner specifies lower and/or upper dose limits per volume of interest, together with a corresponding importance weight for each limit. Then, when a dose-calculation point receives a dose that is lower than the lower dose limit or higher than the upper dose limit, a penalty is given. Finally, all penalties are multiplied by the corresponding importance weight, and the final score of a plan is the sum of all these weighted penalties (Lessard and Pouliot 2001, Lahanas *et al* 2003). The optimization software within IPSA and HIPO then searches for a plan with the lowest total score.

We defined a lower dose limit on the target volumes (prostate and vesicles), and an upper dose limit on the OARs (bladder, rectum, and urethra). Setting an upper dose limit on the target volumes was found to be deteriorating plan quality in preliminary experiments on a limited number of patients, and these limits were thus not set. This setup results in two parameters that need to be specified per volume of interest: a dose-limit value and its corresponding weight. Tuning both the dose limit and weight at the same time complicates parameter tuning, as, for example, decreasing the upper dose limit has a similar effect as increasing the corresponding weight. Therefore, only one of the two is tuned, and the other is set based on the following reasoning.

For the target volumes, the planning criteria are defined in terms of volume-indices, i.e. $V_{100\%}^{prostate}$ and $V_{80\%}^{vesicles}$ (see table 1). The maximum value for these indices (100% of the volume) can be achieved if the lower dose limit is set to the corresponding dose, i.e. 100% and 80% of the planning-aim dose respectively, and the corresponding weight is tuned to be *large enough*.

For the OARs, however, the planning criteria are defined in terms of dose-indices, e.g. $D_{2\text{cm}^3}^{bladder} < 74\%$. When one would set the bladder upper dose limit to 74% of the planning-aim dose and the corresponding weight large, the best value that can be achieved is $D_{2\text{cm}^3}^{bladder} = 74\%$, but never $D_{2\text{cm}^3}^{bladder} < 74\%$, as the underlying dose-penalty model does not assign penalties below the upper dose limit. Therefore, for planning criteria based on dose-indices, the upper dose limit needs to be tuned, which we denote with L_b , L_r , and L_u for respectively the bladder, rectum, and urethra. Note that there are two planning criteria specified for the bladder and rectum in the clinical protocol (table 1), but only one upper dose limit can be specified per volume of interest in the dose-penalty model. Therefore, this dose limit needs to capture both criteria.

We setup the dose-penalty model as similar as possible to the behavior of the bi-objective planning model, to be able to transfer as good as possible the strengths of the bi-objective planning model to the dose-penalty model, and to allow for a fair comparison to directly optimizing the bi-objective planning model. Therefore, similar to the bi-objective planning model, where the trade-off between coverage and sparing criteria is explicitly taken into account, a single weight W_{OARs} was used for all OARs and a separate single weight was used for all targets. The latter was fixed to a value of 1, such that W_{OARs} models the importance of the OARs with respect to the targets. Then, for example, when $W_{\text{OARs}} = 0.5$, the target volumes are two times as important as the OARs, and when $W_{\text{OARs}} = 3$, the OARs are three times as important as the target volumes.

2.3.1. IPSA setup

IPSA is an inverse planning method that uses simulated annealing to optimize the underlying dose-penalty model (Lessard and Pouliot 2001). We set up IPSA as shown in table 2. Upper dose limits on OARs were only specified on the surface. As there are no dwell positions activated within the OARs, the highest dose will occur at the surface, and reducing surface dose will result in reduced total dose. The other methods discussed in this work optimize the *active* dwell positions. IPSA however handles dwell position activation internally, and the user must specify activation margins. We set these activation margins as currently specified in our clinical protocol (see section 2.1), to minimize the difference between methods. However, internal computations by IPSA might still result in a slightly different set of activated dwell positions. The number of dose-calculation points that IPSA uses internally cannot be controlled. Finally, a dwell time deviation constraint (DTDC) can be specified for IPSA. The DTDC aims to increase dose homogeneity by constraining dwell-times differences within a catheter. Dose homogeneity is not a planning criterion in our clinical protocol, and enabling the DTDC potentially deteriorates plan quality (Balvert *et al* 2015). The DTDC is therefore disabled.

2.3.2. HIPO setup

HIPO is an inverse planning method that uses the limited-memory Brodyen–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm to optimize the underlying dose-penalty model (Milickovic *et al* 2002). The parameter setup used for HIPO is shown in table 2. Dose limits are applied to both the surface and the volume depending on the sample point settings. According to our clinical practice, for each OAR, 500 surface points and 500 volume points were used. For each target, we used 1000 volume points, plus surface points with a density of 8cm^{-2} . Compared to IPSA, HIPO requires two additional dose limits; an upper dose limit on normal tissue dose and an upper dose limit on the primary target volume (prostate). To reduce the differences between the optimization aims in IPSA and HIPO, we effectively disabled these by setting the dose limit to 400% and the corresponding weight to 0.001, the lowest value possible. Furthermore, HIPO does not accept overlapping volumes. Therefore, the intersection between the prostate and the urethra was set to only belong to the urethra. This is different from the interpretation by IPSA, Oncentra Brachy, and our dose engine, as there, intersections between volumes are considered part of both volumes. Similar to the DTDC of IPSA, the dwell time gradient restriction (DTGR) in HIPO was disabled.

Table 2. Used parameter settings for the inverse planning methods IPSA and HIPO. For both methods, four parameters remain to be set by automatic parameter tuning: the dose limits L_b , L_r , and L_u for respectively the bladder, rectum, and urethra, and an overall weight W_{OARs} representing the importance of the OARs with respect to the targets. Dose limits are given in percentage of the planning-aim dose of 13 Gy. Weights have no unit.

	Volume	Usage	Lower dose		Upper dose	
			Weight	Limit	Limit	Weight
IPSA	Prostate surface & volume	Primary target	1	100%		
	Bladder surface	Organ			L_b^{IPSA}	$W_{\text{OARs}}^{\text{IPSA}}$
	Rectum surface	Organ			L_r^{IPSA}	$W_{\text{OARs}}^{\text{IPSA}}$
	Urethra surface	Organ			L_u^{IPSA}	$W_{\text{OARs}}^{\text{IPSA}}$
	Vesicles surface & volume	Target	1	80%		
HIPO	Prostate	Primary Target	1	100%	400%	0.001
	Bladder	Organ			L_b^{HIPO}	$W_{\text{OARs}}^{\text{HIPO}}$
	Rectum	Organ			L_r^{HIPO}	$W_{\text{OARs}}^{\text{HIPO}}$
	Urethra	Organ			L_u^{HIPO}	$W_{\text{OARs}}^{\text{HIPO}}$
	Vesicles	Target	1	80%		
	Normal tissue	—			400%	0.001

2.4. The bi-objective planning model

In the bi-objective planning model, the planning criteria are grouped into one coverage objective and one sparing objective (Luong *et al* 2018a, Maree *et al* 2019), that are referred to as the least coverage index (LCI) and least sparing index (LSI), as follows,

$$\text{LCI} = \min \left\{ V_{100\%}^{\text{prostate}} - 95, V_{80\%}^{\text{vesicles}} - 95 \right\},$$

$$\text{LSI} = 13 \text{ Gy} \times \min \left\{ \begin{array}{l} 86 - D_{1 \text{ cm}^3}^{\text{bladder}}, 74 - D_{2 \text{ cm}^3}^{\text{bladder}}, \\ 78 - D_{1 \text{ cm}^3}^{\text{rectum}}, 74 - D_{2 \text{ cm}^3}^{\text{rectum}}, \\ 110 - D_{0.1 \text{ cm}^3}^{\text{urethra}}, 150 - D_{50\%}^{\text{prostate}}, \\ 200 - D_{20\%}^{\text{prostate}} \end{array} \right\}.$$

Both LCI and LSI are constructed by combining the coverage or sparing criteria in table 1 in a worst-case manner. For example, maximizing the LCI implies that the least-fulfilled criterion is improved. The criterion that is least fulfilled changes during the optimization process, thus, ultimately, all criteria are improved. This is similar to how planners were observed to typically perform manual optimization of clinical plans, by trying to improve the least-fulfilled criterion iteratively (Maree *et al* 2019). An important property of the bi-objective planning model is that when both $\text{LCI} > 0\%$ and $\text{LSI} > 0 \text{ Gy}$, the plan is satisfactory. Optimization of the bi-objective planning model thus aims directly for satisfactory plans. This is a property that does not hold for the weighted-sum approaches of IPSA and HIPO, where unfulfilled criteria can be compensated by fulfilled criteria.

Adding the sparing criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ directly to the LSI would lead to a comparison of planning criteria with different quantities (volume and dose), which could give unexpected results as these are not directly comparable. In previous related work (Maree *et al* 2019), these two criteria were not included in the LSI, but added as hard constraints. It was found that these hard constraints did not influence the results, as these were never found to be violated in the considered patient set, which is a subset of the patients considered in this work. However, it cannot be guaranteed that these hard constraints are not violated on additional patients and with different planning strategies considered in this work. For completeness, the planning criteria are therefore not added as hard constraints but rewritten in terms of dose-indices so that they can be directly added to the LSI. The sparing criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ are therefore replaced by respectively $D_{50\%}^{\text{prostate}} < 150\%$ and $D_{20\%}^{\text{prostate}} < 200\%$.

2.5. Trade-off set generation

2.5.1. Trade-off set generation via automatic bi-objective parameter tuning

Using the setup of the dose-penalty model in table 2, four parameters need to be tuned, i.e. L_b , L_r , L_u , and W_{OARs} for either IPSA or HIPO (distinguished by a superscript). These four parameters were automatically tuned by optimizing the bi-objective planning model, i.e. by finding a set of plans with different high-quality trade-offs in the LCI and LSI. For this, the evolutionary algorithm MAMaLGaM (Rodrigues *et al* 2014) was used (full-covariance version). By automatic parameter tuning that entails running IPSA or HIPO, the dose distribution can no longer be quickly updated, and therefore, MAMaLGaM, a similar type of algorithm as GOMEA, has better performance (Luong *et al* 2018a). MAMaLGaM was used to perform the automatic bi-objective parameter

tuning throughout this work, while GOMEA was only used for direct optimization of the bi-objective planning model.

Per patient, MAMaLGaM was run for 25 iterations with a standard population size of 410 solutions, resulting in about 10^4 runs of each of the inverse planning method. This number of iterations was determined based on preliminary testing on a limited set of patients, and obtained trade-off curves were found to be non-improving for the last five iterations in all runs. The typical values for W_{OARs} that were found by preliminary testing were between 0 and 2. Therefore, for the start of the automatic parameter tuning, W_{OARs} was initialized randomly in the interval $[0, 2]$. However, to make sure that this range is not too restrictive, we set a higher upper bound (or maximum value) of 10. Weights cannot be negative. This means that initially, dose penalties for the OARs were at most twice as important as the target coverage penalties. The upper dose limits for the OARs were initialized around the target value stated in our clinical protocol in table 1. That is, L_b and L_r are initialized randomly in the range 60%–100%, and L_u in the range 100%–110%. Dose limits were bounded between 0% and 400%. The initialization ranges for the bladder and rectum are relatively large because both these volumes have two planning criteria specified in the clinical protocol (table 1). Additionally, these planning criteria are specified for a larger volume compared to that for the urethra, i.e. 1 cm^3 and 2 cm^3 compared to 0.1 cm^3 , which implies that a higher dose than the planning criterion is allowed, as long as the volume receiving that higher dose is not too large. To incorporate this in the dose-penalty model, we used a larger initialization range.

MAMaLGaM is a stochastic algorithm, that, depending on the initialization of the random number generator, finds a different trade-off set every run, although previous work shows the variation in outcome to be small (Luong *et al* 2018a). We reduced the effect of this randomness by running MAMaLGaM ten times per patient for each inverse planning method, and since we aim to compare to the best possible trade-off set, the resulting ten trade-off sets were combined into a single trade-off set by maintaining only the best plans. We refer to the resulting trade-off sets obtained by as *IPSA-tuned* and *HIPO-tuned*.

2.5.2. Trade-off set generation via random sampling

As alternative to trade-off set generation via automatic parameter tuning, we generated 1000 treatment plans by randomly sampling values for the parameters of the dose-penalty model (i.e. L_b, L_r, L_u , and W_{OARs}) for both IPSA and HIPO, similar as in (Bélanger *et al* 2019). Parameter values were sampled uniformly random in the same ranges as those that were used for the initialization of automatic bi-objective parameter tuning. The resulting trade-off set is then the subset of plans with only the best trade-offs. As no further optimization is performed in this approach, we did not repeat it multiple times like we did for MAMaLGaM. We refer to the trade-off sets generated via this approach as *IPSA-random* and *HIPO-random*.

2.5.3. Trade-off set generation with GOMEA

The bi-objective planning model is non-convex, non-linear, and non-smooth, and to optimize plans according to it, the state-of-the-art multi-objective real-valued GOMEA, was used (Luong *et al* 2018a, Bouter *et al* 2019). GOMEA exploits that the dose distribution can be quickly updated when only few dwell times change. GOMEA does not require any patient-specific parameters to be set, and a single run results in a trade-off set of hundreds of plans. For the experiments in this work, GOMEA was run for 10 min on a GPU (NVIDIA Titan X). In clinical practice, it was shown that a runtime of 30 s was sufficient (Bouter *et al* 2019). A longer runtime was however chosen here since we aimed to show maximally achievable plan quality, and because we increased the number of dose-calculation points, see section 2.2.

GOMEA is a stochastic algorithm, like MaMaLGaM, and therefore, parameter tuning was repeated ten times per patient, and the resulting trade-off sets were combined into a single trade-off set, which we refer to as *GOMEA-direct*.

2.6. Method comparisons

For each of the 39 patients, trade-off sets were automatically constructed using the three approaches discussed above, which results in five trade-off sets per patient (IPSA-tuned, HIPO-tuned, IPSA-random, HIPO-random, and GOMEA-direct). These trade-off sets were visually compared as trade-off curves, together with the clinically used plan, in the same figure. We compared the best obtained results of IPSA and HIPO (IPSA-tuned and HIPO-tuned) with GOMEA-direct. Additionally, we investigated the added value of performing parameter tuning by comparing IPSA-tuned with IPSA-random, and similarly HIPO-tuned with HIPO-random.

2.6.1. Trade-off performance measure: hypervolume

When a trade-off set is visualized as a trade-off curve, the area under this curve can be used as a performance indicator which is known as the hypervolume (Zitzler *et al* 2003, Auger *et al* 2009). A larger area (hypervolume) implies that the planning method achieved better trade-offs. The hypervolume is a performance measure that in itself may be hard to interpret, because its value depends on a reference point, i.e. the lower cutoff values of the

trade-off curve, and because it has no interpretable unit (in our case, dose times volume). It does however allow for a straightforward comparison of multiple trade-off curves, where a larger hypervolume is better. Differences in hypervolume were tested for statistical significance by a Wilcoxon signed-rank test with $\alpha = 0.05$.

2.6.2. Random parameter sampling versus automatic parameter tuning

All plans corresponding to the randomly sampled parameters were visualized as a scatterplot in the same figure (and not just the plans in the trade-off set, that have the best trade-offs). This gives insight in how many of the resulting plans are satisfactory (i.e. plans that satisfy all planning criteria), and whether the plans satisfy a large variety of different trade-offs. Further, the percentage of satisfactory plans was computed, which is an estimate of the probability that a randomly sampled parameter set results in a satisfactory plan, similar as in (Bélangier *et al* 2019).

2.6.3. Plan comparison

We aimed to investigate maximum performance of each inverse planning method, and therefore, IPSA-tuned, HIPO-tuned, and GOMEA-direct are compared in more detail. To assess quality of individual plans, three reference plans were selected from each trade-off set. For this, the following selection rules were used: a left (L) plan with an LCI-value of approximately -2% , a middle (M) plan with an LCI-value of approximately 1% , and a right (R) plan with an LCI-value of approximately 4% , corresponding with plans that have a target coverage of at least 93% , 96% , and 99% respectively. Of these selected plans, dose-volume indices of the planning criteria on which the LCI and LSI are based (section 2.4) were compared. Differences were tested for statistical significance by a Wilcoxon signed-rank test with $\alpha = 0.05$, and Bonferroni correction was applied for the 60 pairwise tests performed (IPSA/HIPO versus GOMEA for plans L/M/R for each of 10 dose-volume indices), resulting in a corrected $\alpha = 0.0008$.

2.7. Constructing class solutions

Class solutions, i.e. standard parameter sets for IPSA or HIPO, can be used to potentially overcome patient-specific parameter tuning, or to function as a good starting point for manual parameter tuning. A class solution for a specific inverse planning method can be easily constructed from a set of plans of previously treated patients, when these plans are constructed by that inverse planning method, and the corresponding parameter values are known, by taking e.g. for each parameter its average value.

We used plans obtained via automatic parameter tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned) as reference plans to construct multiple class solutions for IPSA and HIPO. By averaging the four parameters (L_b , L_r , L_w , and W_{OARs}) of the reference plans L, M, and R over all patients, we constructed three class solutions CS-L, CS-M, and CS-R, representing different trade-offs. Plan quality of these class solutions was compared to the automatic per-patient tuned results (IPSA-tuned or HIPO-tuned). We recorded how many of the resulting plans were satisfactory, and analyzed the obtained dose-volume index values. To prevent overfitting to the current patient set, leave-one-out cross validation was applied. A class solution was constructed by removing one patient from the set of 39 patients, and constructing a class solution from the remaining 38 patients. This class solution was then validated on the patient that was left out. We repeated this procedure by leaving each patient out once. Only results on the validation sets are shown. The dataset is rather small with only 39 patients. Therefore, leave-one-out cross validation is a good choice, as the training set is kept as large as possible (38 patients). Leave-one-out cross validation is however often time consuming as it needs to be repeated for each data point (patient), thus 39 times. Since constructing class solutions by taking means is extremely fast, as there is no need to re-do the IPSA/HIPO-tuning, this was not an issue.

3. Results

3.1. Trade-off curves

Trade-off curves based on the trade-off sets generated by automatic parameter tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned), and by GOMEA (GOMEA-direct) are shown in figure 1.

Parameter tuning of IPSA had a median runtime of 3 h per patient (range 1.5 h–5 h) and of HIPO 23 h (range 11.5 h–42 h), using the clinical software implementation of both methods. Computations were performed on different machines, which affected runtime, but computation time is patient dependent and relates mainly to the number of active dwell positions, i.e. the number of variables that IPSA or HIPO need to optimize.

For all patients, GOMEA obtained plans with the best trade-offs, i.e. in no case a plan obtained by automatic parameter tuning of IPSA or HIPO outperforms a plan obtained by GOMEA in both objectives simultaneously. For many patients, results obtained by automatic parameter tuning of HIPO were close or similar to results obtained by GOMEA, but larger differences can be observed for a few patients (4, 8, 10, 12, 23, 25, and 33). Only for patients 10 and 33, automatic parameter tuning of IPSA achieved somewhat better trade-offs than that of




Figure 1. Trade-off curves based on the trade-off sets obtained by automatic parameter tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned) for all 39 patients, compared to maximally achievable plan quality obtained by optimizing the bi-objective planning model with GOMEA (GOMEA-direct). Each obtained plan is visualized as a single dot, however, for many patients, individual dots cannot be distinguished. Triangles are the plans obtained by the class solutions (CS) L (<), M (^) and R (>), in the corresponding color of the method. The black squares represent the trade-off of the clinical plan, outlined in the color of the inverse planning method from which treatment planning was initialized (IPSA in blue, HIPO in red). Some clinical plans and class solution plans fell outside the figure range because it violated one or multiple planning criteria too much.

HIPO for lower LCI values, i.e. on the left side of the trade-off curve. For all other patients, trade-offs obtained by automatic parameter tuning of IPSA were inferior to trade-offs obtained by both automatic parameter tuning of HIPO and by GOMEA. The hypervolume of GOMEA-direct was, on average over all patients, 29.7 (range: 7.7–37.9), while the hypervolume of HIPO-tuned and IPSA-tuned were significantly lower ($p < 0.01$) with average values of respectively 27.6 (range: 7.0–37.0) and 24.6 (range: 0.2–35.8).

Obtained trade-off curves differ clearly for different patients in shape and maximally achievable plan quality (see figure 1). For 28/39 patients (72%), all three methods obtained satisfactory plans. Plans obtained by automatic tuning of HIPO were satisfactory for two additional patients (8 and 31), resulting in 30/39 (77%)




Figure 2. Plans obtained via random parameter sampling are illustrated by scattered dots (IPSA/HIPO-random), compared to the maximally achievable trade-offs obtained by automatically tuning IPSA/HIPO (IPSA/HIPO-tuned), and the maximally achievable trade-offs for that patient, obtained by GOMEA (GOMEA-direct). Illustrated for patients 8 (top) and 36 (bottom). Squares indicate the clinical plan. To highlight the differences between these approaches, the areas under the curves are shaded. Note that the areas overlap, e.g. the area under the curve of GOMEA-direct also includes the area shaded in blue/red.

patients with satisfactory plans. GOMEA obtained satisfactory plans for four additional patients (8, 14, 24, and 31), resulting in satisfactory plans for 32/39 (82%) patients. Finally, for seven (18%) patients (4, 18, 19, 20, 23, 30, and 34), none of the methods obtained satisfactory plans. Worst plans were observed for patient 19, in one region of the prostate, the catheters were not placed deep enough. To be able to fully cover the prostate, the planning criteria on the bladder or urethra were largely violated. There was no structural difference observed in prostate volume between patients without satisfactory plans (median 29.3 cm³) and the rest (median 31.5 cm³).

For almost all patients, all three methods outperformed the clinical plan by obtaining plans with a better trade-off than the trade-off of the clinical plan. Exceptions are patients 22 and 29, where automatic parameter tuning of IPSA was not able to outperform the clinical plan, while automatic tuning of HIPO and direct optimization with GOMEA did obtain a better trade-off.

3.2. Random parameter sampling versus automatic parameter tuning

For two example patients (8 and 36) that are representative for the patient set, plans obtained by randomly sampling parameters are shown in figure 2, together with the corresponding trade-off sets IPSA-random and HIPO-random, in which only the plans with the best trade-offs were maintained. A gap can be observed between the IPSA-tuned and IPSA-random trade-off curves, and similar for HIPO-tuned and HIPO-random. The mean hypervolume for IPSA-random was 22.2 (range: 0.0–33.7), which is lower than the hypervolume of IPSA-tuned ($p < 0.01$). Similar, the hypervolume for HIPO-random was 25.4 (range 5.4–36.0), which is lower than the hypervolume of HIPO-tuned ($p < 0.01$). Comparing the 28 patients for which both with automatic parameter tuning of IPSA and HIPO satisfactory plans were obtained, a median of 7.5% (range: 0%–63%) of the IPSA-random plans were satisfactory. For HIPO-random, this was 14.5% (range: 0%–78.5%).

3.3. Plan comparison

Table 3 shows the average dose-volume index values over all patients obtained by directly optimizing the bi-objective planning model by automatic tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned), and GOMEA (GOMEA-direct) for the three reference plans L, M, and R. In all cases, plans obtained by GOMEA spared the urethra and bladder better while similar prostate coverage $V_{100\%}^{prostate}$ was achieved, compared to IPSA-tuned and HIPO-tuned plans. Plans resulting from the automatic tuning of HIPO have lower high-dose regions within the prostate compared to plans resulting from direct optimization by GOMEA or by automatic tuning of IPSA. IPSA-tuned and HIPO-tuned plans exhibit significantly better vesicle coverage $V_{80\%}^{vesicles}$, albeit all three methods achieved values well above the planning criterion of 95%. The values obtained for the planning criteria on the rectum were well below their aim for all patients by all three methods.

3.4. Class solutions

The three class solutions (i.e. standard sets of parameters) obtained for both IPSA and HIPO, based on the three selected plans from the trade-off sets IPSA-tuned and HIPO-tuned are shown in table 4. The obtained dose-volume index values, as a result of leave-one-out cross-validation, are shown in table 5. Further, the corresponding trade-offs are visualized in figure 2. Combined, the three IPSA class solutions resulted in satisfactory plans for four patients. HIPO class solutions resulted in satisfactory plans for 18 patients in total, with class solutions based on plan M as the single best class solution with satisfactory plans for 13 patients. The class solutions based on plan R, aimed to result in plans with 99% coverage, violated the sparing criteria in all but one case.

For both IPSA and HIPO, standard deviations between obtained class solution parameters were larger for bladder and rectum limits L_b and L_r , and notably for the IPSA bladder limit, while the standard deviation for the urethra limit L_u was small. Compared to IPSA, HIPO class solutions had lower limits L_b and L_r for all plans, but somewhat higher L_u . In figure 1, it can be seen that the ordering of the class solutions such that $LCI(\text{plan L}) \leq LCI(\text{plan M}) \leq LCI(\text{plan R})$ was largely maintained. The class solutions were aimed to result in plans with a coverage of 93%, 96%, and 99% respectively, which was achieved on average, but the standard deviation of $V_{100\%}^{prostate}$ (and $V_{80\%}^{vesicles}$) was found to be large compared to the standard deviations in table 3.

4. Discussion

We introduced an automatic bi-objective parameter-tuning approach for inverse planning methods such as IPSA and HIPO. With this approach, the maximally achievable plan quality is automatically obtained, and the resulting set of plans provides the possibility for insightful comparison of the patient-specific trade-offs between the coverage of the targets and the sparing of the OARs. We showed that, when HIPO is properly tuned by use of this approach, plans with fairly similar trade-offs can be obtained as was maximally achievable by direct optimization of the bi-objective planning model using GOMEA for most, but not all, patients.

Plans obtained by automatic tuning of HIPO and by direct optimization of the bi-objective planning model with GOMEA outperformed the trade-off of the clinical plan for all patients, while plans obtained by automatic tuning of IPSA outperformed the clinical plan for all but two (22 and 29) patients. We would like to note that according to our clinical practice, planners often choose for a clinical plan with a higher coverage than strictly required to satisfy the coverage criteria, at the cost of (slightly) violating one or more sparing criterion, even though satisfactory plans might have been achieved for those patients by manual parameter tuning. For seven patients, it was not possible to obtain satisfactory plans with any of the approaches, which was caused by unfavorable implant geometry. There was no relation found between prostate volume and satisfactory plans.

The algorithms that are used for optimization in this work, GOMEA, MAMaLGaM, and also the simulated annealing algorithm within IPSA and L-BFGS within HIPO, are heuristic search algorithms that have no mathematical guarantee that maximum plan quality is reached. However, these methods do aim to obtain a high-quality solution within a reasonable amount of time and effort. Previous work indicates that the obtained results of GOMEA and MAMaLGaM on these and similar problems, are indeed of high quality and therefore (near-) maximally achievable quality. Throughout the article we, for simplicity, referred to this as maximally achievable plan quality achievable (Rodrigues *et al* 2014, Luong *et al* 2018a, Bouter *et al* 2019).

The random sampling approach we used was similar to the approach used in (Bélanger *et al* 2019), albeit on a different protocol and patient set. In that work, no notion of distance to maximally attainable values was given. Here, we showed that plan quality can be further improved by performing automatic parameter tuning.

The ultimate goal of treatment planning is to obtain a treatment plan that represents the best trade-off between (local) tumor control and normal tissue complications, and ideally, that is what is directly optimized for. The bi-objective parameter tuning approach presented in this work does not depend on the objectives that are optimized. Therefore, instead of using the LCI and LSI as objectives, the tumor control probability (TCP) and normal tissue complication probability (NTCP) could be used. To compute the TCP and NTCP, biological

Table 3. Dose-volume indices as described in section 2.1 of selected plans from IPSA-tuned and HIPO-tuned. Mean and standard deviation (between brackets) shown over all 39 patients. (+) indicates a statistically significant better value compared to GOMEA plans, (–) indicates a significant worse value. Volume-indices (*V*) in volume percentage, dose-indices (*D*) in Gy.

Plan	Method	$V_{prostate}^{100\%}$	$V_{vesicles}^{80\%}$	$D_{bladder}^{1\text{ cm}^3}$	$D_{bladder}^{2\text{ cm}^3}$	$D_{rectum}^{1\text{ cm}^3}$	$D_{rectum}^{2\text{ cm}^3}$	$D_{urethra}^{0.1\text{ cm}^3}$	$D_{prostate}^{50\%}$	$D_{prostate}^{20\%}$	$D_{prostate}^{90\%}$
	<i>Criteria</i>	> 95%	> 95%	< 11.2 Gy	< 9.6 Gy	< 10.1 Gy	< 9.6 Gy	< 14.3 Gy	< 19.5 Gy	< 26 Gy	> 13 Gy
L	GOMEA	93.0(0.1)	97.0(2.5)	9.6(0.8)	8.6(0.7)	8.7(0.9)	7.7(0.9)	13.5(0.5)	15.8(0.6)	19.9(1.4)	13.3(0.1)
	IPSA	93.0(0.4)	98.6(2.2) +	9.9(0.6) –	9.0(0.5) –	9.0(1.1) –	8.0(1.0) –	13.8(0.7) –	16.2(0.9) –	20.8(2.3) –	13.3(0.2) +
	HIPO	93.0(0.2)	98.1(2.2) +	9.8(0.6) –	8.9(0.6) –	9.0(0.9) –	8.0(0.9) –	13.6(0.5) –	15.4(0.5) +	19.4(1.6) +	13.2(0.1) –
M	GOMEA	96.0(0.1)	98.2(1.5)	10.2(1.0)	9.1(0.9)	9.1(1.1)	8.0(1.1)	14.0(0.7)	16.1(0.7)	20.3(1.6)	13.6(0.3)
	IPSA	95.9(0.8)	99.2(1.3) +	10.5(0.7) –	9.5(0.7) –	9.3(1.2) –	8.3(1.2) –	14.4(0.9) –	16.8(1.0) –	21.7(2.4) –	13.8(0.3) +
	HIPO	96.0(0.0)	98.8(1.3) +	10.4(0.8) –	9.4(0.7) –	9.2(1.1)	8.2(1.1)	14.1(0.7) –	15.7(0.6) +	19.9(1.8) +	13.5(0.2) –
R	GOMEA	98.9(0.4)	99.7(0.4)	11.4(1.1)	10.3(1.0)	9.6(1.3)	8.6(1.3)	15.0(1.0)	16.8(0.9)	21.4(2.0)	14.2(0.5)
	IPSA	98.6(1.4) –	99.7(0.7)	11.7(0.9) –	10.5(0.8)	9.9(1.4)	8.8(1.3)	15.6(0.9) –	18.1(1.0) –	23.4(2.4) –	14.8(0.5) +
	HIPO	98.9(0.5)	99.7(0.4)	11.5(1.0)	10.4(0.9) –	9.9(1.2)	8.9(1.2) –	15.1(0.9) –	16.4(0.8) +	20.9(2.1) +	13.9(0.4) –

Table 4. Obtained class solutions (CS) for IPSA and HIPO by averaging selected plans obtained by automatically tuning of respectively IPSA and HIPO (trade-off sets IPSA-tuned and HIPO-tuned). Dose limits L_b , L_r , and L_u corresponding to the bladder, rectum and urethra respectively, are expressed in percentage of the planning-aim dose of 13 Gy. W_{OARs} is expressed as relative importance of the OARs to the target volumes, see table 2. Standard deviations of the corresponding parameters over all patients are provided between brackets. ‘Satisfactory plans’ indicates for how many of the 39 patients, plans that satisfy all planning criteria were obtained with the given class solution.

CS		L_b	L_r	L_u	W_{OARs}	Satisfactory plans
L	IPSA	107(91)	104(65)	102(5)	1.7(2.0)	4
	HIPO	69(17)	93(24)	106(3)	1.4(1.0)	11
M	IPSA	129(105)	106(58)	106(6)	1.4(1.7)	3
	HIPO	78(34)	100(37)	110(5)	0.9(0.6)	13
R	IPSA	181(116)	127(75)	114(7)	0.8(1.6)	0
	HIPO	84(39)	121(50)	117(6)	0.7(1.5)	1

models have been developed (Baumann and Petersen 2005). These models however depend on a number of radiobiological parameters which are unknown or uncertain in practice. Although these models give a correct description of the main characteristics of the radiation response, caution has to be taken if these models are to be applied to patients (Brahme 2001, Karger 2006, Dinkla *et al* 2013). When the clinical planning criteria (table 1) are formulated in terms of TCP/NTCP, this can be rather straightforwardly incorporated into our automatic parameter tuning. However, to obtain the best treatment plans in terms of the current clinical protocol that is formulated in terms of D and V indices, it is logical to directly optimize for the latter, which is achieved by the current formulation of the bi-objective planning model.

Other (dose-volume based) multi-objective approaches exist (Milickovic *et al* 2002, Lahanas *et al* 2003, Cui *et al* 2018a, 2018b), but the bi-objective planning model specifically allows for a direct optimization of the dose-volume-based planning criteria without having too many objectives. Optimizing all planning criteria as separate objectives results in a many-objective optimization problem, and solving these problems by presenting a representative trade-off set is difficult and time consuming, even with state-of-the-art algorithms (Luong *et al* 2018b). Additionally, visualization of the trade-off set when using four or more objectives is no longer straightforward. In practice, besides the planning criteria, planners also look at the 3D dose distribution to assess plan quality. It was shown in a retrospective observer study that optimizing the bi-objective planning model with GOMEA resulted in plans that were considered clinically acceptable (Maree *et al* 2019). However, because of fundamental differences in how plans are generated by IPSA and HIPO, structural differences in the dose distribution might have occurred that are harder to quantify, besides the observed differences in the obtained values for the planning criteria.

Well-spread high-quality trade-off sets were obtained when applying automatic bi-objective parameter tuning to both IPSA and HIPO. This suggests that the chosen parameter setup, with only four parameters of the dose-penalty model that were automatically tuned, was sufficiently flexible. We chose in this work not to include upper dose limits for the prostate in IPSA and HIPO, mainly based on prior observations with using GOMEA to solve the bi-objective problem formulation that the planning criteria on the $V_{150\%}^{\text{prostate}}$ and $V_{200\%}^{\text{prostate}}$, as specified in our clinical protocol (table 1), are rather easy to satisfy (Maree *et al* 2019). The results in table 3 show that, even without upper dose limits on the prostate, the planning criteria are well below their aim. Hence, these planning criteria do not play a role in the automatic parameter tuning process of IPSA and HIPO, nor in the optimization process of GOMEA. For HIPO, the obtained $V_{150\%}^{\text{prostate}}$ and $V_{200\%}^{\text{prostate}}$ values are even better than with the other methods. This shows that it is indeed not required for IPSA and HIPO to furthermore add upper dose limits for the prostate, and that the other upper dose limits reduce high-dose regions in the prostate below their aim. However, if the planning criteria on the $V_{150\%}^{\text{prostate}}$ and $V_{200\%}^{\text{prostate}}$ would have been stricter in our clinical protocol, it might have been required to use upper dose limits for the prostate. Despite this observation, including more parameters of the dose-penalty model in the automatic tuning process could potentially improve resulting plan quality. However, this increases computation time of the automatic tuning process, as more parameters need to be tuned. It would furthermore complicate the tuning process, as multiple parameters have the same effect on the obtained plan, of which we gave an example in section 2.3. Increasing planning complexity increases the risk of ultimately achieving inferior plans, both for manual and automatic parameter tuning. When applying automatic bi-objective parameter tuning to a different clinical protocol, for a fair comparison between IPSA, HIPO, and GOMEA, the parameter setup needs to be reconsidered. The reasonings and results in this work could be used as guidelines to do so.

Automatic bi-objective parameter tuning resulted for IPSA in inferior trade-offs compared to HIPO for many patients, although a similar setup of the dose-penalty model was used. For this work, we used the clinical

Table 5. Dose-volume index values of class solutions (CS) for IPSA and HIPO. Results were obtained by leave-one-out cross validation. Average and standard deviation (between brackets) computed over all 39 patients, volume-indices (V) in volume percentage, and dose-indices (D) in Gy.

CS		$V_{prostate}^{100\%}$	$V_{vesicles}^{80\%}$	$D_{bladder}^{1\text{ cm}^3}$	$D_{bladder}^{2\text{ cm}^3}$	$D_{rectum}^{1\text{ cm}^3}$	$D_{rectum}^{2\text{ cm}^3}$	$D_{urethra}^{0.1\text{ cm}^3}$	$D_{prostate}^{50\%}$	$D_{prostate}^{20\%}$	$D_{prostate}^{90\%}$
Criteria		> 95%	> 95%	< 11.2 Gy	< 9.6 Gy	< 10.1 Gy	< 9.6 Gy	< 14.3 Gy	< 19.5 Gy	< 26 Gy	> 13 Gy
L	IPSA	93.7(4.5)	97.3(9.0)	10.6(1.2)	9.6(1.3)	9.6(1.3)	8.6(1.3)	14.2(1.4)	16.6(0.7)	21.8(1.8)	13.4(0.7)
	HIPO	93.2(3.0)	98.8(1.3)	9.5(0.7)	8.6(0.7)	9.2(1.1)	8.3(1.1)	13.6(0.1)	15.4(0.4)	19.7(1.2)	13.2(0.3)
M	IPSA	95.5(4.3)	97.5(8.9)	11.3(1.5)	10.2(1.4)	9.7(1.4)	8.7(1.3)	14.6(1.3)	17.1(0.7)	22.5(1.9)	13.8(0.7)
	HIPO	96.1(2.4)	99.4(0.8)	10.1(0.7)	9.3(0.8)	9.6(1.2)	8.7(1.2)	14.0(0.2)	15.7(0.4)	20.2(1.3)	13.4(0.2)
R	IPSA	97.7(3.5)	98.9(4.0)	12.3(1.8)	11.0(1.7)	10.4(1.7)	9.2(1.6)	15.8(2.0)	18.4(1.1)	24.4(2.5)	14.7(0.9)
	HIPO	98.1(1.7)	99.6(0.6)	10.6(0.7)	9.7(0.8)	10.2(1.4)	9.2(1.4)	14.9(0.2)	16.4(0.6)	21.4(1.5)	13.8(0.3)

software implementation of IPSA and HIPO, and there are multiple fundamental differences between those two methods that could be the cause of these differences. In HIPO, a quadratic dose-penalty model is used, while in, IPSA a linear dose-penalty model is used. Because of these model differences, different solvers were used to optimize them. In HIPO, L-BFGS is used, while the simulated annealing algorithm is used in IPSA to optimize the dose-penalty model. Later improvements on IPSA showed that linear programming is capable of solving this dose-penalty model with higher accuracy, with up to about one percent improvement in dose-volume indices of the prostate and urethra for some patients, but this version is not clinically available (Alterovitz *et al* 2006). Other differences between the two methods are that for IPSA, dose-calculation points were only sampled on the surface of OARs, as only surface dose limits were specified, whereas for HIPO dose-calculation points were also sampled within the OAR volumes. Further, in IPSA the delineated contours specified by the planner were directly used to sample dose-calculation points, while in HIPO, the contours were converted in a 3D mesh of the organ surface. Even when there are many delineated contours, the latter approach is closer to how dose-volume indices are evaluated within Oncentra Brachy, and in our in-house-developed dose engine. Different dose-point sampling and volume reconstruction methods have been shown to affect computed plan quality (van der Meer *et al* 2019).

When looking into the dose-volume index values of the optimized plans, it was noted that direct optimization of the bi-objective planning model with GOMEA resulted in a lower vesicle dose than automatic tuning of IPSA and HIPO, albeit all obtained values were well above their aim. When GOMEA optimizes the bi-objective planning model, it applies an internal weighting mechanism of the planning criteria aimed to prevent that the second-worst objective is not further optimized when the worst objective cannot be improved (Bouter *et al* 2019). This mechanism was not applied in the automatic parameter tuning of IPSA and HIPO. However, it seems not to be required, as the obtained vesicle dose was found to be higher than obtained with GOMEA. This might be a result of how the dose-penalty model was set up in the automatic parameter tuning approach, by using a single weight for both target volumes.

Automatic patient-specific parameter tuning, i.e. its application to daily clinical cases, as performed in this work, is not possible in current clinical practice due to the required computational effort. For that, GPU-based versions are required, which are starting to become available (Bélanger *et al* 2019, Bouter *et al* 2019). A recent study showed that a GPU-based version of IPSA is able to compute 1000 plans within 10 s (Bélanger *et al* 2019), which suggests that our automatic parameter-tuning approach can be run within approximately 4 min, and most likely even faster. GOMEA was shown to be able to optimize the bi-objective planning model in 30 s on a GPU (Bouter *et al* 2019), or within 5 min on a central processing unit (CPU) (Luong *et al* 2019). This is all well within the time limits of the clinical practice, especially since manual graphical optimization often takes 30 min or more (Maree *et al* 2019).

For clinical practice today, class solutions could therefore be used for IPSA or HIPO. With the construction of class solutions based on the obtained HIPO-tuned (or IPSA-tuned) plans, we showed that multi-objective, i.e. trade-off representing, class solutions can effectively be obtained in a principled way. However, subsequent (automatic) patient-specific parameter tuning is still clearly required for IPSA and HIPO to obtain the highest-possible quality plans. Construction of class solutions can be further extended by splitting the patient set based on anatomical features, e.g. prostate size, and constructing a class solution for each subset might increase plan quality. Alternatively, patient features could be correlated to the method's parameters, as was done for IPSA in recently published works (Cui *et al* 2018a, 2018b). Also then, high-quality reference plans are required for all patients that are generated by the inverse planning method of interest, as the corresponding parameters must be known. In practice, treatment plans may have been adapted manually, so that the corresponding parameter values are not known, or the plans may be suboptimal because of the manual treatment planning process that has to be performed in limited time (Maree *et al* 2019). Our bi-objective parameter-tuning approach is then still an essential tool that can generate high-quality reference plans automatically.

The three obtained class solutions for HIPO (table 4) are currently being introduced in our clinical practice for which we will follow the following strategy: first, class solution M is applied. If the coverage of the resulting plan is too low, class solution R is applied, else, if the resulting coverage is too high, class solution L is applied. Then, manual parameter tuning or graphical optimization is performed to further improve plan quality. If manual parameter tuning is performed by a planner, it is recommended to follow the same strategy of adapting only the four parameters in table 2. However, the observed difference in performance between the results from automatic parameter tuning and random parameter sampling do indicate that this is likely still not trivial to do.

5. Conclusion

We successfully developed an automatic bi-objective parameter tuning approach that when applied to IPSA and HIPO, results in a set of plans with a wide range of trade-offs. By tuning HIPO, better trade-offs were obtained than by tuning IPSA. For most patients, automatic tuning of HIPO resulted in plans close to the maximally achievable plan quality obtained by optimizing the bi-objective planning model directly with GOMEA.

Automatic parameter tuning was furthermore shown to improve plan quality significantly compared to random parameter sampling. Finally, from the automatically tuned plans, three class solutions were successfully constructed representing different trade-offs that are being introduced in our clinical practice.

Acknowledgments

This work is part of the research programme IPPSI-TA with project number 628.006.003, which is financed by the Netherlands Organisation for Scientific Research (NWO) and Elekta. The authors would like to thank J Schillings and B van Veelen from Elekta, Veenendaal, The Netherlands, for providing an interface for using HIPO and IPSA, and directly accessing data in Oncentra Brachy. We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan X GPU used for this research.

References

- Alterovitz R, Lessard E, Pouliot J, Hsu I-C J, O'Brien J F and Goldberg K 2006 Optimization of HDR brachytherapy dose distributions using linear programming with penalty costs *Med. Phys.* **33** 4012–9
- Auger A, Bader J, Brockhoff D and Zitzler E 2009 Theory of the hypervolume indicator *Proc. of the Genetic and Evolutionary Computation Conf.—GECCO '07* pp 87–102
- Bachand F, Martin A G, Beaulieu L, Harel F and Vigneault É 2009 An eight-year experience of HDR brachytherapy boost for localized prostate cancer: biopsy and PSA outcome *Int. J. Radiat. Oncol. Biol. Phys.* **73** 679–84
- Balvert M, Gorissen B L, den Hertog D and Hoffmann A L 2015 Dwell time modulation restrictions do not necessarily improve treatment plan quality for prostate HDR brachytherapy *Phys. Med. Biol.* **60** 537–48
- Baumann M and Petersen C 2005 TCP and NTCP: a basic introduction *Rays—Int. J. Radiol. Sci.* **30** 99–104
- Bélanger C, Cui S, Ma Y, Després P, Adam M, Cunha J and Beaulieu L 2019 A GPU-based multi-criteria optimization algorithm for HDR brachytherapy *Phys. Med. Biol.* **64** 105005
- Bouter A, Alderliesten T, Pieters B R, Bel A, Niatsetski Y and Bosman P A N 2019 GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy *Med. Phys.* **46** 3776–87
- Bouter A, Luong N H, Witteveen C, Alderliesten T and Bosman P A N 2017 The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm *Proc. of the Genetic and Evolutionary Computation Conf. on—GECCO '17* ed G Ochoa et al (New York: ACM Press) pp 537–44
- Brahme A 2001 Individualizing cancer treatment: biological optimization models in treatment planning and delivery *Int. J. Radiat. Oncol. Biol. Phys.* **49** 327–37
- Cui S, Després P and Beaulieu L 2018a A multi-criteria optimization approach for HDR prostate brachytherapy: I. Pareto surface approximation *Phys. Med. Biol.* **63** 205004
- Cui S, Després P and Beaulieu L 2018b A multi-criteria optimization approach for HDR prostate brachytherapy: II. Benchmark against clinical plans *Phys. Med. Biol.* **63** 205005
- Dinkla A M, van der Laarse R, Kaljouw E, Pieters B R, Koedooder K, van Wieringen N and Bel A 2015 A comparison of inverse optimization algorithms for HDR/PDR prostate brachytherapy treatment planning *Brachytherapy* **14** 279–88
- Dinkla A M, Pieters B R, Koedooder K, Van Wieringen N, Van Der Laarse R, Van Der Griet J N, Rasch C R, Koning C C and Bel A 2013 Improved tumour control probability with MRI-based prostate brachytherapy treatment planning *Acta Oncol.* **52** 658–65
- Fatya M, Williamson J F, Dogan N, Todor D, Siebers J V, George R, Barani I and Hagan M 2009 A comparison of HDR brachytherapy and IMRT techniques for dose escalation in prostate cancer: a radiobiological modeling study *Med. Phys.* **36** 3995–4006
- Gorissen B L, den Hertog D and Hoffmann A L 2013 Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR brachytherapy *Phys. Med. Biol.* **58** 1041–57
- Holm A, Larsson T and Carlsson Tedgren A 2012 Impact of using linear optimization models in dose planning for HDR brachytherapy *Med. Phys.* **39** 1021–8
- Hoskin P J, Colombo A, Henry A, Niehoff P, Paulsen Hellebust T, Siebert F A and Kovacs G 2013 GEC/ESTRO recommendations on high dose rate afterloading brachytherapy for localised prostate cancer: an update *Radiother. Oncol.* **107** 325–32
- Karger C P 2006 Biological models in treatment planning *New Technologies in Radiation Oncology* ed W Schlegel, T Bortfeld, A L Grosu (Berlin: Springer) pp 221–35
- Lahanas M, Baltas D and Zamboglou N 2003 A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy *Phys. Med. Biol.* **48** 399–415
- Lessard E and Pouliot J 2001 Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function *Med. Phys.* **28** 773–9
- Luong N H, Alderliesten T, Bel A, Niatsetski Y and Bosman P A N 2018a Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment *Swarm Evol. Comput.* **40** 37–52
- Luong N H, Alderliesten T and Bosman P A N 2018b Improving the performance of MO-RV-GOMEA on problems with many objectives using Tchebycheff scalarizations *Proc. of the Genetic and Evolutionary Computation Conf.—GECCO '18* pp 705–12
- Luong N H, Alderliesten T, Pieters B R, Bel A, Niatsetski Y and Bosman P A N 2019 Fast and insightful bi-objective optimization for prostate cancer treatment planning with high-dose-rate brachytherapy *Appl. Soft Comput.* **84** 105681
- Maree S C, Luong N H, van Wieringen N, Bel A, Hinnen K A, Westerveld G H, Pieters B R, Bosman P A N and Alderliesten T 2019 Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy—a retrospective observer study *Brachytherapy* **18** 396–403
- van der Meer M C, Bosman P A N, Pieters B R, Niatsetski Y, van Wieringen N, Alderliesten T and Bel A 2019 Sensitivity of dose-volume indices to computation settings in high-dose-rate prostate brachytherapy treatment plan evaluation *J. Appl. Clin. Med. Phys.* **20** 66–74
- Milickovic N, Lahanas M, Papagiannopoulou M, Zamboglou N and Baltas D 2002 Multiobjective anatomy-based dose optimization for HDR-brachytherapy with constraint free deterministic algorithms *Phys. Med. Biol.* **47** 2263–80

- Morton G C, Sankrecha R, Halina P and Loblaw A 2008 A comparison of anatomy-based inverse planning with simulated annealing and graphical optimization for high-dose-rate prostate brachytherapy *Brachytherapy* **7** 12–6
- Pieters B R, van der Grient J N B, Blank L E C M, Koedooder K, Hulshof M C C M and de Reijke T M 2006 Minimal displacement of novel self-anchoring catheters suitable for temporary prostate implants *Radiother. Oncol.* **80** 69–72
- Rivard M J, Coursey B M, DeWerd L A, Hanson W F, Saiful Huq M, Ibbott G S, Mitch M G, Nath R and Williamson J F 2004 Update of AAPM Task Group No. 43 Report: a revised AAPM protocol for brachytherapy dose calculations *Med. Phys.* **31** 633–74
- Rodrigues S, Bauer P and Bosman P A N 2014 A novel population-based multi-objective CMA-ES and the impact of different constraint handling techniques *Proc. of the 2014 Conf. on Genetic and Evolutionary Computation—GECCO '14* pp 991–8
- RTOG 2016 *RTOG 0924 Androgen Deprivation Therapy and High Dose Radiotherapy With or Without Whole-Pelvic Radiotherapy in Unfavorable Intermediate or Favorable High Risk Prostate Cancer: a Phase III Randomized Trial*
- Siauw T, Cunha A, Atamtürk A, Hsu I-C, Pouliot J and Goldberg K 2011 IPIP: a new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices *Med. Phys.* **38** 4045–51
- Venselaar J, Meigooni A S, Baltas D and Hoskin P J 2012 *Comprehensive Brachytherapy: Physical and Clinical Aspects* (London: Taylor and Francis)
- Zitzler E, Thiele L, Laumanns M, Fonseca C M and Da Fonseca V G 2003 Performance assessment of multiobjective optimizers: an analysis and review *IEEE Trans. Evol. Comput.* **7** 117–32