Physics in Medicine & Biology

Robust optimization for HDR prostate brachytherapy applied to organ reconstruction uncertainty

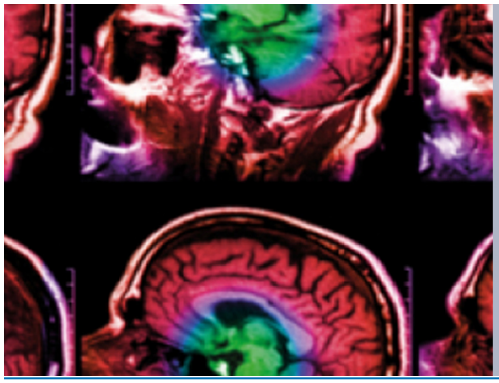

Physics in Medicine & Biology

IPEM Institute of Physics and Engineering in Medicine

OPEN ACCESS

RECEIVED
13 November 2020

REVISED
14 January 2021

ACCEPTED FOR PUBLICATION
27 January 2021

PUBLISHED
12 February 2021

PAPER

Robust optimization for HDR prostate brachytherapy applied to organ reconstruction uncertainty

Marjolein C van der Meer[1], **Peter A N Bosman**[2], **Yury Niatsetski**[3], **Tanja Alderliesten**[4], **Bradley R Pieters**[1] **and Arjan Bel**[1]

[1] Department of Radiation Oncology, Amsterdam UMC, location Academic Medical Center, University of Amsterdam, Amsterdam, The Netherlands
[2] Life Sciences and Health research group, Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
[3] Physics and Advanced Development, Elekta, Veenendaal, The Netherlands
[4] Department of Radiation Oncology, Leiden University Medical Center, Leiden, The Netherlands

E-mail: marjolein.vandermeer@amsterdamumc.nl

Keywords: robust optimization, HDR brachytherapy, prostate neoplasms, treatment planning, dose-volume indices

Abstract

Purpose. Recently, we introduced a bi-objective optimization approach based on dose-volume indices to automatically create clinically good HDR prostate brachytherapy plans. To calculate dose-volume indices, a reconstruction algorithm is used to determine the 3D organ shape from 2D contours, inevitably containing settings that influence the result. We augment the optimization approach to quickly find plans that are robust to differences in 3D reconstruction. *Methods.* Studied reconstruction settings were: interpolation between delineated organ contours, overlap between contours, and organ shape at the top and bottom contour. Two options for each setting yields 8 possible 3D organ reconstructions per patient, over which the robust model defines minimax optimization. For the original model, settings were based on our treatment planning system. Both models were tested on data of 26 patients and compared by re-evaluating selected optimized plans both in the original model (1 organ reconstruction, the difference determines the cost), and in the robust model (8 organ reconstructions, the difference determines the benefit). *Results*. Robust optimization increased the run time from 3 to 6 min. The median cost for robust optimization as observed in the original model was −0.25% in the dose-volume indices with a range of [−0.01%, −1.03%]. The median benefit of robust optimization as observed in the robust model was 0.93% with a range of [0.19%, 4.16%]. For 4 patients, selected plans that appeared good when optimized in the original model, violated the clinical protocol with more than 1% when considering different settings. This was not the case for robustly optimized plans. *Conclusions*. Plans of high quality, irrespective of 3D organ reconstruction settings, can be obtained using our robust optimization approach. With its limited effect on total runtime, our approach therefore offers a way to account for dosimetry uncertainties that result from choices in organ reconstruction settings that is viable in clinical practice.

1. Introduction

In current clinical practice in high-dose-rate (HDR) prostate brachytherapy the dwell times are optimized in a treatment planning system. Two commonly used optimization algorithms are inverse planning simulated annealing (IPSA) (Lessard and Pouliot 2001) and hybrid inverse planning optimization (HIPO) (Karabis *et al* 2009). In order to perform optimization, the clinical requirements for determining the quality of a treatment plan have to be given as input. These requirements are among others often defined in terms of dose-volume indices. In both IPSA and HIPO these dose-volume indices are used indirectly to optimize the dwell times of a treatment plan.

In the past years, multiple approaches have been proposed for optimizing directly on dose-volume indices (Siauw *et al* 2011, Deist and Gorissen 2016, Bouter *et al* 2019). One of these approaches, explored in our group, uses the multi-objective real-valued gene-pool optimal mixing evolutionary algorithm (MO-RV-GOMEA) (Bouter *et al* 2019). In this approach, the clinical requirements are combined into two optimization objectives, one for covering the targets and one for sparing the organs at risk (OARs). By optimizing directly on dose-volume indices, it was shown that with MO-RV-GOMEA better trade-offs between target coverage and organ sparing could be obtained than with IPSA and HIPO, both in a clinical-practice situation as well as in a direct comparison between the optimization algorithms (Maree *et al* 2019, 2020). As of March 2020, the bi-objective dwell time optimization has been introduced clinically at our hospital.

While the dose-volume indices of a treatment plan are uniquely defined, computation of the values for these indices in practice, is not. A key reason is that dose-volume indices are computed from 3D organ volumes. However, clinical delineations usually comprise 2D organ contours on medical scans. The 3D organ shapes are then reconstructed from these 2D contours. Such a reconstruction is however not uniquely defined. Reconstruction algorithms can be configured in various ways, depending on a particular implementation in a treatment planning system. The resulting dose-volume indices of a treatment plan therefore depend on the treatment planning system in which they are evaluated (Kirisits *et al* 2007, van der Meer *et al* 2019).

It has been shown that several key choices (settings) in the reconstruction algorithm can have a large influence on the calculated values for the dose-volume indices (van der Meer *et al* 2019). Therefore, optimizing directly on these indices likely results in treatment plans being overfit to the particular settings of the treatment planning system that is used. If in reality another setting resulted in shapes closer to the true shapes of the volumes of interest, then the true plan quality may be different from what is shown in the system in which the plan was optimized. In general, the settings of the organ reconstruction algorithm leading to shapes closest to the real organ shapes are unknown and may differ per patient.

An approach to deal with these uncertainties is to take them into account in the optimization problem. Robust optimization has previously already been applied in radiotherapy by means of stochastic programming and minimax optimization (e.g. in the context of proton planning (Fredriksson and Bokrantz 2014) and intensity-modulated radiation therapy (Unkelbach *et al* 2018)). Stochastic programming requires knowledge on the probability distribution(s) of the uncertain parameter(s), which is often unavailable. In contrast, minimax optimization only requires a set of optimization scenarios, and is therefore more flexible. In the context of HDR prostate brachytherapy, minimax optimization has previously been applied to delineation uncertainties (Balvert *et al* 2019).

In this work, we will apply minimax optimization over different (key) organ reconstruction settings, meaning that during optimization the quality of a plan is determined not for a single, but for multiple organ reconstruction settings. This way, treatment plans can be obtained with good quality for all settings, or at least, the potentially large negative effect of overfitting to one particular setting will be minimized. In this work, we extend MO-RV-GOMEA in an efficient way to perform such a robust optimization, and compare this extension to the original optimization. Beyond this novel extension itself, the aim is to study the added value of such robust dwell time optimization.

2. Materials and methods

2.1. Patient data

Patient data (from a previous study (van der Meer *et al* 2019)) was used retrospectively. The patient data consisted of 26 consecutive patients who underwent prostate HDR brachytherapy between February 2015 and April 2017. Magnetic resonance imaging (MRI) scans were acquired with a resolution in the axial planes of 0.52×0.52 mm and a slice thickness of 3.3 mm (including a 0.3 mm gap). A median number of 16 catheters were placed (range 14–20). Delineation of catheters and organs (prostate, bladder, rectum, urethra, base of seminal vesicles) was performed manually on nearly all of the axial slices, using the delineations suggested by the interpolation algorithm of the clinical treatment planning system (Oncentra Brachy, version 4.5, Elekta, Sweden) on the other axial slices.

A treatment plan was created with either IPSA (for 5 patients treated before mid-2015) or HIPO (after mid-2015). The mHDR-v2 source description of the 192-Iridium source was used in clinical treatment planning. Manual fine-tuning was then performed using graphical optimization. The median number of activated dwell positions was 454.5 with a range of 250–668.

For 6 patients (patients 3 and 14–18), a second set of manual delineations was available, created retrospectively by an independent observer and corrected and approved by an experienced radiation oncologist (to examine the inter-observer variability).

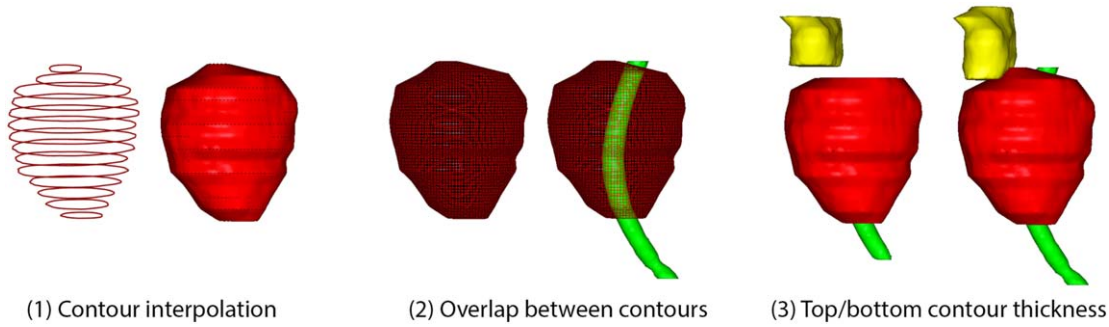

Figure 1. An illustration of the three organ reconstruction settings studied. Examples of organ shapes for prostate (red), urethra (green), and seminal vesicles (yellow) are shown. Setting 1 shows how interpolation between 2D delineated organ contours (left) results in a 3D organ shape (right). Setting 2 shows how overlap between contours can be ignored (left) or taken into account (right). Setting 3 shows an exaggerated case of how the organ shape at the top and bottom contour can use partial slice thickness (left) or full slice thickness (right).

2.2. Organ reconstruction settings

Software for computing the dose-volume indices of a patient was previously in-house developed and validated (van der Meer *et al* 2019). The input of our in-house developed software was the treatment date, delineated contours, catheter information, and source information, all retrieved from the DICOM data. In this work, we study three different organ reconstruction algorithm settings (van der Meer *et al* 2019), as illustrated in figure 1. Different treatment planning systems make different choices for these settings. For each of the settings, there are two options:

1. Delineated contours fill the volume spanned by their axial slice to reconstruct the 3D organ volume, or interpolation is used to add contours in between slices. In previous work, one interpolated contour was added half-way between each consecutive pair of delineated contours (van der Meer *et al* 2019). In a later study, three interpolated contours were added between consecutive pairs of contours (Bouter *et al* 2019), so this approach is followed.
2. The urethra is either considered to be part of the organs it intersects, or not. For our patient data, the urethra was delineated as the part of the urinary catheter through the prostate into the bladder, since part of the urethra itself is often not well visible on the MRI. Therefore, intersections exist between the prostate and the urethra, as well as between the bladder and the urethra.
3. Top/bottom (superior/inferior) contours fill the volume spanned only by the half-slice-thickness towards the other contours, or the full-slice thickness.

In our clinical treatment planning system, interpolation is applied between slices (1), the urethra is considered as part of the prostate (2), and top/bottom contours are extended with a half-slice-thickness (3). It is possible for differently reconstructed shapes of two organs to slightly overlap with each other.

2.3. Original optimization model

In the previous works regarding optimization with MO-RV-GOMEA (Luong *et al* 2018, Bouter *et al* 2019, Maree *et al* 2019, Maree *et al* 2020), the optimization model was based on the clinical protocol used at our clinic at the time. Before undergoing brachytherapy with the planning-aim dose of 13 Gy (1 fraction), patients had received 20 external beam radiation therapy fractions of 2.20 Gy. The following notation is used:

$V^o_{x\%}$: the volume of organ o that receives at least x% of the planning-aim dose.

$D^o_{x\text{cm}^3}$: the lowest dose to the most irradiated x cm^3 of organ o.

The constraints on the EQD2 of the combined treatment, calculated for an α/β-ratio of 3 Gy and a half-time repair of 1.5 h, were $D^{\text{bladder}}_{1\text{ cm}^3} < 78$ Gy, $D^{\text{bladder}}_{2\text{ cm}^3} < 70$ Gy, $D^{\text{rectum}}_{1\text{ cm}^3} < 73$ Gy, $D^{\text{rectum}}_{2\text{ cm}^3} < 70$ Gy, and $D^{\text{urethra}}_{0.1\text{ cm}^3} < 95$ Gy. Based on these constraints, the corresponding planning-aims for brachytherapy are shown in table 1.

There is an inherent trade-off between the planning aims of the targets (i.e. the prostate and the seminal vesicles should receive enough dose) and those of the OARs (i.e. the bladder, the rectum, and the urethra should not receive too much dose). The dose-volume indices of the protocol were therefore combined into two objectives: least coverage index (LCI) and least sparing index (LSI), and a hard optimization constraint C, as follows.

Table 1. The dose-volume indices (planning-aims) and the clinical protocol for HDR prostate brachytherapy used at our center. The unit of each dose-volume index is either percentage of total organ volume for volume indices V, or percentage of planning-aim dose (13 Gy) for dose indices D.

Targets		OARs		
Prostate	Seminal vesicles	Bladder	Rectum	Urethra
$V_{100\%} > 95\%$ $D_{90\%} > 100\%$ $V_{150\%} < 50\%$ $V_{200\%} < 20\%$	$V_{80\%} > 95\%$	$D_{1\text{cm}3} < 86\%$ $D_{2\text{cm}3} < 74\%$	$D_{1\text{cm}3} < 78\%$ $D_{2\text{cm}3} < 74\%$	$D_{0.1\text{cm}3} < 110\%$

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

$$\text{LSI} = \min\{86 - D_{1\ \text{cm}^3}^{\text{bladder}},\ 74 - D_{2\ \text{cm}^3}^{\text{bladder}},\ 78 - D_{1\ \text{cm}^3}^{\text{rectum}},\ 74 - D_{2\ \text{cm}^3}^{\text{rectum}},\ 110 - D_{0.1\ \text{cm}^3}^{\text{urethra}}\}$$

$$C = \min\{50 - V_{150\%}^{\text{prostate}},\ 20 - V_{200\%}^{\text{prostate}}\}.$$

Only plans with a positive C were considered feasible. For feasible plans, the clinical protocol was satisfied if the LCI and LSI were positive.

In order to perform optimization, first a number of dose-calculation points were uniformly randomly sampled in each region of interest (ROI). These points were used to calculate a dose-rate matrix according to the TG-43 formalism (Beaulieu *et al* 2012), describing the dose rate in a water medium from each dwell position to each dose-calculation point. The optimization could then be performed with the evolutionary algorithm MO-RV-GOMEA, which previously showed excellent results for the optimization of the bi-objective model (Luong *et al* 2018) and has recently been GPU-accelerated, allowing dwell time optimization on 100 000 dose-calculation points in 3 min (Bouter *et al* 2019).

In this bi-objective approach, optimization results in a set of plans, each of which has a different trade-off between the LCI and the LSI. This set of plans is called the (approximation) front. The selection of a single treatment plan from this front is left to the physician and can be based on additional information such as tumor characteristics, co-morbidity, and previous treatments.

In the original optimization model, the organ reconstruction settings were based on the standard of our treatment planning system. If multiple sets of manual delineations were available, then the set made at the time of clinical treatment planning was used.

2.4. Robust optimization model

The goal of robust optimization is to take into account multiple organ reconstructions per patient. Let N be the number of organ reconstructions, then there are N combinations of (LCI, LSI, C) values per plan. We defined the robust optimization model as

$$\text{LCI} = \min_{i=1,\ldots,N}\{\text{LCI}_i\},\ \text{LSI} = \min_{i=1,..,N}\{\text{LSI}_i\},\ C = \min_{i=1,\ldots,N}\{C_i\}.$$

A straightforward implementation would be to generate N dose-rate matrices, calculate the dose in all dose-calculation points N times, and compute the LCI, LSI, and C a total of N times. This would lead to approximately N times more computational effort, as calculating the dose in all dose-calculation points is the most time-consuming component of the optimization approach. Therefore, to reduce computation time, advantage was taken of the large overlap between the different organ reconstructions. A single dose-rate matrix was calculated, where for each ROI the dose-calculation points were sampled in the union of all organ reconstructions. In addition, a reconstruction matrix was created, which described for each dose-calculation point to which subset of the N organ reconstructions it belonged (i.e. for which organ reconstructions the point was inside). When evaluating the quality of a treatment plan, the dose in all dose-calculation points was calculated only once, instead of the previously described N times.

For volume indices, the number of dose-calculation points with a dose above a certain threshold inside that volume has to be counted. Therefore, this counting was performed N times, each time based only on the dose-calculation points that belonged to that particular organ reconstruction. For dose indices, dose-calculation points have to be sorted on dose value in descending order. This sorting was performed only once, keeping track of which dose-calculation points belonged to which organ reconstructions. Next, for a given organ reconstruction and dose index $D_{v\text{cm}^3}^{o}$, we take the first n dose-calculation points in organ o that together comprise an equivalent volume of v cm^3, where the equivalent volume of a dose-calculation point that does not belong to organ o in this reconstruction is zero. Then the value of $D_{v\text{cm}^3}^{o}$ is equal to the dose in the nth point.

We extended the GPU-acceleration of the original optimization implementation (Bouter *et al* 2019) to the robust case in CUDA (NVIDIA Corporation, Toolkit v8.0.61). In the parallelization, the dose in all dose-

calculation points of all organ reconstructions of multiple treatment plans was calculated in parallel, exactly as in the original optimization. Moreover, the evaluation of volume and dose indices was performed in parallel, in particular the sorting and counting of dose-calculation points. Code was added to keep track of which dose-calculation points belonged to which organ reconstruction during these operations. Optimization was performed on an NVIDIA Titan Xp, which contained 12 GB of memory.

When applied to organ reconstruction settings, all combinations of the organ reconstruction settings were taken into account, which yielded 8 possible 3D organ reconstructions per patient.

2.5. Experiments

In this section, a distinction is made between the model and the optimization. The original model is the evaluation of a treatment plan which takes into account a single organ reconstruction per patient, the original optimization is the application of MO-RV-GOMEA to the original model. The robust model is the evaluation of a treatment plan which takes into account multiple organ reconstructions per patient, the robust optimization is the application of MO-RV-GOMEA to the robust model.

Both the original optimization approach and the robust optimization approach over organ reconstruction settings were tested on the patient data. Optimization was performed on 100 000 dose-calculation points, and final re-evaluation was performed on 500 000 points to avoid overfitting and obtain more accurate dose-volume indices. The original optimization was run for 3 min. For robust optimization, performing a fixed number of treatment plan evaluations was approximately twice as slow as for the original optimization. A time limit of 6 min was therefore used, as no substantial improvements were found for longer time limits. As MO-RV-GOMEA is a stochastic algorithm, each front that is shown represents the best solutions of 30 runs.

For each treatment plan in an optimization front, the minimum of LCI and LSI can be calculated. For example, if the minimum of LCI and LSI of a treatment plan is 1%, then the $V_{100\%}^{\text{prostate}}$, $V_{80\%}^{\text{vesicles}}$ are both at least 1% higher than the planning aim in the clinical protocol, and the $D_{1\text{ cm}^3}^{\text{bladder}}$, $D_{2\text{ cm}^3}^{\text{bladder}}$, $D_{1\text{ cm}^3}^{\text{rectum}}$, $D_{2\text{ cm}^3}^{\text{rectum}}$, $D_{0.1\text{ cm}^3}^{\text{urethra}}$ are all at least 1% lower than the planning aim in the clinical protocol. For each optimization front, we select the treatment plan for which the minimum of LCI and LSI is the highest, i.e. the treatment plan for which the dose-volume indices are furthest from the planning aims in the clinical protocol (higher for volume indices and lower for dose indices). The minimum of LCI and LSI for this selected plan is called L. To compare the two models, all fronts were re-evaluated in two ways. First, the fronts were re-evaluated in the original model, taking into account a single organ reconstruction per patient. Second, the fronts were re-evaluated in the robust model, taking into account multiple organ reconstructions per patient. This way, the cost and benefit of robust optimization can be calculated by comparing the value L from the original optimization front with the value L from the robust optimization front.

First, we consider the value L as calculated in the original optimization model (i.e. with a single organ reconstruction setting). This value, calculated for the front resulting from the original optimization, can be compared to the value calculated for the front resulting from the robust optimization. The cost of robust optimization was defined as the difference in L between the robust optimization and the original optimization as evaluated in the original model (a negative number).

Second, we consider the value L as calculated in the robust model (i.e. where each dose-volume index is the worst case over all 8 organ reconstruction settings). This value, calculated for the front resulting from the original optimization, can be compared to the value calculated for the front resulting from the robust optimization. The benefit of robust optimization was defined as the difference in L between the robust optimization and the original optimization as evaluated in the robust model (a positive number).

Finally, to study the flexibility of the robust optimization approach, for the 6 patient cases with a second set of delineations, the optimization software was also run over multiple delineations instead of multiple organ reconstructions settings. In both scenarios, the organ reconstruction settings were based on the standard of our treatment planning system. A more in-depth study of robust optimization over delineations has already been performed in recent work (Balvert *et al* 2019).

3. Results

The resulting fronts for the organ reconstruction settings are shown in figure 2. For each patient, four fronts are shown: the results of the original optimization re-evaluated in the original model (original–original) and re-evaluated in the robust model (original–robust), and the results of the robust optimization re-evaluated in the original model (robust–original) and re-evaluated in the robust model (robust–robust). Two main groups of cases can be distinguished from the results. For some patients, all four fronts are close to each other. This indicates that the influence of the organ reconstruction settings was small. For other patients, three fronts are

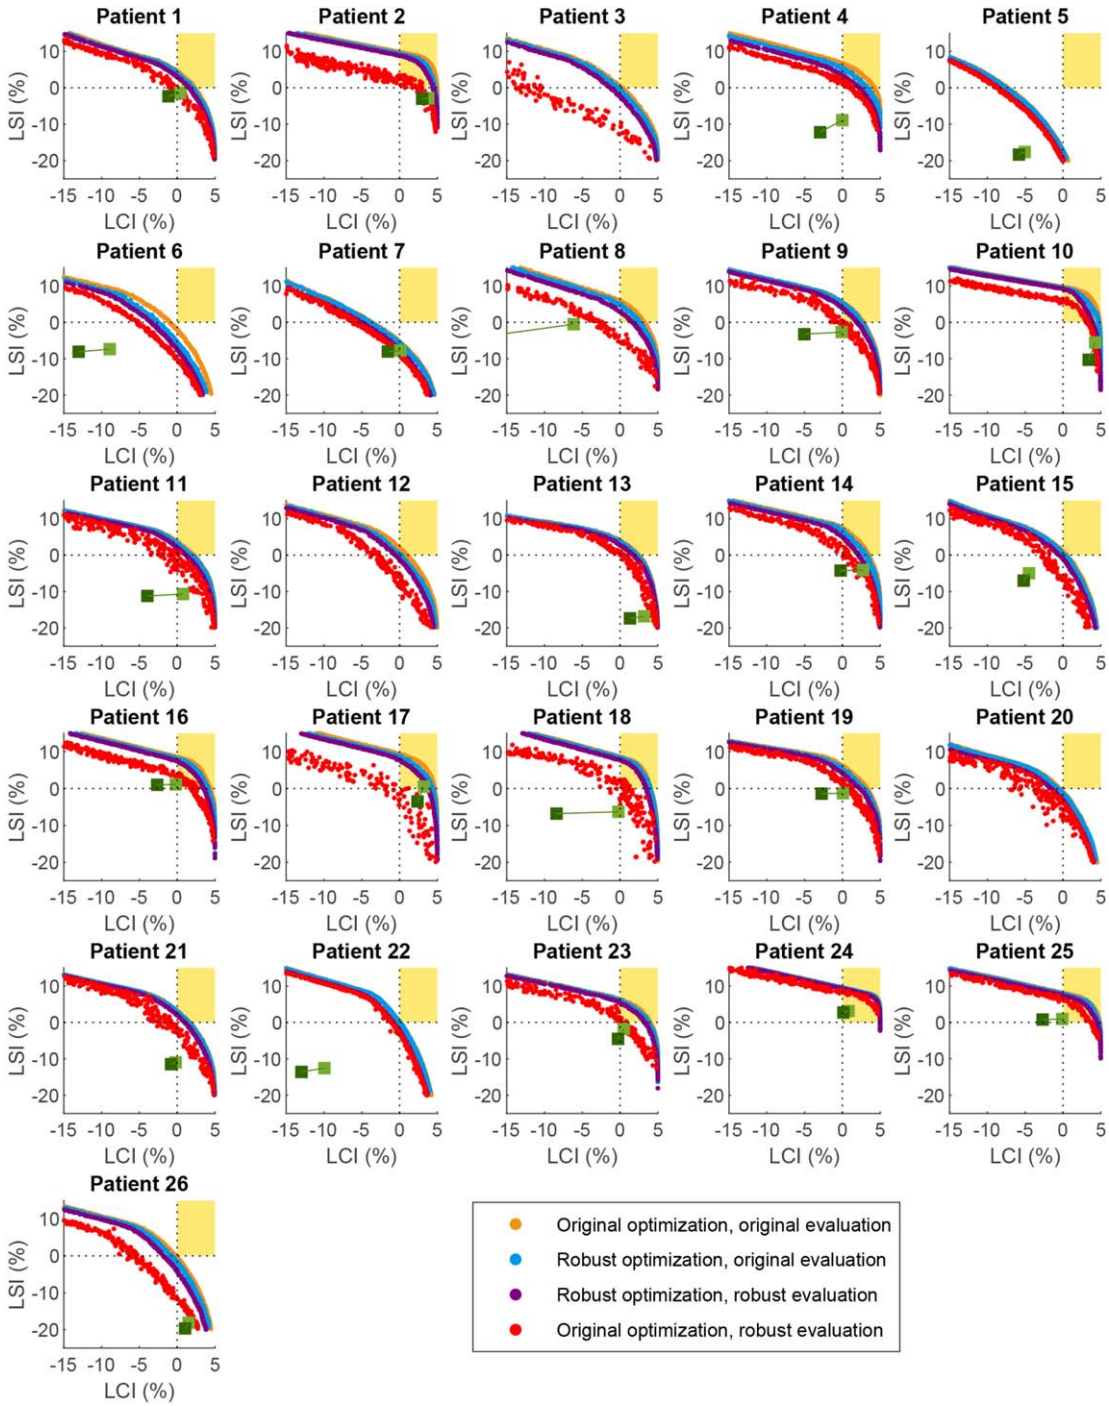

Figure 2. Original and (new) robust bi-objective HDR prostate brachytherapy planning results over organ reconstruction settings for 26 patient cases. All plans, optimized for either the original or robust bi-objective model, are re-evaluated in the original and the robust bi-objective model. Plans in the so-called Golden Corner where the LCI and LSI are both larger than zero satisfy the clinical protocol. The green squares show the clinically used treatment plan, evaluated in the original (light green) and the robust model (dark green, connected by a line for improved visualization). For some patients (3, 12, and 20), the clinical plan was out of range of the figure.

close to each other, and the original–robust front is significantly lower. This indicates that the influence of the organ reconstruction settings was large, and the original optimization was sensitive to it.

The corresponding costs (i.e. the difference between robust–original and original–original) and benefits (i.e. the difference between robust–robust and original–robust) of robust optimization are shown in figure 3. The median cost for robust optimization as observed in the original model was −0.25% with a range of [−0.01%, −1.03%]. The median benefit of robust optimization as observed in the robust model was 0.93% with a range of [0.19%, 4.16%]. For 4 patients (patient 3, 8, 12, and 15), selected plans that appeared to be good when optimized in the original model, violated the clinical protocol with more than 1% when considering different settings. This was not the case for robustly optimized plans.

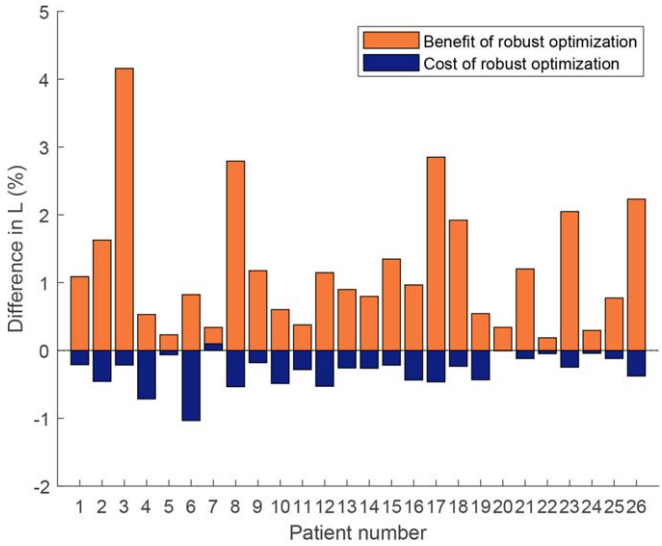

Figure 3. The cost and benefit of robust optimization determined by comparing selected plans from the original and robust bi-objective HDR prostate brachytherapy planning. The cost is defined as the difference between the two plans, evaluated in the original optimization model. The benefit is defined as the difference between the two plans, evaluated in the robust optimization model. For patient 7, the cost is positive, which indicates that the selected plan from the original planning was suboptimal (caused by the stochastic nature of the optimization).

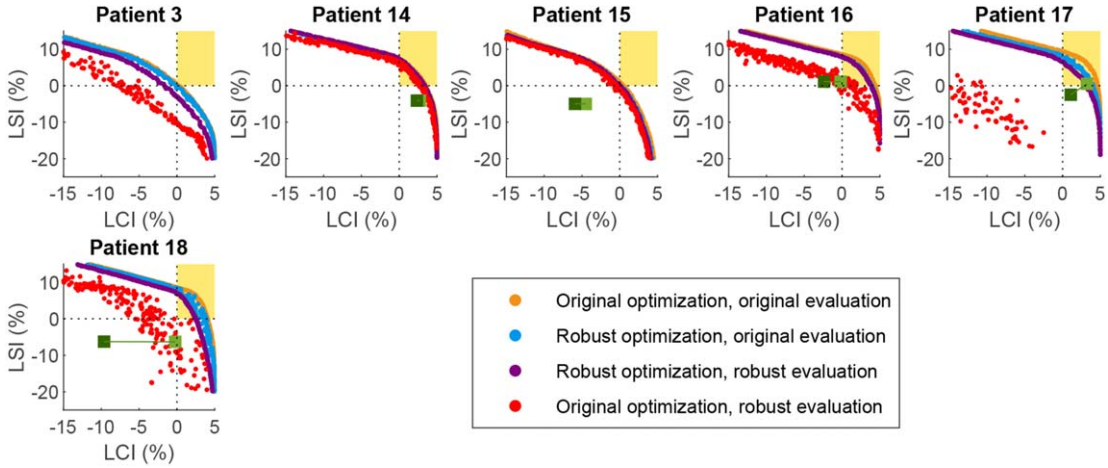

Figure 4. To study the flexibility of the robust optimization approach, for the 6 patient cases for which a second set of delineations was available, the optimization software was also run over multiple delineations instead of multiple organ reconstruction settings. The green squares show the clinically used treatment plan, evaluated in the original (light green) and the robust model (dark green). For patient 3, the clinical plan was out of range of the figure.

The resulting fronts for the multiple organ delineations are shown in figure 4. They show that the technique of robust optimization is more generally applicable than only for the uncertainty of the organ reconstruction settings, since it works for delineation uncertainties as well. The largest benefit of robust optimization is observed for patient 17, mainly due to differences in the delineation of the base of the seminal vesicles.

4. Discussion

We studied robust dwell time optimization for HDR prostate brachytherapy over multiple organ reconstruction settings, including its costs and benefits. For 5 patients, the benefit was larger than 2%, ranging up to 4.2% (which would be 0.5 Gy for a dose index). Hence, for some patients, the clinical benefit will be large. For all patients, the (absolute value of the) cost was at most 1%.

Patients with a larger clinical benefit are those for whom the dose-volume indices are more sensitive to organ reconstruction settings. The indications for a large sensitivity in the dose-volume indices to organ reconstruction settings were studied in previous work (van der Meer *et al* 2019). In particular, a small volume of

the seminal vesicles in combination with a large surface of the top and/or bottom contour, as well as a high amount of radiation at the bottom contour of the bladder, were shown to have a large influence.

For robust optimization on 100 000 dose-calculation points on 8 organ reconstruction settings, the required run time increased from 3 to 6 min. Although the run time increases, the impact on the total treatment time of a patient is minor. If the number of settings further increases, then so does the run time, for multiple reasons. First, the dose-volume indices of a treatment plan have to be calculated for each of the settings. On a single Titan Xp GPU, having already parallelized the calculations of multiple treatment plans and their dose-calculation points, the run time increases approximately linearly with the number of settings. Second, the lower the amount of overlap in the organ reconstructions based on different settings, the more dose-calculation points might be needed, further increasing the run time.

The benefit (and the cost) of robust optimization considering organ reconstructions depends on the influence of the organ reconstruction settings on the dose-volume indices. The use of a smaller slice thickness, combined with delineation of more contours, would lead to a smaller influence of the organ reconstructions settings and would therefore reduce the benefit (and the cost) of robust optimization. However, this would also increase the time and effort required for the manual delineations. With respect to organ intersections, it would be preferable to achieve a consensus within the brachytherapy community on whether or not the urethra should be included in the volume of the prostate, since this is more a question of definition than reconstruction.

The approach of robust optimization is more general than this optimization method (MO-RV-GOMEA). It has been shown to be flexible, working not only for organ reconstruction settings but also for multiple contour delineations. It may be even more beneficial in combination with the recent advancements in automatic organ segmentation algorithms. Extensions to changes in organ positions and/or shape (for example due to swelling) would be interesting as well, but would require more data on the scenarios to be considered.

5. Conclusions

Robust optimization generated plans of high quality, irrespective of organ reconstruction settings. With the limited effect on total runtime that follows from our approach, robust optimization therefore offers a practically viable way to account for dosimetry uncertainties that result from choices in organ reconstruction settings.

Acknowledgments

This work is part of the research program IPPSI-TA with project number 628.006.003, which is financed by the Dutch Research Council (NWO) and Elekta AB (Stockholm, Sweden). The authors gratefully acknowledge the support of the NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

Conflicts of interest

This work was partly funded by Elekta AB (Stockholm, Sweden). Dr Alderliesten, Dr Bel, Prof. Bosman, and Dr Pieters are leaders of projects supported by Elekta. Y Niatsetski is an Elekta employee.

References

Balvert M, den Hertog D and Hoffmann A L 2019 Robust optimization of dose-volume metrics for prostate HDR-brachytherapy incorporating target and OAR volume delineation uncertainties *INFORMS J. Comput.* **31** 100–14

Beaulieu L *et al* 2012 Report of the Task Group 186 on model-based dose calculation methods in brachytherapy beyond the TG-43 formalism: current status and recommendations for clinical implementation *Med. Phys.* **39** 6208–36

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

Deist T M and Gorissen B L 2016 High-dose-rate prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing *Phys. Med. Biol.* **61** 1155–70

Fredriksson A and Bokrantz R 2014 A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning *Med. Phys.* **41** 81701

Karabis A, Belotti P and Baltas D 2009 Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming *World Congress on Medical Physics and Biomedical Engineering* (Berlin: Springer) pp 612–5

Kirisits C *et al* 2007 Accuracy of volume and DVH parameters determined with different brachytherapy treatment planning systems *Radiother. Oncol.* **84** 290–7

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 2018 Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment *Swarm Evol. Comput.* **40** 37–52

Maree S C et al 2019 Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy—a retrospective observer study *Brachytherapy* **18** 396–403

Maree S C et al 2020 Automatic bi-objective parameter tuning for inverse planning of high-dose-rate prostate brachytherapy *Phys. Med. Biol.* **65** 75009

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

Unkelbach J et al 2018 Robust radiotherapy planning *Phys. Med. Biol.* **63** 22TR02

van der Meer M C et al 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