



Robust Evolutionary Bi-objective Optimization for Prostate Cancer Treatment with High-Dose-Rate Brachytherapy

Marjolein C. van der Meer¹(✉), Arjan Bel¹, Yury Niatsetski²,
Tanja Alderliesten³, Bradley R. Pieters¹, and Peter A. N. Bosman⁴

¹ Department of Radiation Oncology, Amsterdam UMC, University of Amsterdam,
Amsterdam, The Netherlands

`marjolein.vandermeer@amsterdamumc.nl`

² Physics and Advanced Development, Elekta, Veenendaal, The Netherlands

³ Department of Radiation Oncology, Leiden University Medical Center,
Leiden, The Netherlands

⁴ Life Sciences and Health Research Group, Centrum Wiskunde & Informatica,
Amsterdam, The Netherlands

Abstract. We address the real-world problem of automating the design of high-quality prostate cancer treatment plans in case of high-dose-rate brachytherapy, a form of internal radiotherapy. For this, recently a bi-objective real-valued problem formulation was introduced. With a GPU parallelization of the Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA), good treatment plans were found in clinically acceptable running times. However, optimizing a treatment plan and delivering it to the patient in practice is a two-stage decision process and involves a number of uncertainties. Firstly, there is uncertainty in the identified organ boundaries due to the limited resolution of the medical images. Secondly, the treatment involves placing catheters inside the patient, which always end up (slightly) different from what was optimized. An important factor is therefore the robustness of the final treatment plan to these uncertainties. In this work, we show how we can extend the evolutionary optimization approach to find robust plans using multiple scenarios without linearly increasing the amount of required computation effort, as well as how to deal with these uncertainties efficiently when taking into account the sequential decision-making moments. The performance is tested on three real-world patient cases. We find that MO-RV-GOMEA is equally well capable of solving the more complex robust problem formulation, resulting in a more realistic reflection of the treatment plan qualities.

Keywords: Evolutionary Algorithms · Robust optimization · Multi-objective optimization · Empirical study · Radiation oncology

Supported by Elekta, Sweden.

© Springer Nature Switzerland AG 2020
T. Bäck et al. (Eds.): PPSN 2020, LNCS 12270, pp. 441–453, 2020.
https://doi.org/10.1007/978-3-030-58115-2_31

1 Introduction

Brachytherapy is a form of internal radiotherapy that can be used for treating prostate cancer. The treatment involves intraoperative placement of a number of very thin needles, called catheters, inside the patient, for a radioactive source to be moved through. The entire procedure is a two-stage sequential decision-making process. The first part is determining how to place the catheters, after which the catheters are actually placed. The second part is determining how to move the radioactive source through the catheters, which is performed after catheter placement. Each catheter has a fixed set of positions, called dwell positions, where the radioactive source can pause for certain amounts of time, called dwell times. The workflow is illustrated in Fig. 1.




Fig. 1. The simulation workflow used to study robust optimization. Purple blocks indicate decision-making stages. (Color figure online)

Key quality indicators can be formulated for treatment plans, enabling optimization to support decision making. For the first part, both catheter position variables and dwell time variables play a role. For the second part, only the dwell time variables still play a role. On the one hand, enough catheters should be placed to ensure a good treatment. On the other hand, a larger number of catheters increases the risk of complications for the patient. The number of catheters is therefore an important part of the catheter position optimization. Since catheter placement is performed in the operating room, no changes to the catheters can be made afterwards. The result of catheter position optimization (that includes dwell times) should therefore be representative of what can be achieved in the dwell time optimization.

Optimizing brachytherapy is difficult, for multiple reasons. Firstly, the problem is inherently multi-objective, due to the trade-off between radiation dose to the tumor (which you want to maximize) and to the surrounding tissue (which you want to minimize). Secondly, the objective functions comprising the multi-objective problem are such that there is no gradient information. Thirdly, there are time constraints on the optimization due to the patient waiting for treatment. For solving difficult multi-objective problems, Evolutionary Algorithms (EAs) are the state-of-the-art [3]. Previous work on bi-objective optimization for both dwell times [2] and catheter positions [11] has shown promising results.

Since all problem variables in the second stage are already part of the first stage, arguably there is no need for a two-stage optimization process. After a single optimization, both decisions can be made and executed. However, in practice such a one-shot approach is not sufficient, because brachytherapy involves

a number of uncertainties. Firstly, there is uncertainty in the identified organ boundaries due to the limited resolution of the medical images used for treatment planning. Secondly, catheters always end up (slightly) different in the patient from what was planned. Therefore, to ensure the best possible plans, the dwell times should be re-optimized after actual catheter placement has taken place. To avoid overly optimistic catheter position optimization fronts, that may lead to the wrong conclusion about how many catheters are needed for a particular patient in the first decision phase, these uncertainties should be taken into account in the optimization problem. Since the most time consuming part of the optimization is the calculation of the objective functions of a treatment plan, straightforwardly applying robust optimization would result in clinically infeasible run times.

In this work, we will introduce robust optimization to the full workflow for high-dose-rate prostate brachytherapy, while still keeping the run times low. The aim is for the optimization fronts to be representative of what can be achieved in clinical practice. Specifically, the catheter position optimization fronts should be representative of what can be achieved later in the dwell time optimization. We will evaluate the run time of the optimization, as well as the robustness of the resulting treatment plans.

2 Background

2.1 Insightful Decision Support via Bi-objective Optimization

The ultimate goal is to obtain the highest quality treatment plan to be used for the dose delivery. There are several key evaluation criteria that can be mathematically formulated, enabling the use of optimization methods. For an in-depth explanation of all details involved, we refer the interested reader to related literature [11]. Here, we briefly summarize the most important concepts.

In clinical practice, the evaluation of a treatment plan is based on a clinical protocol, which describes how much radiation the prostate and seminal vesicles should receive as part of the treatment, as well as how much dose is maximally allowed to the surrounding healthy organs to avoid complications. This radiation dose that is prescribed for the prostate is called the *planning-aim dose*. The clinical protocol of the Amsterdam UMC is formulated in terms of so-called dose-volume indices. There are two types of dose-volume indices; volume indices and dose indices. A volume index V_x^o is the volume of organ o that receives at least $x\%$ of the planning-aim dose. A dose index D_x^o is the lowest dose to the most irradiated $x\text{cm}^3$ of organ o .

Single-objective optimization approaches are often based on a simplified version of the clinical protocol [7, 8]. All objectives following from the simplified protocol are combined into a single optimization function by the weighted-sum approach. As a result, optimized treatment plans often require manual improvements by the medical planners [4], which is a time-consuming and little insightful process. Alternatively, optimizing for all these indices would entail solving a

many-objective optimization problem, for which the results are not straightforward to interpret, visualize, and use for decision making.

For this reason, two grouped objectives were defined that have proven effective and insightful for clinical practice [9]. The resulting bi-objective optimization model is based directly on the clinical protocol. Dose-volume index criteria related to the dose coverage of the prostate and the seminal vesicles are combined into the Least Coverage Index (LCI). Criteria related to the sparing of organs at risk, namely rectum, bladder, and urethra, are combined into the Least Sparing Index (LSI). Upper bounds to the amount of radiation the prostate can receive also fall under the LSI. Since it is unknown a priori how to weight the different dose-volume index criteria, the objectives are constructed by combining the criteria in a worst-case manner, which was observed to be much related to how plans are manually improved in clinical practice. This results in the following optimization objectives:

$$\begin{aligned} \text{LCI} &= \min \left\{ V_{100\%}^{\text{prostate}} - 95, V_{80\%}^{\text{vesicles}} - 95 \right\}, \\ \text{LSI} &= \min \left\{ 86 - D_{1\text{cm}^3}^{\text{bladder}}, 74 - D_{2\text{cm}^3}^{\text{bladder}}, 78 - D_{1\text{cm}^3}^{\text{rectum}}, 74 - D_{2\text{cm}^3}^{\text{rectum}}, \right. \\ &\quad \left. 110 - D_{0.1\text{cm}^3}^{\text{urethra}}, 50 - V_{150\%}^{\text{prostate}}, 20 - V_{200\%}^{\text{prostate}} \right\}. \end{aligned} \quad (1)$$

For catheter position optimization, an additional constraint on the healthy tissue immediately surrounding the prostate is necessary [11]. This constraint is based on the number of catheters N :

$$C = \left\{ \begin{array}{ll} V_{200\%}^{\text{healthy tissue}} - 0.125N, & \text{for } \text{LSI} \geq -25 \\ V_{200\%}^{\text{healthy tissue}} - 0.125N \left(1 + \frac{-25 - \text{LSI}}{100} \right), & \text{for } \text{LSI} < -25 \end{array} \right\} \leq 0. \quad (2)$$

The result of solving this problem is a trade-off curve of treatment plans that, when satisfying $\text{LCI} > 0$ and $\text{LSI} > 0$ (and $C \leq 0$) adhere to the clinical protocol. Visualizing this makes the most important trade-offs immediately insightful, as well as whether the clinical protocol can be achieved.

2.2 Problem Variables

As mentioned in the introduction, there are two decision phases. In the first phase, the catheter positions need to be optimized. To use the model of Sect. 2.1, we also need to set the dwell times per catheter. Hence, in the first phase, all catheter positions are optimized at the same time as the dwell times pertaining to these catheters. In the second phase, the catheter positions are fixed, and only the dwell times are to be optimized.

For catheter position optimization, the number of catheters is given as input. Moreover, constraints are added to the optimization model describing which catheter positions are feasible. Catheters have to be inside either the prostate (with a -1 mm margin) or the seminal vesicles. Catheters are not allowed to intersect with either rectum or urethra (both with a 1 mm margin). Finally, the distance between the surfaces of each pair of catheters has to be at least 1 mm. For an in-depth explanation of all details involved, we refer the interested reader to related literature [11].

2.3 Evolutionary Optimization

In a comparison between different EAs, the best performing EA for the problem at hand was the Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA) [9]. A key reason is that MO-RV-GOMEA is capable of exploiting gray-box settings where problem-specific enhancements can be readily applied. Specifically, MO-RV-GOMEA makes use of so-called *partial evaluations*. Instead of changing all variables of a potential solution and then performing an evaluation, the variables are changed in multiple steps, and after each step an evaluation is performed. If the solution was improved, the change is kept; if not, the change is reverted. For brachytherapy, these many evaluations can be done efficiently, because the impact of changes to certain dwell times can be computed by considering radiation originating from the corresponding dwell positions only [9]. A similar argument holds for catheter positions [11].

A second reason for its enhanced performance is that MO-RV-GOMEA models the dependencies between variables by using a so-called *Linkage Tree* (LT). At the bottom of this tree, each of the variables is in a singleton set. Higher up, sets are merged together based on the strength of the dependencies between their variables. At the top of the LT, all sets have been merged, resulting in a single set containing all variables. Combining all sets in the LT results in the so-called *Family Of Subsets* (FOS). During optimization, all FOS elements are considered. For every FOS element, a joint Gaussian distribution is estimated, based on a selection of best solutions. Such a Gaussian distribution is known to work well when there is no gradient information and the fitness landscape may not be smooth everywhere, e.g., it is adopted by the state-of-the-art real-valued EA known as CMA-ES [6]. The variables that are in one FOS element are then resampled together based on this distribution. This way, dependencies between variables are taken into account. When applied to brachytherapy, the dependencies between variables are modelled based on the distances between the dwell positions [9, 11].

3 Accounting for Uncertainties via Robust Optimization

3.1 Organ Reconstructions: A Problem-Specific Solution

While the dose-volume indices of a treatment plan are theoretically uniquely defined, computing values for these indices in practice is not. A key reason is that dose-volume indices are computed from 3D (organ) volumes. However, medical scans are usually sets of 2D images. An algorithm is then used to reconstruct the 3D organ shapes from delineations, performed on the 2D images. Due to the limited resolution of the medical images, such a reconstruction is not uniquely defined and differs from one clinical system to another. A solution is to perform robust optimization over different organ reconstructions, to avoid overfitting on one particular reconstruction. Three organ reconstruction settings have previously been studied, for details, see [10]. Combinations of these settings yield 8 possible 3D organ reconstructions per patient. Hence, there are 8 combinations

of (LCI, LSI, C) values per plan. Taking again a worst-case scenario approach to combining objective values (in this case defined for different reconstruction settings), the robust optimization model is defined as

$$\text{LCI} = \min_{i=1,\dots,8} \{\text{LCI}_i\}, \quad \text{LSI} = \min_{i=1,\dots,8} \{\text{LSI}_i\}, \quad \text{C} = \max_{i=1,\dots,8} \{\text{C}_i\}. \quad (3)$$

This model is identical for both dwell time optimization and catheter position optimization. A straightforward implementation would be to compute the LCI, LSI, and C 8 times. This would lead to approximately 8 times more computational effort, as calculating the objective values associated with a treatment plan is the most time-consuming component in the EA. With runtime being important for clinical usability, reducing this additional runtime is important. To do so, advantage is taken of the large volume overlap between different organ reconstructions (i.e., it is at the borders that organ reconstructions differ, not at the interiors). When evaluating the quality of a treatment plan, the dose in each overlapping part of the patient in all scenarios is calculated only once. The parts that do not overlap are small, and evaluated separately, for each reconstruction. After this, the dose-volume indices are calculated 8 times. As a result, performing a fixed number of evaluations is only approximately twice as slow as the original optimization.

3.2 Catheter Displacements: An EA Generic Solution

When catheters are placed inside the patient, they always end up (slightly) different from what was planned. Accounting for this uncertainty requires taking into account the fact that the complete workflow is a sequential decision-making process. Between catheter position optimization and dwell time optimization, there is the catheter placement which causes the uncertainty. We will simulate actual catheter placement by randomly displacing all catheters by 1 mm, where the 1 mm is based on discussions with a clinical expert. After the displacements, dwell times are re-optimized, but catheter positions are fixed.

If these displacements are not taken into account, catheter position optimization fronts will be overly optimistic compared to the dwell time optimization fronts, because optimization will overfit on the one scenario in which catheters are not displaced at all. Hence, a lower number of catheters will appear to be sufficient than is really the case. As a result, optimization will be an ineffective decision support tool because likely not enough catheters would be placed in the patient to ensure a good treatment.

To avoid this, the random catheter displacements should thus be taken into account in the optimization. The most straightforward approach to do so correctly would be to consider many catheter displacements (in the order of 100) each time a set of catheter positions is evaluated. For each catheter displacement, dwell times would be separately optimized, to take into account that dwell time optimization is performed again after the catheters are displaced. Unfortunately, this would be prohibitively computationally expensive, because one full dwell-time optimization takes about 30 s [2] and we have only a few minutes to decide catheter positions in clinical practice.

Alternatively, when a set of catheter positions is evaluated, dwell times could be kept fixed when displacing catheters. This is a conservative lower bound on the real evaluation, since this disregards the dwell time optimization performed after catheter placement. This would therefore result in overly pessimistic catheter position optimization fronts, which is also undesirable from the perspective of a clinical decision support tool. As a result, too many catheters would be placed in the patient, which would increase the risk of complications. Moreover, due to the many scenarios, this approach would still be too computationally expensive.

We therefore propose a third approach that is generic to multi-objective EAs in sequential decision-making processes under uncertainty where the uncertainty between stages involves variable realization (i.e., realizing the actual optimized catheter positions in the clinic). When evaluating a treatment plan for multiple catheter displacements, dwell times are still kept fixed when displacing catheters. However, such evaluations are only used to frequently filter the solutions in the elitist archive. Specifically, every generation, Algorithm 1 is used; outside of the elitist archive, no robust evaluations over catheter displacements are performed. This way, some robustness of the treatment plans to catheter displacements is taken into account, without resulting in too optimistic/pessimistic fronts or clinically infeasible running times. It should be said that in MO-RV-GOMEA, the elitist archive plays a role in providing parent solutions, so the impact of only filtering the elitist archive for robustness this way is potentially larger than for other EAs that employ elitist archives.

Algorithm 1: Filtering of the elitist archive

```

1 Let  $(X, Y)$  be the representation of a catheter position.
2 Make a backup of the catheter positions in the elitist archive.
3 for  $m=1, \dots, 100$  do
4   for all catheters  $i$  do
5     Sample  $\theta$  uniformly in  $[0, 2\pi]$ .
6     for all solutions  $j$  in the elitist archive do
7        $X(j) += 1\text{mm} \cdot \cos(\theta)$ .
8        $Y(j) += 1\text{mm} \cdot \sin(\theta)$ .
9       Apply boundary repair if necessary.
10    end
11  end
12  Evaluate the elitist archive.
13  Restore the backup of the catheter positions in the elitist archive.
14 end
15 for all solutions  $j$  in the elitist archive do
16   Determine the Nadir point of  $j$  of the 100 evaluations.
17 end
18 Filter the elitist archive based on the Nadir points: solutions for which its Nadir
    point is dominated by the Nadir point of another solution, are removed from
    the elitist archive.

```

4 Experiments

In our experiments, we simulate the workflow of clinical practice, including the sequential decision-making steps. The goal is to see whether with our new approach, good plans can still be obtained and, possibly more importantly, whether the predicted quality of plans in the first stage is a realistic representation of plans obtained in the second stage. If so, a properly informed decision can be made about the number of catheters to use for a particular patient.

A simulation of the workflow starts with catheter position optimization. The running time is limited to 15 min. After catheter position optimization, a single treatment plan is selected from the front, with the highest quality in terms of $\min\{\text{LCI}, \text{LSI}\}$. This quality is defined as L :

$$L := \max_{\text{plans } j \text{ in front}} \{\min\{\text{LCI}_j, \text{LSI}_j\}\}. \quad (4)$$

Subsequently, dwell time optimization is performed again separately. The running time is limited to 6 min for the original dwell time optimization, and 15 min for the robust dwell time optimization. We use larger runtimes here than strictly needed in clinical practice because we want to observe also the convergence properties of the EA. Each simulated workflow is applied to the data of 3 patient cases for 16, 10, and 4 catheters. Due to the randomness in the EA and the catheter displacements, 10 runs are performed of each simulated workflow.

Three approaches are compared. Catheter position optimization is always followed by catheter displacements and robust dwell time optimization over organ reconstructions. The first approach uses the original catheter position optimization, where no uncertainties are considered at all. The second approach uses the robust catheter position optimization over only organ reconstructions. The third approach uses the robust catheter position optimization over both organ reconstructions and catheter displacements, using elitist archive filtering as in Algorithm 1.

The difference between the results of catheter position optimization and dwell time optimization is tested with a paired samples t-test on L for each of the patients, numbers of catheters, and versions of catheter position optimization separately, whereby the difference was considered to be statistically significant if $p < 0.00185$. This includes a Bonferroni correction for 27 test, i.e., $p < 0.05/27$.

To study the convergence of MO-RV-GOMEA, we use the well-known hypervolume metric [12], i.e., the area in the bi-objective space that is covered by the front and a so-called *reference point*. Here, we choose the reference point $(-30, -30)$ and only consider solutions in the front dominating this point.

For all code, a GPU-acceleration was implemented in CUDA (NVIDIA Corporation, Toolkit v8.0.61), based on previous work [2]. Optimization was performed on an NVIDIA Titan Xp, which contained 12 GB of memory.

5 Results

The results of the original catheter position optimization are shown in Fig. 2. The part of the objective space where all aims in the clinical protocol are

satisfied, i.e. $LCI > 0$ and $LSI > 0$, is highlighted. The influence magnitude of the uncertainties depends on the patient and the number of catheters. Except for patient 2 with 16 catheters, there is a statistically significant difference between the catheter position and dwell time fronts. This means that the catheter position fronts are not realistic, as they are higher than what is obtained when taking into account the uncertainties. This shows the need for robust optimization.

The results of robust catheter position optimization over only organ reconstruction settings are shown in Fig. 3. For 16 catheters, for all patients, there is no statistically significant difference between the catheter position and dwell time fronts. Hence, in these cases, only robust optimization over organ reconstruction settings is needed. It should be noted that the catheter position fronts have dropped towards the dwell time fronts, but the dwell time fronts themselves did not improve. Hence, taking into account organ reconstruction settings during catheter position optimization results in more realistic fronts, but not necessarily in better catheter positions.

For patient 1 with 4 catheters, and for patient 2 with 10 and 4 catheters, there is still a statistically significant difference between the catheter position and dwell time fronts. Hence, in these cases, taking into account only organ reconstruction settings is not sufficient to also obtain robustness to catheter displacements. This shows the need for robust optimization over both uncertainties.

The results of robust catheter position optimization over both organ reconstruction settings and catheter displacements are shown in Fig. 4. For all patients and numbers of catheters, there is no statistically significant difference between the catheter position and dwell time fronts. Hence, the catheter position fronts are now realistic. It can be seen that with 16 catheters, for all patients, plans exist that satisfy all constraints in the clinical protocol (i.e., $LCI > 0$ and $LSI > 0$). This is sometimes the case for 10 catheters, and never for 4 catheters.

Finally, the hypervolume values of the fronts of the three types of catheter position optimization over time are shown in Fig. 5. For these patients, even with robust optimization, convergence is still achieved quickly, indicating that in clinical practice we may very well use only 5 min instead of 15, which is clinically acceptable.

6 Discussion

In this paper, we introduced preclinical work on robust optimization for high-dose-rate prostate brachytherapy. By performing robust optimization over both organ reconstruction settings and catheter displacements, the catheter position and dwell time fronts obtained in the first and second stages of the sequential decision-making process become virtually the same. Hence, decisions based on the catheter position fronts are now more representative of the resulting dwell time fronts. The larger part of this robustness appears to be due to the robust optimization over organ reconstructions, rather than over catheter displacements. It should still be studied whether this also holds for different patients and numbers of catheters.

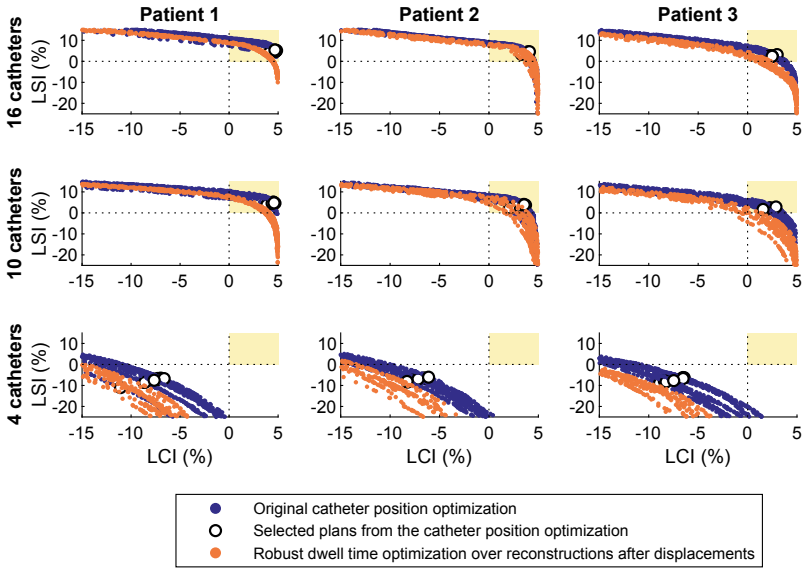


Fig. 2. The original (non-robust) catheter position optimization (blue), for selected plans (white circles) followed by catheter displacements and robust dwell time optimization over organ reconstruction settings (orange). Ten runs are shown. (Color figure online)




Fig. 3. The robust catheter position optimization over organ reconstruction settings (blue), for selected plans (white circles) followed by catheter displacements and robust dwell time optimization over organ reconstruction settings (orange). Ten runs are shown. (Color figure online)

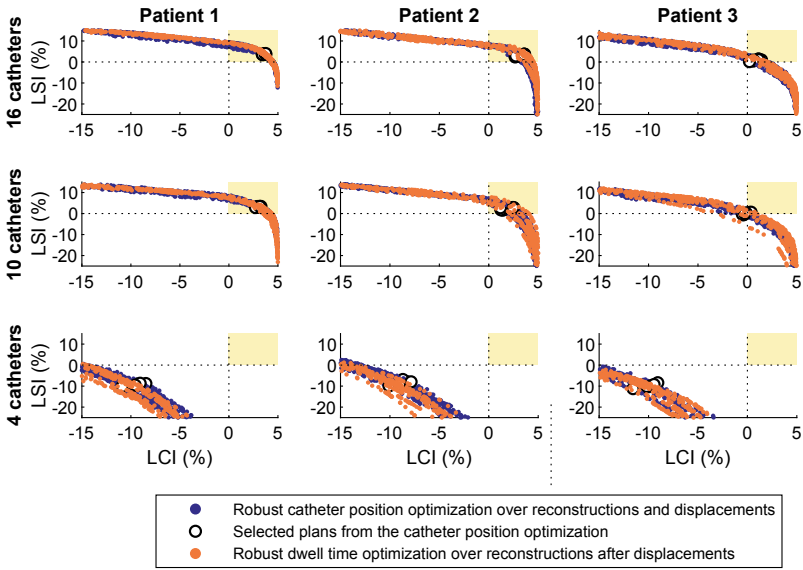


Fig. 4. The robust catheter position optimization over organ reconstruction settings and catheter displacements (blue), for selected plans (white circles) followed by catheter displacements and robust dwell time optimization over organ reconstruction settings (orange). Ten runs are shown. (Color figure online)

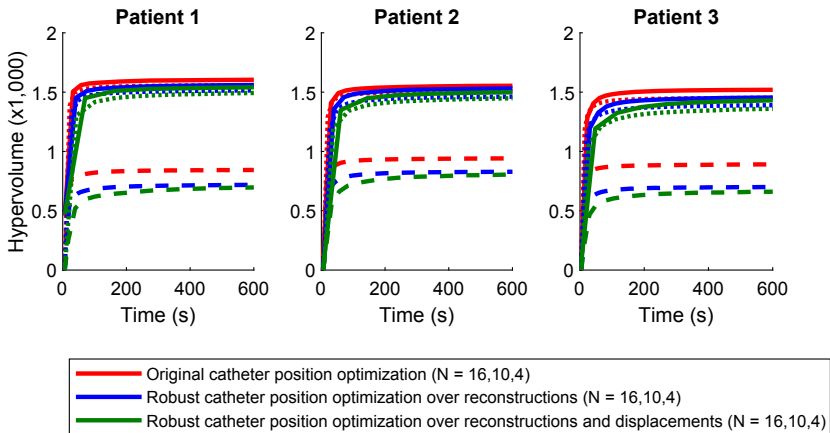


Fig. 5. The hypervolume values of the fronts of the three types of catheter position optimization (colors) over time. For all patients and numbers of catheters (16 is solid, 10 is dotted, 4 is dashed), the average over ten runs is shown. (Color figure online)

In the simulated workflow, only a single set of catheter positions was selected from each catheter position front. After random catheter displacements, dwell time optimization was sufficient to obtain a front of equally good plans again. Combined with the fact that different catheter position configurations are indeed obtained along the front of the first-stage optimization, this suggests that the problem of robust catheter position optimization itself is highly redundant, e.g. due to many (almost) equally good local optima. Arguably, positioning itself could be considered to be single-objective: the objective of maximizing $\min\{\text{LCI}, \text{LSI}\}$ would have been sufficient. However, a more in-depth analysis with physicians is needed of the different catheter position configurations that are obtained to see if there are any other reasons to deviate from this.

The proposed techniques for robust optimization are more general than this optimization method (MO-RV-GOMEA) or these uncertainties (organ reconstruction settings and catheter displacements). Besides the generality of the elitist archive filtering for sequential multi-objective decision making under uncertainty, it is for instance likely that the technique for re-using intersections of organs will also work for different uncertainties related to organ shape and catheter positions (such as uncertainties in delineations [1] and catheter angles [5]). This could be explored in future work.

7 Conclusion

We showed how a recently introduced state-of-the-art evolutionary bi-objective optimization approach for high-dose-rate prostate brachytherapy can be extended to include robust optimization, without requiring a prohibitively large running time when optimized with MO-RV-GOMEA. Two types of uncertainty were considered: one with a fixed set of scenario's, and one with a stochastic component. Using a different approach for each type of uncertainty, both were included directly in the optimization. The results show that more realistic fronts of catheter position optimization can now be obtained. This way, the optimization can be used more reliably in clinical practice as a basis for making such important clinical decisions as how many catheters to use for a particular patient and where to place them. Moreover, additional insights into the optimization can now be obtained. Specifically, we have learned that a promising approach that may well improve run time further may be to robustly optimize catheter positions single-objectively, by optimizing the minimum of the two objectives in the original optimization model.

Acknowledgements. This work is part of the research program IPPSI-TA, which has project number 628.006.003 and is financially supported by the Netherlands 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.

References

1. Balvert, M., den Hertog, D., Hoffmann, A.L.: Robust optimization of dose-volume metrics for prostate HDR-brachytherapy incorporating target and OAR volume delineation uncertainties. *INFORMS J. Comput.* **31**(1), 100–114 (2019)
2. Bouter, A., Alderliesten, T., Pieters, B.R., Bel, A., Niatsetski, Y., Bosman, P.A.N.: GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy. *Med. Phys.* **46**(9), 3776–3787 (2019)
3. Deb, K.: Multi-objective optimisation using evolutionary algorithms: an introduction. In: Wang, L., Ng, A., Deb, K. (eds.) *Multi-objective Evolutionary Optimisation for Product Design and Manufacturing*, pp. 3–34. Springer, London (2011). https://doi.org/10.1007/978-0-85729-652-8_1
4. Dinkla, A.M., et al.: A comparison of inverse optimization algorithms for HDR/PDR prostate brachytherapy treatment planning. *Brachytherapy* **14**(2), 279–288 (2015)
5. Gorissen, B.L.: Practical robust optimization techniques and improved inverse planning of HDR brachytherapy. Tilburg University, School of Economics and Management, Tech. rep. (2014)
6. Hansen, N., Müller, S.D., Koumoutsakos, P.: Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES). *Evol. Comput.* **11**(1), 1–18 (2003)
7. Karabis, A., Belotti, P., Baltas, D.: Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming. In: Dössel, O., Schlegel, W.C. (eds.) *World Congress on Medical Physics and Biomedical Engineering*, pp. 612–615. Springer, Heidelberg (2009). <https://doi.org/10.1007/978-3-642-03474-9>
8. Lessard, E., Pouliot, J.: Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Med. Phys.* **28**(5), 773–779 (2001)
9. Luong, N.H., Alderliesten, T., Bel, A., Niatsetski, Y., Bosman, P.A.N.: Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment. *Swarm Evol. Comput.* **40**, 37–52 (2018)
10. van der Meer, M.C., et al.: Sensitivity of dose-volume indices to computation settings in high-dose-rate prostate brachytherapy treatment plan evaluation. *J. Appl. Clin. Med. Phys.* **20**(4), 66–74 (2019)
11. van der Meer, M.C., Pieters, B.R., Niatsetski, Y., Alderliesten, T., Bel, A., Bosman, P.A.N.: Better and faster catheter position optimization in HDR brachytherapy for prostate cancer using multi-objective real-valued GOMEA. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1387–1394 (2018)
12. Zitzler, E., Thiele, L.: Multiobjective optimization using evolutionary algorithms— a comparative case study. In: Eiben, A.E., Bäck, T., Schoenauer, M., Schwefel, H.-P. (eds.) *PPSN 1998. LNCS*, vol. 1498, pp. 292–301. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0056872>