

Better and Faster Catheter Position Optimization in HDR Brachytherapy for Prostate Cancer using Multi-Objective Real-Valued GOMEA

Marjolein C. van der Meer
Academic Medical Center
Amsterdam, The Netherlands
marjolein.vandermeer@amc.uva.nl

Bradley R. Pieters
Academic Medical Center
Amsterdam, The Netherlands
b.r.pieters@amc.uva.nl

Yury Niatsetski
Elekta
Veenendaal, The Netherlands
yury.niatsetski@elekta.com

Tanja Alderliesten
Academic Medical Center
Amsterdam, The Netherlands
t.alderliesten@amc.uva.nl

Arjan Bel
Academic Medical Center
Amsterdam, The Netherlands
a.bel@amc.uva.nl

Peter A.N. Bosman
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
peter.bosman@cw.nl

ABSTRACT

The recently-introduced Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) family has been shown to be capable of excellent performance on academic benchmark problems, outperforming other state-of-the-art EAs, especially when efficient partial evaluations are possible. This holds true also for the latest extension, the Multi-Objective Real-Valued GOMEA (MO-RV-GOMEA). In this paper, we apply MO-RV-GOMEA to the real-world multi-objective optimization problem of catheter placement in High-Dose-Rate (HDR) brachytherapy for prostate cancer, a problem that is non-trivial to solve and has high real-world importance and relevance. Due to the underlying geometric structure of the real-valued variables, partial evaluations can be performed, allowing MO-RV-GOMEA to exploit this structure. The performance of MO-RV-GOMEA is tested on three real-world patient cases and compared to a recent state-of-the-art mixed-integer EA that is aimed at a restricted version of the problem. We find that with MO-RV-GOMEA better solutions can be found much faster, making our proposed approach much more realistic to be used in clinical practice, and enabling new insights into both catheter placement for prostate brachytherapy and on objectives used for automated treatment planning. First results indicate that richer problem models are needed to better match real-world clinical preferences.

CCS CONCEPTS

• **Applied computing** → **Health informatics**; • **Mathematics of computing** → *Evolutionary algorithms*;

KEYWORDS

Medicine, Evolutionary algorithms, Multi-objective optimization, Empirical study

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '18, July 15–19, 2018, Kyoto, Japan

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5618-3/18/07...\$15.00

<https://doi.org/10.1145/3205455.3205505>

ACM Reference Format:

Marjolein C. van der Meer, Bradley R. Pieters, Yury Niatsetski, Tanja Alderliesten, Arjan Bel, and Peter A.N. Bosman. 2018. Better and Faster Catheter Position Optimization in HDR Brachytherapy for Prostate Cancer using Multi-Objective Real-Valued GOMEA. In *GECCO '18: Genetic and Evolutionary Computation Conference, July 15–19, 2018, Kyoto, Japan*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3205455.3205505>

1 INTRODUCTION

Brachytherapy is a form of internal radiotherapy used for treating prostate cancer, the second most common type of cancer among men worldwide [14]. The radiation, although targeted at tumor cells, poses a risk for healthy organs around the prostate. Therefore, it is important to construct a treatment plan with the best possible trade-off between radiation to the targets and radiation to organs at risk (OARs).

In brachytherapy, a radioactive source is moved into or close to the tumor. For prostate brachytherapy, this is achieved by placing a number of catheters into the prostate through the area between the scrotum and the anus. The radioactive source is then moved through each catheter and is paused at fixed positions in the catheters, called dwell positions, for certain amounts of time, called dwell times. The longer the dwell time at a dwell position, the more radiation is delivered to the surrounding volume. By varying the dwell times, different treatment plans are possible. Finding a clinically acceptable treatment plan (such as shown in Figure 1) is a complex task, and involves many clinical requirements that need to be satisfied, which are formulated in a clinical protocol.

In order to search for a good treatment plan, an optimization problem can be formulated, see e.g., [3, 7–9]. Because there are multiple factors of importance, however, it is not trivial to design a formulation and propose a solver that together lead to clinically desirable results. At the hospital involved in this study, clinically available software collapses the multiple requirements in the clinical protocol into a single-objective optimization problem. Often, subsequently time needs to be spent manually adjusting the resulting treatment plan until satisfaction.

The quality of the treatment plans achievable with the optimization of dwell times depends on the placement of the catheters. In general, if more catheters are used, then better treatment plans are possible. However, to minimize the possibility of complications, one

should use as few catheters as possible [4, 5, 13]. Even for a given number of catheters, there are still many placements possible and for each placement many treatment plans are possible, resulting in a large search space. At the hospital involved in this study, typically 16-18 catheters are used. The placement is done mostly based on clinical expertise (such as [10]).

Simultaneous optimization of catheter positions and dwell times could lead to better treatment plans. Multiple works already exist on this, e.g., [2, 6], using single-objective optimizers. Recently, a fully multi-objective approach, the Multi-Objective Genetic Algorithm for Model-Based mixed-Integer opTimization (MO-GAMBIT), has been applied to this optimization problem [11]. Three objectives were considered, namely number of catheters, coverage of targets, and sparing of OARs, and promising results were found. However, the optimization model was restricted in the sense that only pre-determined catheter positions could be used. Moreover, ultimately, the required running times were too long for use in clinical practice. In this work, we extend the problem formulation to allow more freedom in the positioning of catheters. Instead of a fixed set of pre-determined catheter positions, catheters can take any position within a certain domain. This changes the problem from mixed integer to fully real valued. While the search space is factually larger, it is also important to reduce the required running time. We therefore consider the Multi-Objective Real-Valued Gene-Pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA, [1]), which has recently been shown to achieve good results sufficiently fast on the optimization of a treatment plan *after* catheter placement [9]. In this paper, we extend the application of MO-RV-GOMEA to tackle the brachytherapy problem including catheter positioning, with the aim of finding better treatment plans with lower running time than was possible with MO-GAMBIT.




Figure 1: Graphical representation of the anatomy and implanted catheters of patient 2 (left), with the isodose lines (lines on which all points receive equal radiation dose) corresponding to the clinically accepted treatment plan for this patient on an MRI scan (right).

2 MO-RV-GOMEA

The Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA, [12]) has in recent years been shown on various academic benchmark problems, as well as real-world problems, to be an efficient and effective optimizer, especially in the case where certain problem-specific knowledge can be exploited, which is the case in many real-world problems. The algorithm has been extended to problems with multiple objectives and real-world variables in MO-RV-GOMEA

[1]. An important characteristic of GOMEA is that it exploits the dependencies between variables in order to obtain offspring solutions in an efficient way. These dependencies can be learned from the population, or be determined problem-specifically and hard-coded into GOMEA. In the latter case, the efficiency of GOMEA is typically increased.

The performance of GOMEA can in particular be improved substantially if partial evaluations are possible. In GOMEA, offspring solutions are created by iteratively improving on each solution in the population with consecutive local changes. After each local change, an evaluation is performed, checking for improvement. If no improvement was obtained, the local change is reversed; otherwise, it is maintained. If the change in objectives resulting from a few changed variables can be computed faster than if all variables were changed, i.e., if a partial evaluation is faster than a full evaluation, then the evaluation times in GOMEA are reduced.

3 BRACHYTHERAPY

The first step in brachytherapy is the implantation of a number of catheters into the prostate, based on live ultrasound images. These catheters are implanted through the patient’s perineum skin (between scrotum and anus) based on the geometry of the involved organs, clinical experience (such as [10]), and optionally a preplan made in clinical software. At the hospital involved in this study, typically 16-18 catheters are used. The catheter implant will determine the quality of the possible treatment plans.

Once the catheters have been placed, Magnetic Resonance Imaging (MRI) scans are made on which the catheters and the relevant tissues are fully visible. Not only the prostate, but also the seminal vesicles are deemed targets, as the vesicles could also contain tumor cells. Other organs which are nearby, i.e., bladder, rectum, and the urethra passing through the prostate, are deemed OARs and should be irradiated as little as possible. Using clinical software, the catheters, targets, and OARs are delineated. This defines the dwell positions, i.e., the positions in the catheters at which the radioactive source can pause for a given dwell time. Only dwell positions within 5mm of prostate and vesicles, and at least 1mm away from the urethra are used.

The next step is the optimization of dwell times. A trivial treatment plan, in which dwell times take on very large values, gives enough radiation to the targets but irradiates the OARs too much. On the other hand, if all dwell times are zero, then the OARs are fully spared, but no radiation is delivered to the targets either. The difficulty lies in finding a treatment plan that results in both enough coverage of targets and enough sparing of OARs, while at the same time also adhering to several other criteria.

3.1 Plan evaluation

At the start of a High-Dose-Rate (HDR) prostate brachytherapy treatment, a radiation dose is prescribed for the prostate, called the *planning aim dose*. For the patients involved in this study, the planning aim dose was 13 Gy. The seminal vesicles are prescribed to receive at least 80% of this 13 Gy. For the OARs, there are upper bounds on the amount of radiation they can receive. Moreover, there are also upper bounds on the amount of radiation that the prostate itself can receive.

A guideline in the plan evaluation is the clinical protocol, which is however not absolutely strict. The current clinical protocol at the hospital involved in this study is formulated in terms of so-called Dose-Volume Indices (DVIs). A DVI describes how large the volume of an organ covered by a certain dose is. There are two types of DVIs: volume and dose indices. Volume indices, i.e., the sub volume of an organ that receives at least / at most a specific dose, are useful for describing the amount of radiation delivered to targets. Dose indices, i.e., the lowest dose to the most irradiated sub volume of a certain size of an organ, are useful for describing the amount of radiation delivered to OARs. We use the following notation: V_x^o is the volume of organ o that receives at least $x\%$ of the planning aim dose; D_x^o is the lowest dose to the most irradiated $x\text{cm}^3$ of organ o . A full description of the clinical protocol in terms of DVIs at the hospital involved in this study is presented in Table 1.

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

Table 1: Brachytherapy treatment planning in clinical practice at the involved hospital. Volume indices V are in percentage of the total organ volume, dose indices D are in percentage of the planning aim dose of 13 Gy.

In the clinical software, the computation of DVIs is performed by calculating the dose in a number of points in the organ of interest. These points are called *dose-calculation points*. The more points are used, the more precise the resulting values of the DVIs are. The strength of the radioactive source, in combination with the position of a dose-calculation point relative to a dwell position, determines the *dose rate*, which is the amount of radiation that is delivered from the source to the dose-calculation point per second, in Gy/s. Multiplying the dose rate with the dwell time gives the absolute amount of radiation that is delivered, in Gy. For a given catheter, the total radiation that each dose-calculation point receives from that catheter is the combined dose delivered from all the dwell positions in that catheter. Finally, the total radiation at each dose-calculation point is obtained by summing the dose from each separate catheter.

Let D be the set of all dose-calculation points. Let C be the set of all catheters in a treatment plan. For every $c \in C$, let T_c be the set of all dwell positions of that catheter. With a certain source strength, R_c is a matrix where entry (i, j) indicates the dose rate associated with dwell position $j \in T_c$ and dose-calculation point $i \in D$. Let \vec{t}_c be the vector of dwell times at all dwell positions in catheter c . The vector \vec{d} of the amounts of radiation received at all dose-calculation points can be computed as

$$\vec{d} = \sum_{c \in C} R_c \vec{t}_c. \quad (1)$$

Hence, the computation of the dose in the dose-calculation points requires for each catheter c a dose-rate matrix R_c in combination with a vector of dwell times \vec{t}_c .

3.2 Clinical software

Current clinical treatment planning software can quickly provide a treatment plan. This is done by solving a simplified optimization

model of the problem with local search methods. The simplification is used because it is difficult to directly optimize DVIs, due to their discrete nature. For example, maximizing the $V_{100\%}^{\text{prostate}}$ means maximizing the number of dose-calculation points inside the prostate that receives at least the planning aim dose. These simplifications however typically make it hard to ensure that all criteria are satisfied and that the treatment plan proposed by the software is the best possible one for the patient case at hand.

The optimization in the clinical software is single-objective. All objectives following from the clinical protocol are combined into a single optimization function by the weighted-sum approach. For each setting of the weighting-coefficient vector, there is a single optimal solution. The proper setting of these weights for a given desired trade-off in tumor coverage versus organ sparing is patient-specific and difficult to determine a priori, resulting in the need to run multiple trial-and-error runs of the optimization.

Often, the resulting treatment plan is not immediately approved by the physicians, in which case this proposed plan has to be manually improved by medical planners. An improvement on one criterion can however result in a deterioration of another criterion, making this a time-consuming and little insightful process.

The plan is adjusted until the physicians are satisfied. In some cases, no treatment plan can be found that satisfies all clinical criteria. This depends on the catheter implant and the geometry of the involved organs. In this case, the physicians need to decide which criteria are more important to be satisfied and which criteria can be compromised. This leads to a final approved treatment plan, which might (slightly) violate some criteria in the clinical protocol.

3.3 Research software

Brachytherapy treatment planning is intrinsically a multi-objective optimization problem, trading off target coverage and sparing of OARs. By formulating treatment planning as such, and using an a posteriori multi-objective approach, the decision of the desired trade-off for a specific patient can be determined by the physician by insightfully selecting a plan from the obtained set of plans, i.e., the Pareto front. A clinically relevant multi-objective formulation combined with a multi-objective optimizer that is capable of dealing with nonlinearities and the lack of gradient information such as in the counting-based DVIs, can therefore be very helpful for clinical practice. Recent results show that with MO-RV-GOMEA excellent results can indeed be obtained for the case where catheters are already in place using a bi-objective optimization model [9].

In principle, each criterion in the protocol can be taken to be an objective. However, solving the resulting many-objective optimization model does not necessarily readily lead to insightful results. Moreover, although superior results have been reported for MO-RV-GOMEA on multi-objective optimization problems with few objectives [1], it is currently unknown if this performance scales well with the number of objectives. For this reason, a bi-objective formulation was constructed by identifying that the clinical protocol in terms of DVIs in Table 1 consists of two types of criteria. The first type is related to the coverage of the targets, namely $V_{100\%}^{\text{prostate}}$ and $V_{80\%}^{\text{vesicles}}$. The larger the coverage of the prostate and the vesicles, the better the treatment plan. Hence, for these criteria large DVI values are preferable. This can be achieved by increasing the dwell

times. For this first type of criteria, the optimum would be 100% coverage for both prostate and vesicles.

The second type of DVI criteria is related to the sparing of OARs, namely the DVIs of rectum, urethra, and bladder. The lower the dose to the OARs, the better the treatment plan, so small DVIs are preferable, which can be achieved by decreasing the dwell times. The optimum would be 0% for all these criteria.

The two objectives in the bi-objective problem formulation correspond to the two conflicting types of coverage criteria and sparing criteria. Specifically, nine DVI criteria in the clinical protocol are combined in the *Least Coverage Index* (LCI), which corresponds to the worst DVI in the coverage criteria, and the *Least Safe Index* (LSI), which corresponds to the worst DVI in the sparing criteria.

In this paper, the problem formulation is changed in two ways, in order to avoid that all treatment plans with a small number of catheters are infeasible.

Firstly, the criteria on the $V_{150\%}^{prostate}$ and $V_{200\%}^{prostate}$ are included in the LSI, instead of set as hard constraints. For a number of catheters for which treatment plans exist that satisfy all clinical criteria, the criteria on the $V_{150\%}^{prostate}$ and $V_{200\%}^{prostate}$ are usually easily satisfied, making the difference between the two problem formulations small. The resulting objectives are shown below.

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

$$LSI = \min\{50 - V_{150\%}^{prostate}, 20 - V_{200\%}^{prostate}, 86 - D_{1cm^3}^{bladder}, 74 - D_{2cm^3}^{bladder}, 78 - D_{1cm^3}^{rectum}, 74 - D_{2cm^3}^{rectum}, 110 - D_{0.1cm^3}^{urethra}\} / 100$$

Secondly, the feasible search space consists of all treatment plans with at least 75% coverage of both prostate and vesicles, with the LSI unconstrained (instead of a hard constraint of $LSI \geq -0.2$).

4 CATHETER POSITION OPTIMIZATION

For catheter position optimization, three objectives are considered, namely *Number of Catheters* (NC), coverage of targets, and sparing of OARs. Optimizing all three objectives at the same time is difficult because the Pareto-optimal catheter configurations for a given NC are not necessarily related to the catheter configurations for $NC + 1$. The optimization is therefore sliced, running a single instance of MO-RV-GOMEA for each NC. The larger NC, the larger the search space, and the better the theoretically optimal Pareto front. We consider a range for NC of 1 to 30. In this section, we outline the key changes and required additions compared to [9] to include the catheter positions into the optimization.

Feasible catheter configurations. For every patient case, a vector \vec{v} is determined which describes the angle at which catheters can be placed in the patient. This vector is used in two ways.

Firstly, a two-dimensional grid of 90 parallel catheter configurations is placed on prostate and vesicles combined, where the direction of each catheter is equal to \vec{v} . This is the largest possible grid that can be placed in the clinical software. An example of such a grid is shown in Figure 2. The clinical software then provides information on these catheter configurations and their corresponding dwell positions, which is used to calculate the dose-rate matrix of each catheter in the grid. Moreover, for each catheter, it is calculated which dwell positions require penetration of the bladder.




Figure 2: Graphical representation of the grid of catheter positions of patient 2.

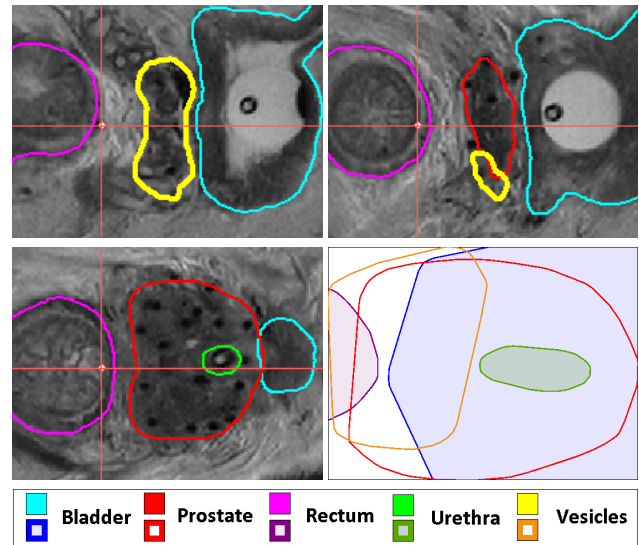


Figure 3: Three slices of an MRI scan with delineated contours of patient 2 at different heights, together with the two-dimensional projections of all organs (bottom right figure), projected along the direction of a catheter.

Secondly, the vector \vec{v} is used to create the projections of all organs. For this, a fixed plane called the *projection plane* is selected, which is parallel to the planes in which the organ contours were delineated. Next, the contours of all organs are projected along the direction of \vec{v} onto this projection plane. Finally, for each organ, the convex hull of the projected contours is taken, resulting in one contour as the projection of this organ. An example is shown in Figure 3.

During optimization, catheters can in principle take all positions of the catheters in the grid, but also all positions in between grid catheters. However, there are two constraints on the positions and depths of insertion.

Firstly, catheters are not allowed to intersect with either rectum or urethra. In the optimization a slightly stronger constraint is used, which is that catheters are not allowed to intersect with the projections of the rectum and urethra. Since all catheters are parallel to each other, the position of a catheter can be described as a point on the projection plane. This way, the intersection of a catheter with rectum or urethra can be calculated as intersection of a point with

a two-dimensional contour, which is a relatively cheap operation. To keep catheters within the vicinity of the targets, catheters have to intersect with the projection of either prostate or vesicles.

Secondly, catheters are only allowed to use dwell positions up to the bladder. This corresponds to the clinical practice of inserting catheters only as far as the bladder. Given a catheter position, the four catheters surrounding this catheter in the grid are determined. Only dwell positions in the current catheter which do not require intersection with the bladder in all the surrounding four catheters are allowed to contribute to the dose distribution. The dwell times of other dwell positions are treated as being zero.

Initialization. The catheters are initialized in feasible positions by first sampling random points inside the union of the projections of prostate and vesicles, outside of the projections of rectum and urethra. Next, the positions are divided into NC clusters using the k-means clustering algorithm. For each individual in the population, the position of the i th catheter $i \in \{1, \dots, NC\}$ is initialized as a position from cluster i . This way, the different catheters in two individuals are ordered similarly.

Dwell positions which require intersection with the bladder in the surrounding four grid catheters are initialized with zero dwell time. Dwell positions outside the bladder are initialized with a random real number in the range $[0, W_{NC}]$, where W_{NC} is an upper bound depending on NC. The values of W_{NC} were determined empirically such that the initial population contains a few feasible treatment plans with different amounts of coverage, i.e., different values of $LCI \in [-0.2, 0.05]$.

Evaluation. The formulas for evaluating the DVIs of a treatment plan can be found in [9]. The protocol evaluation requires a dose-rate matrix for each catheter, which is calculated through interpolation. Given a catheter position, again the four catheters surrounding this catheter in the grid are determined. For a given dwell position in the current catheter and a given dose-calculation point, the entry in the dose-rate matrix is calculated using bilinear interpolation between the corresponding entries in the dose-rate matrices of the four surrounding catheters. This is faster than directly calculating the dose-rate matrix, and results in only a small loss of precision. Interpolated dose-rate matrices are recalculated for each evaluation.

Partial evaluations are possible and compatible with MO-RV-GOMEA. For catheter position optimization, there are two cases. If a changed variable is a dwell time, then the catheter containing the corresponding dwell position is determined. For the dose-rate matrix of this catheter, only the entries corresponding to the dwell positions have to be computed. If a changed variable is a catheter-position variable, then two full dose-rate matrices of this catheter have to be computed, containing all entries. The first dose-rate matrix corresponds to the previous catheter position, the second dose-rate matrix corresponds to the current catheter position. By multiplying both matrices with the dwell times, the change in dose for each dose-calculation point can be computed.

Variable dependencies. Dependencies between variables are usually computed either a priori from problem-specific information, or online from the population. Here, we compute the dependencies between variables every generation, using both the population and problem-specific knowledge. Because we have two types of

variables (catheter positions and dwell times), in each generation, the dependency information between variables is based on two different distance matrices. Each matrix uses the average catheter positions in this generation.

The first matrix contains the distance between variables that describe the positions of the NC catheters. The distance between two catheter-position variables x_1 and x_2 is the following. If x_1 and x_2 together describe the position of a single catheter, then the distance is zero. If x_1 and x_2 belong to two different catheters, then the distance is equal to the physical distance between these two catheter points in the projection plane.

The second matrix contains the distance between dwell-time variables. The distance between two dwell-time variables is equal to the physical distance between the corresponding dwell positions.

The matrices are separately used to describe the dependencies between variables of the corresponding type.

5 EXPERIMENTS

Clinical data from three treated patients is used for optimization. In order to allow a comparison with GAMBIT, the patient data used in [11] was requested and used as the patient cases. We follow the approach used for GAMBIT regarding the number of dose-calculation points. The final Pareto fronts are evaluated with 100,000 randomly chosen dose-calculation points. To speed up optimization, 20,000 of those are used for fitness evaluation.

For each patient, a separate run of GOMEA is performed for each of the numbers of catheters $NC \in \{1, 2, \dots, 9, 10, 12, 14, \dots, 28, 30\}$, each performing bi-objective optimization. The base population size is 300, the base number of clusters is 5. The larger the number of catheters, the more difficult the problem, resulting in longer required running times. The duration of optimization of each GOMEA instance is therefore limited to $1800 + 900 * NC$ seconds (e.g., 45 minutes for 1 catheter and 8 hours for 30 catheters). Each Pareto front, that we show, represents the best solutions of 30 runs (i.e., the Pareto front of Pareto fronts; for the empirical standard deviation of MO-RV-GOMEA, see [9]). Experiments of up to and including 10 catheters were run on an Intel(R) Xeon(R) processor E5-2630v4 (2.2 GHz), experiments of more than 10 catheters were run on an Intel(R) Xeon(R) processor E5-2699v4 (2.2-3.0 GHz).

The results are compared to the results obtained with GAMBIT in [11]. These results were reported to have been obtained with a duration of optimization limited to 48 hours per catheter number. Also in this paper, each Pareto front represents the best solutions of 30 runs. The experiments were reportedly run on two different types of processors, an AMD Opteron(tm) Processor 6386 SE and an Intel(R) Xeon(R) CPU E5-2699v4, the slowest of which is approximately twice as slow as the processors used in our experiments.

Moreover, we consider also the treatment plans that were clinically used for these patients, which all used 16 catheters. For optimization of dwell times with a fixed catheter placement using MO-RV-GOMEA, run times of 1 hour for 16 catheters were previously shown to be sufficient [9]. To assess the potential of the clinically placed catheters, for each patient MO-RV-GOMEA is run for 1 hour to optimize the dwell times for the clinically used catheter implant, resulting in a single Pareto front. In this case, the base population size is 30, the base number of clusters is 3. Each Pareto front,

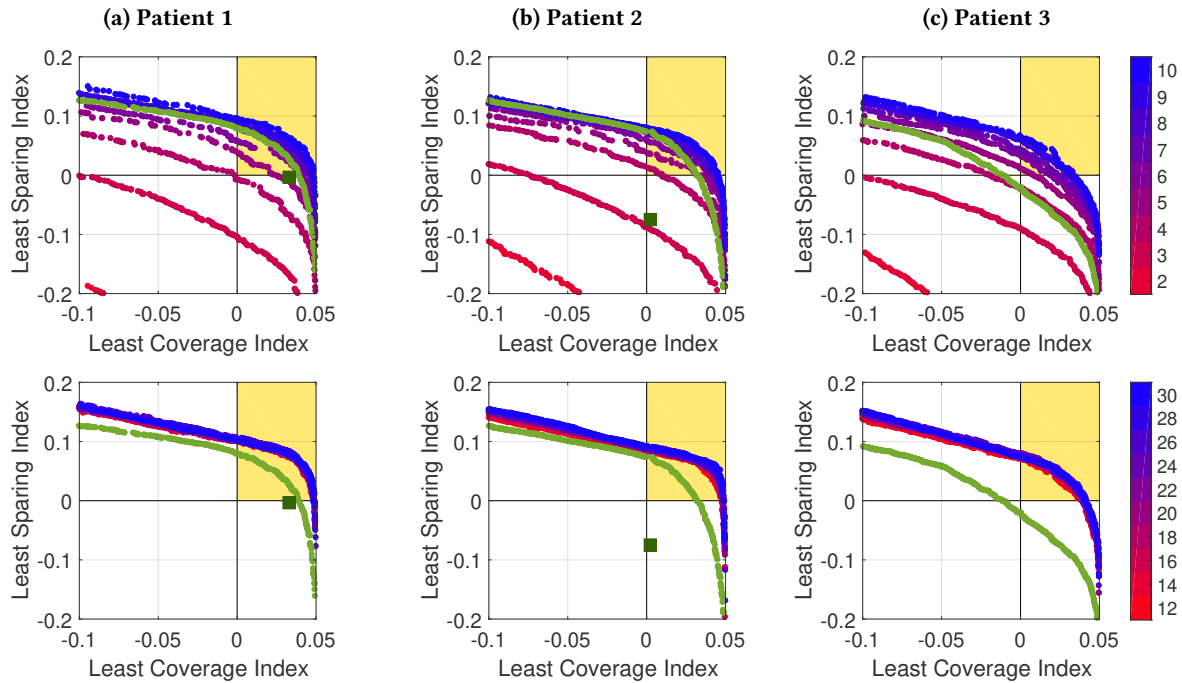


Figure 4: The Pareto fronts of different numbers of catheters (top row: 1-10 catheters, bottom row: 12-30 catheters). For each patient the Pareto front of 1 catheter is below the left bottom corner of the figure, all Pareto fronts with 12 catheters or more go through the so-called Golden Corner. The green square shows the clinically used treatment plan with 16 catheters (which has a Least Coverage Index below -0.1 for patient 3), the green dots show the Pareto front of the optimization of dwell times on these catheters. (Color versions of graphs available online.)

that we show, again represents the best solutions of 30 runs. The experiments were run on the Intel(R) Xeon(R) processor E5-2630v4 (2.2 GHz).

6 RESULTS

The Pareto fronts for $NC \in \{1, 2, \dots, 9, 10, 12, 14, \dots, 28, 30\}$ are shown in Figure 4. Treatment plans which satisfy all clinical criteria are visualized as being in the so-called *Golden Corner* where $LCI > 0$ and $LSI > 0$. For each of the three patients, the Golden Corner can be reached. This requires 5, 4, and 5 catheters for patients 1, 2, and 3, respectively. For comparison, the Pareto fronts presented in [11] required 7, 7, and 14 catheters, respectively.

The treatment plans that were clinically used for these patients are all outside the Golden Corner. The clinically used plan for patient 3 is even outside the range of the plots of the Pareto fronts. For this patient, the coverage of the vesicles is only 72.9%, resulting in a low LCI value. For patient 1 and 2, the clinical treatment plans are both dominated by a Pareto front of a number of catheters less than the 16 used in the clinical plan. The minimum number of catheters for which the clinical plan is dominated, is 6 for patient 1 and 4 for patient 2.

The Pareto fronts obtained by optimizing the dwell times of the clinically used catheters using MO-RV-GOMEA all dominate the clinically used plan, in-line with recent results reported in [9]. For patient 1, the clinically used treatment plan is close to the results of MO-RV-GOMEA. For patient 2, optimizing the dwell times resulted

in treatment plans much closer to the Pareto front of 16 catheters for which the positions were optimized. For patient 3, optimizing the dwell times also gave a large improvement, but the treatment plans obtainable with this clinical catheter implant are still far from the Pareto front of 16 catheters for which the positions were optimized.

In general, it can be seen that the more catheters are used, the better the obtainable treatment plans, as is to be expected. Also, the improvement obtainable by adding a catheter reduces as NC increases. However, the quality of the Pareto front of a given NC and the improvement obtainable by adding a catheter are also patient specific. For example, the Pareto front of 10 catheters for patient 2 is better than for patient 3. This could be due to the fact that patient 3 has a large urethra and a small prostate, making the volume of the urethra relatively large, complicating covering the prostate while still sparing the urethra enough.

Based on the run time limits and the fact that the results with MO-RV-GOMEA are also better within that time span, the improved Pareto fronts were obtained with a speed-up of at least a factor 3. The lower the number of catheters, the faster the results could be obtained, showing a possible additional speed-up for less than 30 catheters.

A few selected catheter positions present in the Pareto fronts are shown in Figure 5. For a given patient and NC, all plans in the Pareto front actually have very similar catheter configurations, and hence they differ mostly in dwell times. Moreover, for a small

NC, different patients appear to have similar patterns in catheter configurations, such as a rectangle for four catheters.

The catheter configurations also show differences between patients. For example, for patient 3, it can be seen that covering the vesicles enough requires catheters to pass closely to the rectum, which is less the case for patient 2. Patient 3 however requires more catheters to pass through the vesicles, which has a larger volume than for patient 2.

Finally, Figure 6 shows selected treatment plans in the Golden Corner. The DVIs were re-evaluated in the clinical software used at the hospital involved in this study, and all satisfy the criteria in the clinical protocol.

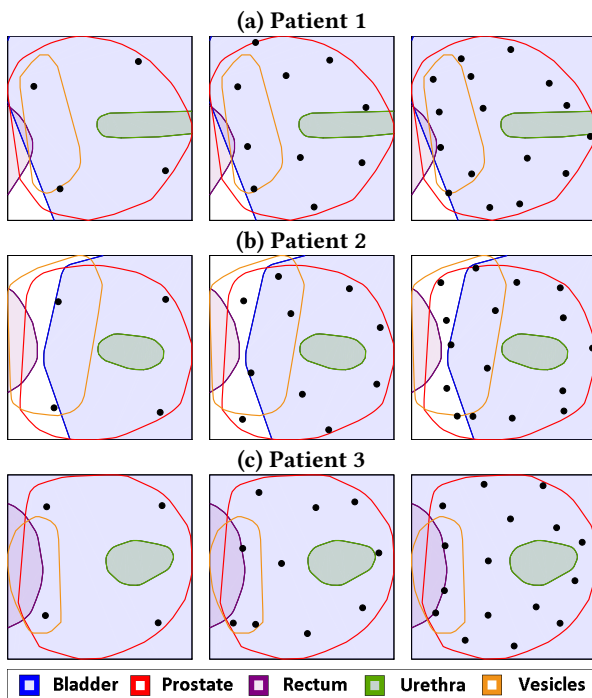


Figure 5: For each patient, for 4, 10, and 16 catheters, the catheter positions of a selected plan inside the Golden Corner from the Pareto front are shown.

7 DISCUSSION

A few selected treatment plans were shown to a physician at the hospital involved in this study. These treatment plans were regarded to be promising, and the catheter positions as feasible. However, there was also room for highly desirable improvement. One of the main factors is related to the problem formulation, which is based on the current clinical protocol. The clinical protocol is formulated based on the delineated organs. However, there is normal tissue in between and around these organs that is also being irradiated during prostate brachytherapy. Since there is no criterion based on this tissue in the clinical protocol, a treatment plan in the Pareto front can have a large amount of radiation to this tissue without deteriorating the LSI, while possibly contributing to the LCI, especially if dwell positions are inside, or close to, this tissue. This

can occur especially for small numbers of catheters. With an effective optimization algorithm such as MO-RV-GOMEA, this gets fully exploited in finding the best possible plans. Even though the clinical protocol allows this, the physician indicated that such plans are not desirable. An important next step would therefore be to determine a constraint on the amount of radiation to normal tissue in between and around the delineated organs, and include it in the optimization, ultimately changing the clinical protocol.

There are multiple possible reasons why this has not played a (major) role in optimizing dwell times of catheters already in place. The first is that in clinical practice, the number of catheters is sufficiently large for each catheter to have low dwell times. Secondly, when optimizing catheter positions, catheters can be placed at the edge of the prostate, thereby covering the prostate but also having part of their high-dose region outside of the prostate. This way, the lack of criterion on normal tissue can be exploited more than in dwell time optimization. Thirdly, the clinical optimization software contains the possibility to put a criterion on the amount of radiation to normal tissue, even though the clinical protocol does not contain such a criterion. Finally, manual optimization could have improved on the dose to normal tissue, to satisfy clinical intuition and experience rather than the clinical protocol.

A second factor which was not taken into account in the optimization is the robustness of a treatment plan. Especially for catheter placement, robustness is of high importance. After the selection of a treatment plan, catheters still have to be placed inside the patient, and could end up in a (slightly) different position from what was optimized. Especially catheter implants where catheters are close to OARs are less desirable, a fact that is not part of the clinical protocol and thus not of our current optimization model. Moreover, if the plan consists of only 4 catheters, then any mis-positioning compared to the planned positions can have a large influence on the LCI and LSI. In contrast, if more catheters are used, then the mis-positioning of one catheter could be compensated more by adjusting the dwell times of the other catheters. To account for this, the robustness of a treatment plan to small catheter mis-positionings could be included in the optimization objectives.

8 CONCLUSIONS

This paper considers the optimization of catheter positions for HDR prostate brachytherapy. A multi-objective problem formulation was chosen, optimizing directly on the criteria in the clinical protocol and considering the trade-offs between three key objectives: covering the targets, sparing the healthy organs close to the targets, and using as few catheters as possible. The non-linearity of objectives and the multi-objective formulation make the problem highly real-world relevant, but non-trivial to solve, for which we turn to state-of-the-art EAs. By applying MO-RV-GOMEA, an EA which has already shown excellent performance on benchmark problems and on the optimization of dwell times for HDR prostate brachytherapy, problem-specific information on the dependencies between variables and the evaluation of a treatment plan could be exploited effectively. Our results indicate that with this new approach, better treatment plans can be found in much shorter running time, compared to a recent state-of-the-art mixed-integer EA that was even aimed at a restricted version of the problem. This allows for new

insights into catheter placement for brachytherapy to be obtained, as well as on objectives for automated treatment planning, where our results have indicated that richer problem models are needed to better match real-world clinical preferences.

ACKNOWLEDGMENTS

This work is part of the research programme IPPSI-TA (project number 628.006.003), which is financed by the Netherlands Organisation for Scientific Research (NWO) and Elekta. The authors gratefully acknowledge the Nijbakker-Morra Stichting for financing a high-performance computing system.

REFERENCES

- [1] A. Bouter, N.H. Luong, C. Witteveen, T. Alderliesten, and P.A.N. Bosman. 2017. The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 537–544.
- [2] T.M. Deist and B.L. Gorissen. 2016. High-dose-rate prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing. *Physics in Medicine & Biology* 61, 3 (2016), 1155.
- [3] A.M. Dinkla, R. van der Laarse, E. Kaljouw, B.R. Pieters, K. Koedooder, N. van Wieringen, and A. Bel. 2015. A comparison of inverse optimization algorithms for HDR/PDR prostate brachytherapy treatment planning. *Brachytherapy* 14, 2 (2015), 279–288.
- [4] L. Eapen, C. Kayser, Y. Deshaies, G. Perry, E. Choan, C. Morash, J.E. Cygler, D. Wilkins, and S. Dahrouge. 2004. Correlating the degree of needle trauma during prostate brachytherapy and the development of acute urinary toxicity. *International Journal of Radiation Oncology • Biology • Physics* 59, 5 (2004), 1392–1394.
- [5] S.K. Kang, R.H. Chou, R.K. Dodge, R.W. Clough, H.L. Kang, M.G. Bowen, B.A. Steffey, S.K. Das, S. Zhou, A.W. Whitehurst, et al. 2001. Acute urinary toxicity following transperineal prostate brachytherapy using a modified Quimby loading method. *International Journal of Radiation Oncology • Biology • Physics* 50, 4 (2001), 937–945.
- [6] A. Karabis, P. Belotti, and D. Baltas. 2009. Optimization of catheter position and dwell time in prostate HDR brachytherapy using HIPO and linear programming. In *World Congress on Medical Physics and Biomedical Engineering, September 7-12, 2009, Munich, Germany*. Springer, 612–615.
- [7] A. Karabis, S. Giannouli, and D. Baltas. 2005. 40 HIPO: A hybrid inverse treatment planning optimization algorithm in HDR brachytherapy. *Radiotherapy and Oncology* 76 (2005), S29.
- [8] E. Lessard and J. Pouliot. 2001. Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Medical Physics* 28, 5 (2001), 773–779.
- [9] N.H. Luong, T. Alderliesten, A. Bel, Y. Niatsetski, and P.A.N. Bosman. 2017. Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment. *Swarm and Evolutionary Computation* (2017). <https://doi.org/10.1016/j.swevo.2017.12.003>
- [10] B. Pierquin, A. Dutreix, C.H. Paine, D. Chassagne, G. Marinello, and D. Ash. 1978. The Paris system in interstitial radiation therapy. *Acta Radiologica: Oncology, Radiation, Physics, Biology* 17, 1 (1978), 33–48.
- [11] K.L. Sadowski, M.C. van der Meer, N.H. Luong, T. Alderliesten, D. Thierens, R. van der Laarse, Y. Niatsetski, A. Bel, and P.A.N. Bosman. 2017. Exploring trade-offs between target coverage, healthy tissue sparing, and the placement of catheters in HDR brachytherapy for prostate cancer using a novel multi-objective model-based mixed-integer evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 1224–1231.
- [12] D. Thierens and P.A.N. Bosman. 2011. Optimal mixing evolutionary algorithms. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 617–624.
- [13] C. Vargas, M. Ghilezan, M. Hollander, G. Gustafson, H. Korman, J. Gonzalez, and A. Martinez. 2005. A new model using number of needles and androgen deprivation to predict chronic urinary toxicity for high or low dose rate prostate brachytherapy. *The Journal of Urology* 174, 3 (2005), 882–887.
- [14] World Cancer Research Fund International. February 2017. *Prostate cancer statistics*. <http://www.wcrf.org/int/cancer-facts-figures/worldwide-data>.

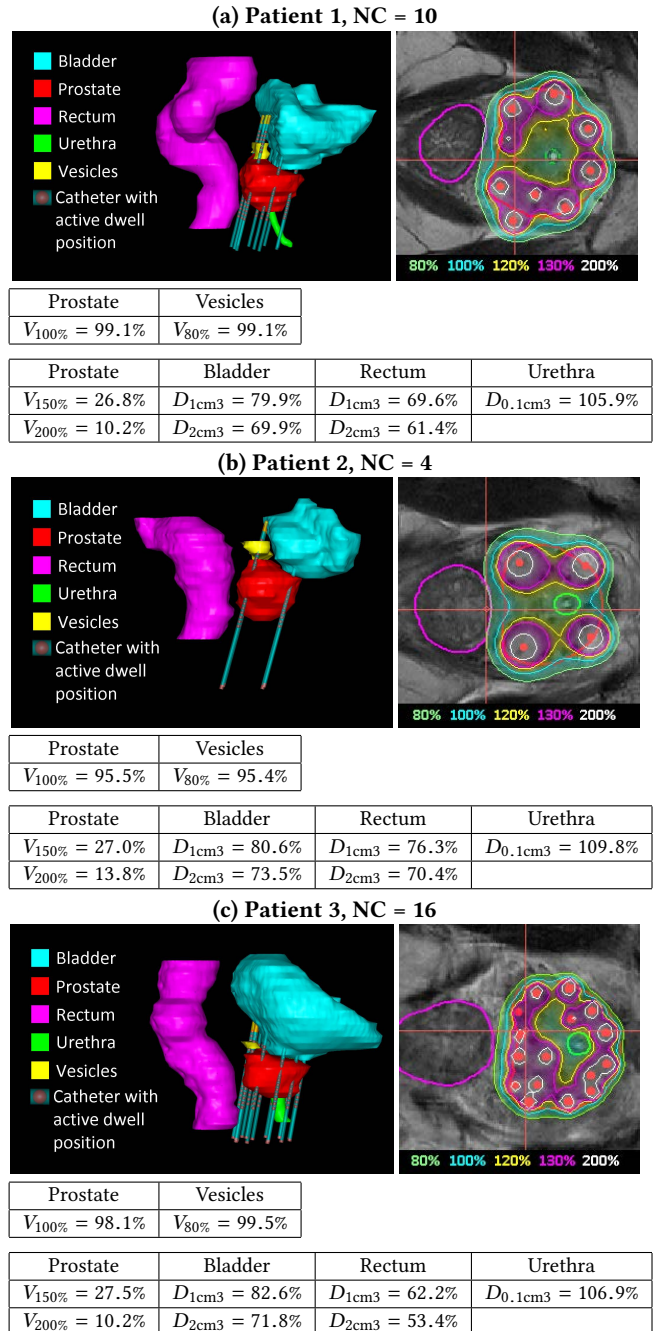


Figure 6: Selected solution for patient 1, 2, and 3 using 10, 4, and 16 catheters respectively, satisfying all clinical criteria. The DVIs were re-evaluated in the clinical software.