



Contents lists available at ScienceDirect

Swarm and Evolutionary Computation

journal homepage: www.elsevier.com/locate/swevo

Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment

Ngoc Hoang Luong^{a,*}, Tanja Alderliesten^b, Arjan Bel^b, Yury Niatsetski^c, Peter A.N. Bosman^a

^a Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands

^b Academic Medical Center (AMC), Amsterdam, The Netherlands

^c Elekta, Veenendaal, The Netherlands

ARTICLE INFO

Keywords:

Linkage learning
Multi-objective optimization
Partial evaluation
Brachytherapy
Radiotherapy
Cancer treatment planning

ABSTRACT

High-Dose-Rate (HDR) brachytherapy (BT) treatment planning involves determining an appropriate schedule of a radiation source moving through a patient's body such that target volumes are irradiated with the planning-aim dose as much as possible while healthy tissues (i.e., organs at risk) should not be irradiated more than certain thresholds. Such movement of a radiation source can be defined by so-called *dwell times* at hundreds of potential *dwell positions*, which must be configured to satisfy a clinical protocol of multiple different treatment criteria within a strictly-limited time frame of not more than one hour. In this article, we propose a bi-objective optimization model that intuitively encapsulates in two objectives the complicated high-dimensional multi-criteria nature of the BT treatment planning problem. The resulting Pareto-optimal fronts exhibit possible trade-offs between the coverage of target volumes and the sparing of organs at risk, thereby intuitively facilitating the decision-making process of treatment planners when creating a clinically-acceptable plan. We employ real medical data for conducting experiments and benchmark four different Multi-Objective Evolutionary Algorithms (MOEAs) on solving our problem: the Non-dominated Sorting Genetic Algorithm II (NSGA-II), the Multi-Objective Evolutionary Algorithm based on Decomposition (MOEA/D), the Multi-objective Adapted Maximum-Likelihood Gaussian Model Iterated Density-Estimation Evolutionary Algorithm (MAMaLGaM), and the recently-introduced Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA). The variation operator that is specific to MO-RV-GOMEA enables performing partial evaluations to efficiently calculate objective values of offspring solutions without incurring the cost of fully recomputing the radiation dose distributions for new treatment plans. Experimental results show that MO-RV-GOMEA is the best performing MOEA that effectively exploits dependencies between decision variables to efficiently solve the multi-objective BT treatment planning problem.

1. Introduction

1.1. Clinical practice

Brachytherapy (BT) [1,2] involves the use of radiation sources to treat cancer by irradiating cancerous tissues from inside the patient's body, whereas External Beam Radiation Therapy (EBRT) delivers radiation by use of beams passing through the patient's body. Because irradiation occurs at the vicinity of target volumes in BT, radiation dose distributions can be fine-tuned better to conform to target volumes while sparing surrounding healthy organs from radiation risks. Further-

more, high-strength radiation sources, so-called High-Dose-Rate (HDR) sources, can be employed, resulting in fewer treatment sessions compared to EBRT. In this article, we address HDR-BT for prostate cancer but the methodology can be easily adapted to other types of BT, such as Low-Dose-Rate (LDR) or Pulse-Dose-Rate (PDR), where the employed radiation sources have lower strengths.

The target volumes in prostate cancer treatment are the prostate and (part of) the seminal vesicles while the Organs At Risk (OARs) are the bladder, the rectum, and the urethra. Depending on the sizes and specific locations of the target volumes, about 14–20 *catheters* are inserted into the patient's body through the transperineal skin, reaching (and

* Corresponding author.

E-mail address: hoang.luong@cwi.nl (N.H. Luong).

<https://doi.org/10.1016/j.swevo.2017.12.003>

Received 7 August 2017; Received in revised form 22 November 2017; Accepted 5 December 2017

Available online XXX

2210-6502/© 2017 Elsevier B.V. All rights reserved.




Fig. 1. An example implant of catheters in HDR-BT treatment for prostate cancer.

going through) the target volumes. An example implant of catheters in HDR-BT treatment for prostate cancer is given in Fig. 1. Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) scans (i.e., medical images) of the patient's pelvic area are then acquired and loaded into BT treatment planning software to be used in the treatment planning session. BT treatment planners (i.e., radiation oncologists, radiation therapy technologists, and clinical physicists) then delineate (i.e., contour) the inserted catheters, the target volumes, and the OARs on the obtained medical images. Each catheter contains a number of so-called *dwell positions*, typically with a step size of 2.5 mm. Only dwell positions within the margins of 5.0 mm extended from surfaces of target volumes are activated while the ones located further away from the surface of the target volumes are kept inactive. When passing through the catheters, a radiation source can reside at each activated dwell position for a certain time, termed *dwell time*, before being moved to the next activated dwell position. The longer the source dwells at a position, the more radiation is released toward the surrounding tissues. After the delineation of target volumes and OARs, BT treatment planners proceed to construct a proper treatment plan. Radiation oncologists plan a certain radiation dose level, termed *planning-aim dose*, that is sufficient to sterilize tumor cells. An HDR-BT treatment plan consists of the dwell times at all activated dwell positions. Dwell times should be long enough to ensure that the amount of dose delivered to target volumes at least equals the planning-aim dose while not resulting in excessive dose being delivered to healthy cells. A clinically-acceptable treatment plan should satisfy the clinical protocol which consists of treatment criteria (or requirements) that indicates the effective thresholds for dose to target volumes and the safety thresholds for dose to OARs. After a clinically-acceptable treatment plan is constructed and approved, the inserted catheters are connected to the *afterloader*, which controls the movements of the radiation source. The radiation source is passed through the catheters such that the source dwells at each dwell position for the amount of time as indicated in the approved plan. After the treatment plan is carried out, the source is retrieved back into the afterloader.

1.2. Challenges in solving the BT treatment planning problem

Stated in its most basic form, the BT treatment planning problem involves determining how long each dwell time should be such that the target volumes are properly covered by a certain planning-aim dose while not exposing OARs to radiation levels higher than clinically-acceptable thresholds. In practice, making such BT plans is non-trivial.

The time budget available for planning is highly constrained, which should not be more than one hour, while HDR-BT plans typically consist of hundreds of dwell times that need to be configured to satisfy multiple clinical criteria. Further, the clinical protocol consists of so-called Dose-Volume (DV) criteria, which involves discrete terms (see Section 2), making direct formulations of the optimization problem associated with BT treatment planning difficult to be efficiently solved due to the absence of fine-grained gradient information. Therefore, the problem is often reformulated using relaxed models that can be quickly solved by, e.g., simulated annealing [3], BFGS [4], or linear programming [5]. The optimal solutions of these simplified models are not necessarily (near) the optimal solutions of the original problem and are therefore not guaranteed to satisfy all criteria in the clinical protocol. Several studies have been done into directly handling DV criteria [6,7]; however, the time required to obtain optimal solutions makes the use of exact solvers clinically prohibitive. Detailed surveys on problem formulations and BT treatment planning methods can be found in Refs. [8,9].

Another challenge associated with the BT treatment planning problem lies in its multi-objective nature. On the one hand, target volumes should be covered by the planning-aim dose as much as possible; otherwise, the treatment is not effective because the tumor is under-irradiated. On the other hand, OARs should be exposed to radiation as little as possible; otherwise, over-irradiation may lead to adverse effects. Maximizing target coverage and maximizing OARs sparing are therefore two conflicting objectives, where the *utopian* treatment plan, that covers entire target volumes with sufficient dose while delivering no radiation to OARs, does not exist. Instead, there exists a *Pareto-optimal set of non-dominated solutions* (i.e., treatment plans). The objective value vectors of these non-dominated solutions form the so-called *Pareto-optimal front* that exhibits the possible optimal trade-offs between the objectives such that improving the coverage of targets leads to deterioration in the sparing of OARs, and vice versa. Particulars of specific patient cases that are not in the optimization models (e.g., physical condition or age of the patient) and the preference of the treating physician cause different treatment plans to be deemed the most preferable trade-off to be carried out. Existing optimization tools in clinical BT treatment planning software, however, do not handle the problem in the *true* multi-objective manner [10]. Instead, they employ the weighted-sum method to aggregate all the (reformulated) clinical criteria into a single-objective optimization function. Solving one single-objective formulation for a single combination of weights results in a single solution. Moreover, each such formulation is equally difficult, and therefore, is either time-consuming to solve or requires some form of reformulation as mentioned earlier. The latter is the case in available clinical software. It is moreover difficult to determine *a priori* the desired setting of coefficients because this depends on the geometry of the implant (i.e., the locations of organs and inserted catheters), and again, the specific details of the patient being treated as well as the treating physician's preferences. Hence, for different patients, the same setting of coefficients leads to different trade-offs (i.e., objective value vectors). As a consequence, treatment planners typically need to run clinically available optimization tools several times with different coefficient settings (equivalent to different formulations) to obtain different suggestions and must often further manually adjust the suggested solutions until reaching a desired treatment plan.

1.3. Related works

In this article, we tackle the BT treatment planning problem in the *true* multi-objective optimization manner, i.e., directly approaching the Pareto-optimal front. Moreover, we employ objective formulations that consider directly the DV indices that are clinically used to evaluate the quality of treatment plans. Because the number of possible trade-offs on the Pareto-optimal front can be numerous, or even infinite, and because achieving mathematical optimality in practice is often not the most important aspect, it often suffices to obtain an approxima-

tion set of non-dominated solutions that is *as close as possible* to the Pareto-optimal front and *as diverse as possible* (measured in the objective space). Such an approximation set exhibits and visualizes the possible trade-offs between the coverage of target volumes and OARs sparing, which has high potential to serve as a decision support tool for treatment planners because the treatment plan associated with the desired trade-off can be efficiently determined *a posteriori*.

Evolutionary Algorithms (EAs) have been shown to be among the state of the art to tackle difficult multi-objective optimization problems [11]. EAs typically do not require (much) problem-specific knowledge, such as gradient information, in their operations, making EAs highly suitable for black-box optimization where domain knowledge is unavailable or too complicated to be effectively exploited in a fully problem-specific tailor-made white-box optimization algorithm. Such information, if available, can typically still be used to customize an EA's operations to further enhance their performance without extreme effort. Furthermore, EAs are population-based algorithms, which makes them inherently well-suited for multi-objective optimization in the sense that multiple solutions, that approximate the Pareto-optimal front, can be obtained at the same time in one single run instead of running a single point-based algorithm (e.g., hill-climbing) multiple times with different settings.

Various Multi-Objective EAs (MOEAs) have been used to solve difficult multi-objective formulations of the BT treatment planning problem [8,12,13]: the Niched Pareto Genetic Algorithm (NPGA), the Non-dominated Ranking Genetic Algorithm (NRGA), the Strength Pareto Evolutionary Algorithm (SPEA), and the Non-dominated Sorting Genetic Algorithm (NSGA). However, the problem objectives considered in these works were not formulated directly from the criteria in common clinical protocols. Therefore, it was inconvenient for treatment planners to interpret their resulting approximation sets. A more recent problem formulation based on the clinical criteria was proposed in Ref. [10], where the well-known MOEA Non-dominated Sorting Genetic Algorithm II (NSGA-II) [14] was employed as the optimization algorithm. However, the problem formulation in this study only contained a small subset of commonly-used clinical criteria. The fact that not all clinical criteria were used whereas at the same time the dimensionality of the optimization problem was non-trivial (i.e., 4 objectives), makes it not straightforward for planners to intuitively traverse the obtained results in terms of identifying treatment plans that satisfy all the involved clinical criteria. Also, the experiment in Ref. [10] was demonstrated for only one prostate case, so the obtained results are difficult to generalize. Interested readers are referred to [8,9] for more detailed surveys on BT treatment planning methods.

The recently-introduced Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA) [15] has been shown to be a promising optimizer especially in the case of gray-box optimization where certain assumptions are made about types of problem-specific knowledge that can be exploited, which may well be the case in many real-world problems. MO-RV-GOMEA is considered a member of the GOMEA family of EAs [16]. In the black-box context, GOMEAs employ linkage learning to recognize potential *linkages* between problem variables (i.e., groups of problem variables that exhibit a dependency relation to some degree), which are then effectively exploited for solution variation (i.e., generating offspring solutions) by the Gene-pool Optimal Mixing (GOM) operator [15]. Such capabilities are highly important to the effectiveness of the variation operators and the overall scalability of EAs in general [16]. In the gray-box context, such linkage information may be available directly from the problem formulation, allowing a linkage model to be defined offline (i.e., *a priori*), increasing the efficiency of GOMEAs. For the BT treatment planning problem, the geometry information of the dwell positions in the inserted catheters can be used to build customized linkage models because it is likely that stronger dependencies exist between dwell positions that are close to each other. The performance of GOMEAs can be substantially enhanced when *partial evaluations* are

possible [15,17]. Partial evaluations relate to the idea that the objective values of an offspring solution can be efficiently calculated if the offspring solution only differs from a parent solution in a few problem variable values. Different from existing MOEAs, the MO-RV-GOMEA, like any GOMEA, does not create an entire offspring solution each time, but iteratively improves each existing solution in a step-wise manner. That is, local changes to the existing solution are made and evaluated for improvement checks. The evaluations pertaining to BT treatment planning are computationally expensive due to many required calculations of the radiation dose distribution at target volumes and OARs. Therefore, it is highly beneficial if the evaluation time can be reduced by employing partial evaluations to compute the impact of local changes on existing treatment plans when MO-RV-GOMEA performs solution variation. Indeed, first results of applying MO-RV-GOMEA to the BT treatment planning problem showed much promise on 3 patient cases [18].

1.4. Our contributions

While the problem formulation employed in Ref. [18] took into account commonly-used clinical criteria based on DV indices, many of them were altered to a mathematically equivalent constraint formulation for the sake of computational efficiency and uniformity in formulation. Treatment plans satisfying all clinical criteria could thereby be identified but the objective values were not easy to be interpreted by BT planners (who are medical experts). In this article, we therefore propose a different formulation for the BT treatment planning problem that encapsulates critical treatment requirements in their original form, thus not compromising on the validity nor the interpretability of the obtained optimization results. Similar to [18], we benchmark the performance of multiple MOEAs in solving the BT treatment planning problem: the well-known MOEA NSGA-II [14], the Multi-objective Adapted Maximum-Likelihood Gaussian Model Iterated Density-Estimation Evolutionary Algorithm MAMaLGaM [19], and the recently-introduced MO-RV-GOMEA [15]. This article furthermore extends our pilot publication by additionally considering another well-known MOEA: the Multi-Objective Evolutionary Algorithm based on Decomposition MOEA/D [20]. We further substantially extend the work in Ref. [18] by conducting experiments on a much larger set of patient cases (15 cases vs. 3 as in Ref. [18]), consolidating the generality of the results. As a final addition, we present the results of multiple runs in a probabilistic way using heatmaps, giving a fair indication of expected results when an algorithm is run just once in practice.

The remainder of the article is organized as follows. Section 2 presents the set of clinical criteria that are frequently used for BT treatment planning of prostate cancer. Section 3 proposes our multi-objective formulation of the problem. Section 4 outlines the 4 MOEAs that we consider in this article. Section 5 shows how geometry information of the catheter implant can be exploited for linkage learning and how partial evaluations can be implemented. Section 6 presents and discusses the results of our experiments. Section 7 concludes the article.

2. Dose-Volume criteria for HDR-BT treatment of prostate cancer

Each treatment plan produces a certain radiation dose distribution in the tissues surrounding the catheter implant. At the Academic Medical Center (AMC, Amsterdam, The Netherlands), the hospital involved in this study, a radiation dose of 13 Gy is currently employed for HDR-BT treatment for prostate cancer. A treatment plan is deemed clinically acceptable if it satisfies the protocol that is employed to indicate the effective lower bounds of target volume coverage and the allowable upper bounds of healthy organ radiation exposure. It should however be noted that a plan that does not satisfy this protocol, can still be considered clinically acceptable. The extent to which a deviation is allowed depends on many factors in practice, including, for instance, whether

better plans are not expected to be possible because of, e.g., a strongly sub-optimal placement of catheters. A clinical protocol of BT treatment often consists of the so-called Dose-Volume (DV) indices and their corresponding criteria (or requirements):

- V_d^o criteria specify how large the *cumulative volume* of an organ o receiving at least the radiation dose level d (relative to the planning-aim dose) should be.
- D_v^o criteria specify how high the radiation dose level that covers the most-irradiated *cumulative volume* v of an organ o should be.

For example, the requirement $V_{100}^{prostate} > 95\%$ indicates that the cumulative prostate volume covered by at least 100% of the planning-aim dose should be more than 95% of the prostate [21]. Similarly, the requirement $V_{80}^{vesicles} > 95\%$ indicates that more than 95% of the seminal vesicles volume should be irradiated with at least 80% of the planning-aim dose. To avoid necrosis in the prostate due to over-irradiation, $V_{200}^{prostate} < 20\%$ indicates that the cumulative prostate volume covered by at least 200% of the planning-aim dose should be less than 20% of the prostate. To protect the rectum from excessive radiation exposure, the requirement $D_{1cc}^{rectum} < 78\%$ demands that the most-irradiated 1 cc (cubic centimeter) of the rectum volume receives less than 78% of the planning-aim dose. Similarly, $D_{2cc}^{bladder} < 74\%$ indicates that the most-irradiated 2 cc of the bladder volume should receive less than 74% of the planning-aim dose. Table 1 presents the DV indices and their corresponding requirements currently employed at the AMC.

A key part of evaluating a treatment plan is calculating the DV index values and comparing them with the thresholds as in Table 1. DV index values are often approximated by computing the radiation dose using a large number of randomly distributed *dose calculation points*. The radiation dose at each dose calculation point is the total sum of the dose contributions received from each activated dwell position. Each dose contribution is the product of the dwell time at the corresponding dwell position with the *dose rate*, which is defined as the amount of radiation delivered from the source when it pauses at the dwell position to the dose calculation point per second (i.e., Gy/s). Dose rates are computed following the TG-43 protocol (the American Association of Physicists in Medicine AAPM Task Group No. 43 Report) [22–24], which involves the strength and shape of the radiation source employed for treatment and the distances between dose calculation points and dwell positions.

Let T denote the set of dwell positions, $|T| = n_T$, and D denote the set of dose calculation points, $|D| = n_D$. For a radiation source with a certain strength, an $n_D \times n_T$ dose rate matrix R is computed, where each matrix element R_{ij} indicates the dose rate between dose calculation point i and dwell position j . A treatment plan can be seen as a vector t of dwell times at activated dwell positions, i.e., $t = (t_1, t_2, \dots, t_{n_T})$. The dose distribution d associated with the treatment plan t can be represented as a vector of radiation doses at all dose calculation points, i.e., $d = (d_1, d_2, \dots, d_{n_D})$, which is computed as:

$$d = Rt \quad (1)$$

Let D_o denote the set of dose calculation points inside organ o , $D_o \subset D$, $|D_o| = n_D^o$. The value of a DV index V_d^o is approximated as:

$$V_d^o = \frac{1}{n_D^o} \sum_{i \in D_o} \chi(d_i, d) \quad (2)$$

where d_i is the total radiation at the dose calculation point i , and $\chi(d_i, d)$ is an indicator function:

$$\chi(d_i, d) = \begin{cases} 1 & d_i \geq d \\ 0 & d_i < d \end{cases} \quad (3)$$

Let d^o be the dose distribution in organ o , i.e., the radiation dose levels at all dose calculation points in organ o . Using a descending sorted version of d^o , denoted as $d^{s,o}$, the value of a DV index D_v^o can be defined as:

$$D_v^o = d_{\lfloor \frac{v^o}{n_D^o} \rfloor}^{s,o} \quad (4)$$

where v^o is the volume (in cc) associated with each dose calculation point inside an organ o . Note that, for the convenience of notation, dose values d_i 's (in d , d^o , and $d^{s,o}$), dose levels d 's in V_d^o indices, and the values of D_v^o indices are presented in terms of percentages of the planning-aim dose. Computing the value of a DV index (V_d^o or D_v^o) is based on numbers of dose calculation points satisfying certain criteria. This counting introduces discontinuities in the form of step functions in the landscape of any DV criteria-based optimization problem formulation, making BT treatment planning a difficult problem to be tackled, even when disregarding the additional complexity introduced by its multi-objective nature.

Disregarding uncertainties in organ contouring and catheter delineation, the larger the set of dose calculation points D is, the better the accuracy of a DV index calculation is. In this article, we employ $|D| = 20,000$ (i.e., 4,000 points in each target volume/OAR), giving results of much higher fidelity than those of many existing works tackling DV criteria in the literature [7,25]. Preliminary tests with (far) less points as used in aforementioned literature ($|D| \approx 4,000$) showed that high-quality Pareto fronts with high occupancy especially in regions with large target coverage (e.g., $V_{100}^{prostate}$) cannot be obtained. Note that, however, sets with many more dose calculation points would make treatment plan evaluations during the optimization process prohibitively computationally expensive.

3. Multi-objective optimization problem formulation

BT treatment planners need to consider multiple DV criteria during the planning process and each DV index can be seen as an objective to be improved (i.e., to be optimized). For example, improving target coverage relates to maximizing the DV indices $V_{100}^{prostate}$ and $V_{80}^{vesicles}$. Improving organ sparing relates to minimizing the DV indices associated with OARs, e.g., $D_{1cc}^{bladder}$, D_{2cc}^{rectum} , $D_{0.1cc}^{urethra}$. Optimizing all DV indices in the clinical protocol (see Table 1) at the same time, however, is equivalent to solving a 9-objective optimization problem. Such problem formulation of high dimensionality is difficult to be efficiently solved, and its 9-dimensional Pareto front results would be difficult to visualize, let alone being interpreted and utilized by human experts. Most DV criteria-based approaches in the literature opt to optimize only one objective, often associated with $V_{100}^{prostate}$, while other DV criteria are formulated as hard constraints [6,7,25]. However, such single-objective optimization results in one single plan, and maximizing the prostate coverage is not the sole objective of BT treatment planners. Planners would normally like to consider multiple alternatives and investigate possible trade-offs between DV indices during the planning process. A multi-objective approach with NSGA-II was proposed in Ref. [10], which considered a smaller number of DV criteria and solved 3-4 objective optimization problems following 3-4 DV indices. We argue that the results of such problem formulation are still not guaranteed to satisfy the full clinical protocol of 9 DV indices in Table 1. Also, the visualization of the results in Ref. [10] did not allow the treatment planners to conveniently traverse the Pareto fronts and interpret the non-dominated solutions.

In our recent work [18], we divided DV indices in the clinical protocol into 2 groups:

1. Target Coverage: $V_{100}^{prostate}$ and $V_{80}^{vesicles}$

Table 1

BT treatment planning DV criteria at AMC.

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

2. Organ Sparing: $V_{200}^{prostate}$, $V_{150}^{prostate}$, D_{1cc}^{rectum} , D_{2cc}^{rectum} , $D_{1cc}^{bladder}$, $D_{2cc}^{bladder}$, $D_{0.1cc}^{urethra}$

We employed the *maximin* approach to combine the Target Coverage group into the *Least Coverage Index* (LCI) objective and the Organ Sparing group into the *Least Sparing Index* (LSI) objective. More specifically, for a candidate treatment plan, its LCI objective value corresponds to the worst-scored DV index in the Target Coverage group and its LSI objective value corresponds to the worst-scored DV index in the Organ Sparing group. The Organ Sparing group contains both V_d^o (volume) and D_v^o (dose) indices, which are measured in different units (cc for volume and Gy for dose) and have different ranges of values. Moreover, because computing D_v^o indices involves a sorting operation, which is more computationally expensive, in our previous work we opted to convert all D_v^o criteria into V_d^o criteria so that the whole group can be formulated into one objective (more details in Ref. [18]). While this conversion still maintains the validity of the model in terms of satisfying the clinical protocol, the obtained results were afterward found not to be straightforward for the treatment planners when evaluating against the original clinical protocol. Moreover, converting D_v^o indices into V_d^o indices has the undesirable effect that the size of an organ dictates its importance in the grouping, which is unnatural (to physicians). In this article, we therefore propose a novel problem formulation that directly employs the DV indices in the clinical protocol as follows.

Let $V_d^{o,min}$ be the lower bound threshold for a V_d^o index in the Target Coverage group. The further the value of V_d^o is above this threshold, the better the target o is covered by its required dose level d . The distance δ_v of the V_d^o index value above the threshold $V_d^{o,min}$ can be measured as:

$$\delta_v(V_d^o) = V_d^o - V_d^{o,min} \quad (5)$$

For a treatment plan t , its LCI value is defined as:

$$LCI(t) = \min\{\delta_v(V_{100}^{prostate}), \delta_v(V_{80}^{vesicles})\} \quad (6)$$

Based on the thresholds of DV indices in the Target Coverage group $V_d^{o,min}$ in Table 1, the value of LCI, that we would like to maximize, is in $[-0.95, 0.05]$. Positive values indicate that *both* clinical protocol requirements $V_{100}^{prostate} > 95\%$ and $V_{80}^{vesicles} > 95\%$ are satisfied. The ideal value 0.05 indicates that the whole volume of each target volume is covered by the recommended dose level. Non-positive values indicate that at least one target volume (i.e., prostate or seminal vesicles) is underdosed according to the clinical criteria.

Let $D_v^{o,max}$ be the upper bound threshold for a D_v^o index in the Organ Sparing group. The further the value of D_v^o is below this threshold, the better the organ o is spared from radiation. The distance δ_d of the D_v^o index value below the threshold $D_v^{o,max}$ can be measured as:

$$\delta_d(D_v^o) = D_v^{o,max} - D_v^o \quad (7)$$

The LSI value of a treatment plan t is defined as:

$$LSI(t) = \min\{\delta_d(D_{1cc}^{bladder}), \delta_d(D_{2cc}^{bladder}), \delta_d(D_{1cc}^{rectum}), \delta_d(D_{2cc}^{rectum}), \delta_d(D_{0.1cc}^{urethra})\} \quad (8)$$

Based on the thresholds $D_v^{o,max}$ of D_v^o indices in Table 1, the value of LSI, that we would like to maximize, is in $(-\infty, 0.74]$, which is unbounded on the minimum side because the radiation dose can be made arbitrarily large. Positive values indicate that all D_v^o indices do not exceed their corresponding thresholds, ensuring that all clinical requirements for D_v^o indices in Table 1 are met. A non-positive value indicates that at least one D_v^o index exceeds its threshold $D_v^{o,max}$. The ideal value 0.74 indicates that there is no radiation to OARs since the most-irradiated 2 cc of the bladder volume and the most-irradiated 2 cc of the rectum volume have the lowest bound value of 0.74. In practice, such a situation only happens when all dwell times are 0.0s, which equals no radiation treatment.

The two indices $V_{150}^{prostate}$ and $V_{200}^{prostate}$ are excluded from the Organ Sparing group and are not considered for evaluating the LSI objective, but instead are employed as hard optimization constraints. This

is because these two indices have a different unit, making a combination with the D_v^o indices non-trivial. Moreover, treatment plans must really satisfy these two criteria $V_{150}^{prostate} < 50\%$ and $V_{200}^{prostate} < 20\%$, i.e., the prostate volume should not be over-irradiated. Also, we do not want to consider treatment plans that under-irradiate target volumes too much either. The clinical requirement $V_{100}^{prostate} > 95\%$ and $V_{80}^{vesicles} > 95\%$ means that at most 5% of the prostate volume and at most 5% of the vesicles volume can be irradiated below their recommended dose levels, i.e., 100% and 80% of the planning-aim dose, respectively. How good treatment plans can be depends on the geometry of the organs and the quality of the catheter implant. It might not be possible to satisfy all criteria in Table 1 because, for example, the rectum is too close to the prostate or the number of inserted catheters is too small. We thus introduce some relaxation on the clinical protocol by employing the clinical protocol constraints as optimization objectives. It still makes sense to also exclude treatment plans that result in excessive over-irradiation of OARs or under-irradiation of target volumes. We thus additionally employ the hard optimization constraint $LCI(t) \geq -0.2$ to allow that at most 25% of a target volume can be under-irradiated. Similarly, we employ the optimization constraint $LSI(t) \geq -0.2$ to allow that the most-violated D_v^o can only exceed its corresponding threshold $D_v^{o,max}$ by 20% of the planning-aim dose. To handle constraint violations, when any two solutions need to be compared, we employ the constrained-domination principle [14,26] as follows. A feasible solution is always preferred to an infeasible solution. Between two infeasible solutions, the one with a smaller total constraint violation is preferred. Between two feasible solutions, the one with better objective values is preferred (i.e., the one that Pareto dominates the other in the multi-objective context).

Any treatment plan can be evaluated to these two objectives (LCI, LSI). The *maximin* approach ensures that the values of all DV indices in a group are at least as good as the representative index of that group. A treatment plan with $V_{150}^{prostate} < 50\%$ and $V_{200}^{prostate} < 20\%$ satisfies all other criteria in the clinical protocol in Table 1 if its LCI > 0 and LSI > 0. All the DV indices are considered and treated individually without the need of employing a more cumbersome model of 9 separate objectives. Furthermore, our problem formulation bears resemblance with the planning process in practice because planners normally attempt to adjust the treatment plan in an iterative manner such that the current most-violated DV index is improved first.

4. Multi-Objective Evolutionary Algorithms (MOEAs)

4.1. NSGA-II

NSGA-II [14] is probably the best-known and still the most widely-used MOEA in practical applications. Thus, NSGA-II is often employed for comparison purposes when introducing new multi-objective algorithms. For real-valued optimization, in each generation, NSGA-II creates offspring solutions by applying the simulated binary crossover (SBX) and mutation operators to parent solutions in the current population. The fitness value of a candidate solution in NSGA-II is based on its non-domination rank (regarding the Pareto-dominance relation with other solutions in the population) and its crowding distance value (regarding the density of solutions in the same non-domination rank in the objective space). The population for the next generation is selected from the combined pool of parent solutions and offspring solutions on the basis of their fitness values such that solutions that have a better non-domination rank and reside in less-crowded regions are favored. Details about NSGA-II can be found in Ref. [14].

Note that we additionally implement an adaptive elitist archive [27] for NSGA-II in this work. An elitist archive is an external population that keeps track of the non-dominated solutions obtained so far along the search. The use of an adaptive elitist archive has been found to be important to the performance of MOEAs because the Pareto-optimal fronts can contain numerous (or even an infinite) number of non-dominated solutions while the working populations of MOEAs are

typically limited in size [27]. Moreover, true convergence of an MOEA is not guaranteed otherwise [28].

4.2. MOEA/D

MOEA/D [20] is a more recent MOEA that has obtained many positive results. MOEA/D differs from many available MOEAs in that the multi-objective problem at hand is decomposed into a number of scalar optimization subproblems. Each subproblem is associated with a specific weight vector that is employed to aggregate all the original objective functions into the single-objective optimization function of that subproblem. Each subproblem maintains a candidate solution whose fitness value is defined by the aggregated objective function instead of the original objective functions. Offspring solutions are created from existing (parent) solutions by recombination (using the SBX operator in the real-valued context) and mutation. An offspring solution will replace the candidate solution of a subproblem if its fitness value (evaluated against the aggregated objective function) is better than that of the current solution. To introduce niching, solution variations and replacements are restricted among a neighborhood of a certain size of each parent solution. The working population of MOEA/D, therefore, can be seen as consisting of the best solution obtained so far for each subproblem. The weighted Tchebycheff aggregation approach is often used to ensure that potential non-convex parts of Pareto-optimal fronts can be obtained. The fitness value of each candidate solution is then defined as its weighted Tchebycheff distance toward the *utopian* solution in the objective space. For the bi-objective formulation of the BT treatment planning problem in this work, the utopian value for (LCI, LSI) is (0.05, 0.74). Details about MOEA/D can be found in Ref. [20]. Note that an adaptive elitist archive [27] is also implemented for the MOEA/D in this work.

4.3. MAMaLGaM

Respecting linkage (i.e., dependency) relations among problem variables during solution variations is crucial for the scalability of EAs [29]. Linkage relates to the notion that certain problem variables are dependent on each other to some degree and should thus be jointly treated (as a group) when performing variation. It has been shown that the traditional variation operators typically employed in NSGA-II (e.g., crossover and mutation), which do not take into account problem linkages, are not efficient in solving certain classes of decomposable multi-objective problems that could otherwise be efficiently solved if the linkage information is exploited [30]. Problem-specific knowledge, if available, can be used to specify such linkage information. In the black-box optimization context, linkage information can be inferred from the working population of EAs by linkage learning procedures.

MAMaLGaM [19] is a state-of-the-art MOEA for the real-valued domain. MAMaLGaM belongs to the family of Estimation-of-Distribution Algorithms (EDAs) that differ from traditional EAs in that, for every generation, offspring solutions are generated by sampling the multivariate Gaussian distributions that are estimated over promising candidate solutions selected from the population. An estimated Gaussian distribution in MAMaLGaM consists of the mean vector and the covariance matrix that characterize the dependencies among all problem variables. Note that Pareto-optimal fronts have multiple regions and non-dominated solutions in different regions typically expose different characteristics. For example, in the BT treatment planning context, plans that maximize target coverage typically differ from plans that maximize organ sparing. MAMaLGaM addresses this issue by employing a cluster-based operation. In every generation, the selection set of promising candidate solutions is partitioned into k equal-sized clusters in the objective space, and for each cluster, a Gaussian distribution is estimated over the solutions belonging to that cluster. An equal number of offspring solutions are sampled from each Gaussian distribution of each cluster to create the next population, ensuring that

all regions of the Pareto-optimal front can be evenly approached. Therefore, MAMaLGaM is an MOEDA that employs a mixture model of Gaussian distributions. An elitist archive is also implemented in MAMaLGaM, and elitism is maintained by replacing parts of the population by elitist archive solutions in every generation. Details about MAMaLGaM can be found in Ref. [19].

4.4. MO-RV-GOMEA

While MAMaLGaM effectively exploits the linkage information using multivariate Gaussian distributions that concern all problem variables jointly, it is arguable that the linkage information could be more efficiently exploited in a more fine-grained manner. Using the Family Of Subsets (FOS) concept, the linkage information is described by a linkage model, which can be represented as a set of *linkage sets*. Each linkage set corresponds with a subset of problem variables that depend on each other to some degree. A linkage set can involve only one problem variable, indicating the independent state of that variable, or multiple problem variables, indicating the joint dependency relation of the involved variables, or even all problem variables, indicating the total dependency relation as in MAMaLGaM. Such fine-grained linkage exploitation is the key feature of the recently-introduced family of Gene-pool Optimal Mixing Evolutionary Algorithms (GOMEAs) [16].

MO-RV-GOMEA [15] shares several algorithmic design concepts with MAMaLGaM [19] as well as with the original concept of the GOMEA family [16]. The key features of MO-RV-GOMEA are linkage learning with a linkage model as done in any GOMEA to recognize dependencies between problem variables and exploiting these during solution variation with the Gene-pool Optimal Mixing (GOM) operator to effectively improve existing solutions. The most commonly-adopted linkage model in the GOMEA family is the Linkage Tree (LT) model. Both the LT model and the GOM operator were first employed in the single-objective GOMEA for the discrete domain [16], in MO-GOMEA for multi-objective combinatorial optimization [31], recently in the real-valued domain with RV-GOMEA [17], and, by combining the latter two, in MO-RV-GOMEA [15]. We further give more details on MO-RV-GOMEA to explain its unique capability to exploit problem-specific linkage features as in Section 5.

MO-RV-GOMEA maintains a population P of n solutions. An adaptive elitist archive is also employed to keep track of non-dominated solutions [27]. In every generation, MO-RV-GOMEA performs truncation selection based on non-domination rankings similar to MAMaLGaM to obtain a selection set S of size $\lceil \tau n \rceil$ from P , where $\tau = 0.35$ was found to give good results [15]. The selection set S is clustered into k clusters (with $k > m$, and m is the number of objectives) that have equal sizes $c = \frac{2}{k}|S|$, resulting in overlapping between neighboring clusters, similarly to MO-GOMEA and MAMaLGaM. For each of the m objectives, a dedicated cluster is assigned to perform single-objective optimization by accepting improvements in the GOM phase based solely on that one objective, aiming to approach *extreme* regions on the Pareto-optimal front.

Instead of estimating multivariate Gaussian distributions with full covariance matrices of all variables as in MAMaLGaM, the LT model can be applied in MO-RV-GOMEA to explicitly model the dependencies between variables at different levels. Although different linkage models may be applied, the LT is the most commonly adopted and most generally applicable model. For MO-RV-GOMEA, in the black-box optimization context, the LT can be learned from each cluster separately. The LT is a hierarchical linkage model and can be built in a bottom-up approach. Let L be the set of all problem variable indices, i.e., $L = \{1, 2, \dots, l\}$. An LT is first initialized with leaf nodes of l singleton sets, each contains exactly one variable index $F^i = \{i\}$, $i \in L$. Next, intermediate nodes are created by iteratively merging two existing nodes that are closest/most similar to each other on the basis of a similarity metric. Each time, the new node is added to the LT, and while the two constituent nodes are still kept in the LT, they will

not be considered for further merging. For example, a new node F^i is created by merging two existing nodes F^j and F^k , i.e., $F^i = F^j \cup F^k$ where $F^j \cap F^k = \emptyset$, $|F^j| < |F^i|$, $|F^k| < |F^i|$. This procedure continues until the root node containing all variable indices is formed. Each node of the LT is then a *linkage set*, and problem variables in the same linkage set can be understood as being dependent on each other to some degree, and should thus be jointly considered by variation operators when creating offspring solutions. The LT constructed as above has exactly $2l - 1$ nodes, exhibiting multiple levels of linkages, ranging from the univariate all-independent state in the leaf nodes, to multivariate dependencies in the intermediate nodes, up to the totally-dependent state in the root node. An example LT is shown in Fig. 2.

In the black-box context, we can employ the Mutual Information (MI) of real-valued variables, derived from the sample Pearson correlation coefficient [32], as the similarity metric in constructing the LT. Moreover, only pairwise combinations of variables are considered, in combination with a hierarchical clustering algorithm known as UPGMA, allowing the LT to be learned in $\mathcal{O}(nl^2)$ time (for details, see Refs. [17,33]). If domain knowledge is available and can be exploited, the similarity metric can be defined based on problem-specific information, and the LT can be learned once beforehand rather than be learned anew from the population in every generation. For each linkage set in the LT of each cluster, a multivariate joint Gaussian distribution is estimated based on the solutions in that cluster. The Gaussian distribution associated with a linkage set only concerns the variables that are indicated in the linkage set. Thus, in total, $2l - 1$ Gaussian distributions are maintained for each cluster. This is different from MAMaLgAM, where a single multivariate Gaussian distribution, concerning all variables, is maintained for each cluster. Each Gaussian distribution, like in MAMaLgAM, is then adapted regarding certain search information collected during the optimization process, such as the distance of improvements found from the Gaussian means, in order to counteract the variance-diminishing effect of selection if needed [15,19]. After all linkage sets in the LT of a cluster are considered, all population solutions associated with that cluster are transformed into new offspring solutions using the variation operator GOM that iteratively improves existing solutions in a step-wise manner, where each step corresponds to a local change associated with a linkage set instead of creating an entire offspring solution each time as in other MOEAs. The GOM variation operator samples the learned Gaussian distributions associated with the linkage sets of the LT of each cluster to improve the population solutions associated with that cluster. Per linkage set, new values for variables pertaining to the linkage set are sampled from the associated Gaussian distribution for all solutions associated with the current cluster. These partially-altered solutions are evaluated and only the modifications that result in an improvement are kept. Further details of MO-RV-GOMEA, including pseudo-code, can be found in Ref. [15].




Fig. 2. An example linkage tree with $l = 8$ problem variables.

5. Problem-specific performance enhancement for MO-RV-GOMEA

5.1. Exploiting geometry information for linkage tree building

In the black-box context, the LT can be learned from the working population as described in Section 4.4. Although the LT can be learned efficiently with a time complexity of $\mathcal{O}(nl^2)$ [33], the linkage learning still incurs some computing overhead and certain linkages might not be captured if the population size is not large enough. If domain knowledge is available to be exploited, the LT can be built before running the optimization, and this offline-constructed LT can be fixed during the run. In the context of the BT treatment planning problem, geometry information of the implant is available, and it can be argued that dwell times at neighboring dwell positions have stronger interactions than those at far apart dwell positions. Such neighboring dwell times, therefore, should be jointly considered when creating offspring solutions. The Euclidean distances between pairs of dwell positions can be computed based on their coordinate information, which are acquired from the CT/MRI scans of the patient with the catheter implant. In this work, we will employ such Euclidean distance metric as the basis of the similarity metric used in learning the LT (see Section 4.4).

5.2. Partial evaluations

MO-RV-GOMEA creates an offspring solution by iteratively applying local changes to an existing solution. Each local change alters the current solution in a few variables, and requires a fitness evaluation to compute objective values and constraint violations. For many real-world problems, solution evaluations are typically computationally expensive. If the new solution differs from an existing one in only a (small) number of variable values, as is exactly typically the case in GOMEAs as a consequence of their design, performing a full evaluation is unnecessarily inefficient if the impact of the new variable values can be efficiently computed. For the BT treatment planning problem, a treatment plan evaluation involves the computation of the radiation dose distribution (Equation (1)) and the computation of DV indices (Equations (2) and (4)), in which the former dominates the computing time. Equation (1) involves a matrix-vector multiplication between the dose rate matrix and the vector of dwell times. If only a few elements of the dwell time vector are altered, the impact can be computed by considering only the columns of the dose rate matrix associated with the dwell positions of the altered dwell times. For example, let t denote an existing treatment plan and let d be its dose distribution that has been computed. Let t' denote a new treatment plan that is created by altering a few dwell time values of t , i.e., $t' = t + \Delta t$, in which Δt has many zero elements. The new dose distribution d' can be calculated as:

$$d' = R t' = R(t + \Delta t) = R t + R \Delta t = d + R \Delta t \quad (9)$$

where the i -th component of $R \Delta t$ is:

$$(R \Delta t)_i = \sum_{\substack{j=1 \\ \Delta t_j \neq 0}}^{n_T} R_{ij} \Delta t_j. \quad (10)$$

Thus, $R \Delta t$ involves the multiplication of the columns $R_{:j}$ of R that correspond to non-zero elements Δt_j of Δt . In other words, the entire $R \Delta t$ is not computed, but rather only the compact vector of non-zero elements in Δt associated with the variables of the linkage set at hand.

Partial evaluations can be straightforwardly used for MO-RV-GOMEA thanks to its genetic-local-search-like variation operator GOM. For other MOEAs employed in this work, i.e., NSGA-II, MOEA/D, and MAMaLgAM, a whole new offspring solution is created each time during variation. Such offspring solutions differ completely from the parent solutions (especially for variation in the real-valued domain), requiring a full evaluation for each new offspring solution.

6. Experiments

6.1. Experiment settings

Data of fifteen anonymized HDR-BT cases for prostate cancer from the Academic Medical Center (AMC) were available for conducting our experiments. For each case, the contours of the involved organs (i.e., prostate, seminal vesicles, bladder, rectum, and urethra) and the inserted catheter information were exported by using the BT treatment planning software Oncentra Brachy (research version 4.6, Elekta Brachytherapy). This extracted information is then used as the input data for multi-objective optimization.

We perform experiments with three MO-RV-GOMEA variants employing three linkage models: the Univariate Factorization (UF) model where all dwell times are deemed independent from each other, the Linkage Tree (LT) model which is learned from the population in each generation, and the fixed LT which is constructed *a priori* based on the geometry information of active dwell positions. For each MO-RV-GOMEA variant, we run two settings: 1) black-box full treatment plan evaluations are always carried out to assess candidate treatment plans, and 2) gray-box partial evaluations are enabled to assess partially-altered treatment plans when performing solution variation with GOM. We firstly benchmark different MO-RV-GOMEA variants with different linkage models and the option of performing partial evaluations. We employ the first 3 patient cases to select which MO-RV-GOMEA variant is the most favorable variant to be later used for conducting experiments with the full data set of 15 cases and to show the impact of the ability of exploiting partial evaluations. For the purpose of performance comparison between MO-RV-GOMEA and state-of-the-art MOEAs, we consider NSGA-II [14], MOEA/D [20], and MAMaLGaM [19]. For every patient case, we run each algorithm 30 times independently. Each optimization run is allowed to operate one hour (1 h) to obtain an approximation set of non-dominated plans. The approximation sets from 30 runs of each algorithm will be used to analyze its consistency in terms of variance in the obtained results.

All MOEAs concerned in this work are population-based optimization algorithms, for which the population size parameter setting is crucial to their performance. However, it is impossible to determine *a priori* the optimal population size for each MOEA solving each problem instance in real-world applications because the optimal solutions are unknown. In order to get rid of the need of setting the population size parameter while still putting all MOEAs on an equal footing for a fair performance comparison, we employ an Interleaved Multi-start Scheme

(IMS). This IMS was firstly introduced for a single-objective genetic algorithm [34], and was later adapted for the multi-objective domain in Ref. [15]. In essence, IMS operates multiple populations of different sizes in an interleaved fashion. The first population is started with some (arbitrarily) small size. Populations of larger sizes (often double the size of the preceding one) start later. Each population is run for 1 generation after the preceding population has been run for b generations. Here, we suggest $b = 8$, which was previously found to give good results [15]. The elitist archive is shared between all populations. A population is deemed to be inefficient and should be terminated if it contributes less than 10% of the non-dominated front combined from all populations while a population of larger size manages to contribute more than 10% of the front. More details about IMS can be found in Refs. [15,34]. The elitist archive is maintained to contain about 1,000 non-dominated solutions, which is deemed sufficient for the BT treatment planning problem. When this capacity is exceeded, the archive will be adapted as in Ref. [27].

We use the well-known hypervolume metric [35] to compare the convergence performance of MOEAs. The hypervolume can be intuitively defined as the volume (or area in the case of bi-objective optimization) in the objective space that is covered by a set of non-dominated solutions and a *reference* point, which is a point that can be selected such that it will be dominated by any possible solutions. Here, due to the fact that $LCI \geq -0.2$ and $LSI \geq -0.2$ (see Section 3), we choose $(-0.3, -0.3)$ as the reference point. We employ the Mann-Whitney-Wilcoxon statistical hypothesis test for equality of medians with $p < \alpha = 0.05$ to see whether the final result obtained by one MOEA is statistically significantly different from that of another MOEA.

In this work, to carry out a large number of experiments within a reasonable time budget, we use a server computer with 128 somewhat dated AMD cores (AMD Opteron Processor 6386 SE, 2.8 GHz) to run multiple experiments at the same time. We note that the optimization performance of any individual run in practice could be made at least twice faster by using computers with more recent, powerful CPUs (e.g., recent Intel cores).

6.2. Results

6.2.1. Hypervolume results of MO-RV-GOMEA variants

Fig. 3 shows the graphs of hypervolume development over time (in seconds) averaged over 30 runs of each MO-RV-GOMEA variant for the first 3 patient cases in our data set. It can be seen clearly that par-




Fig. 3. The average hypervolume values of the Pareto fronts of different MO-RV-GOMEA variants over time. The error bars represent the standard deviation.

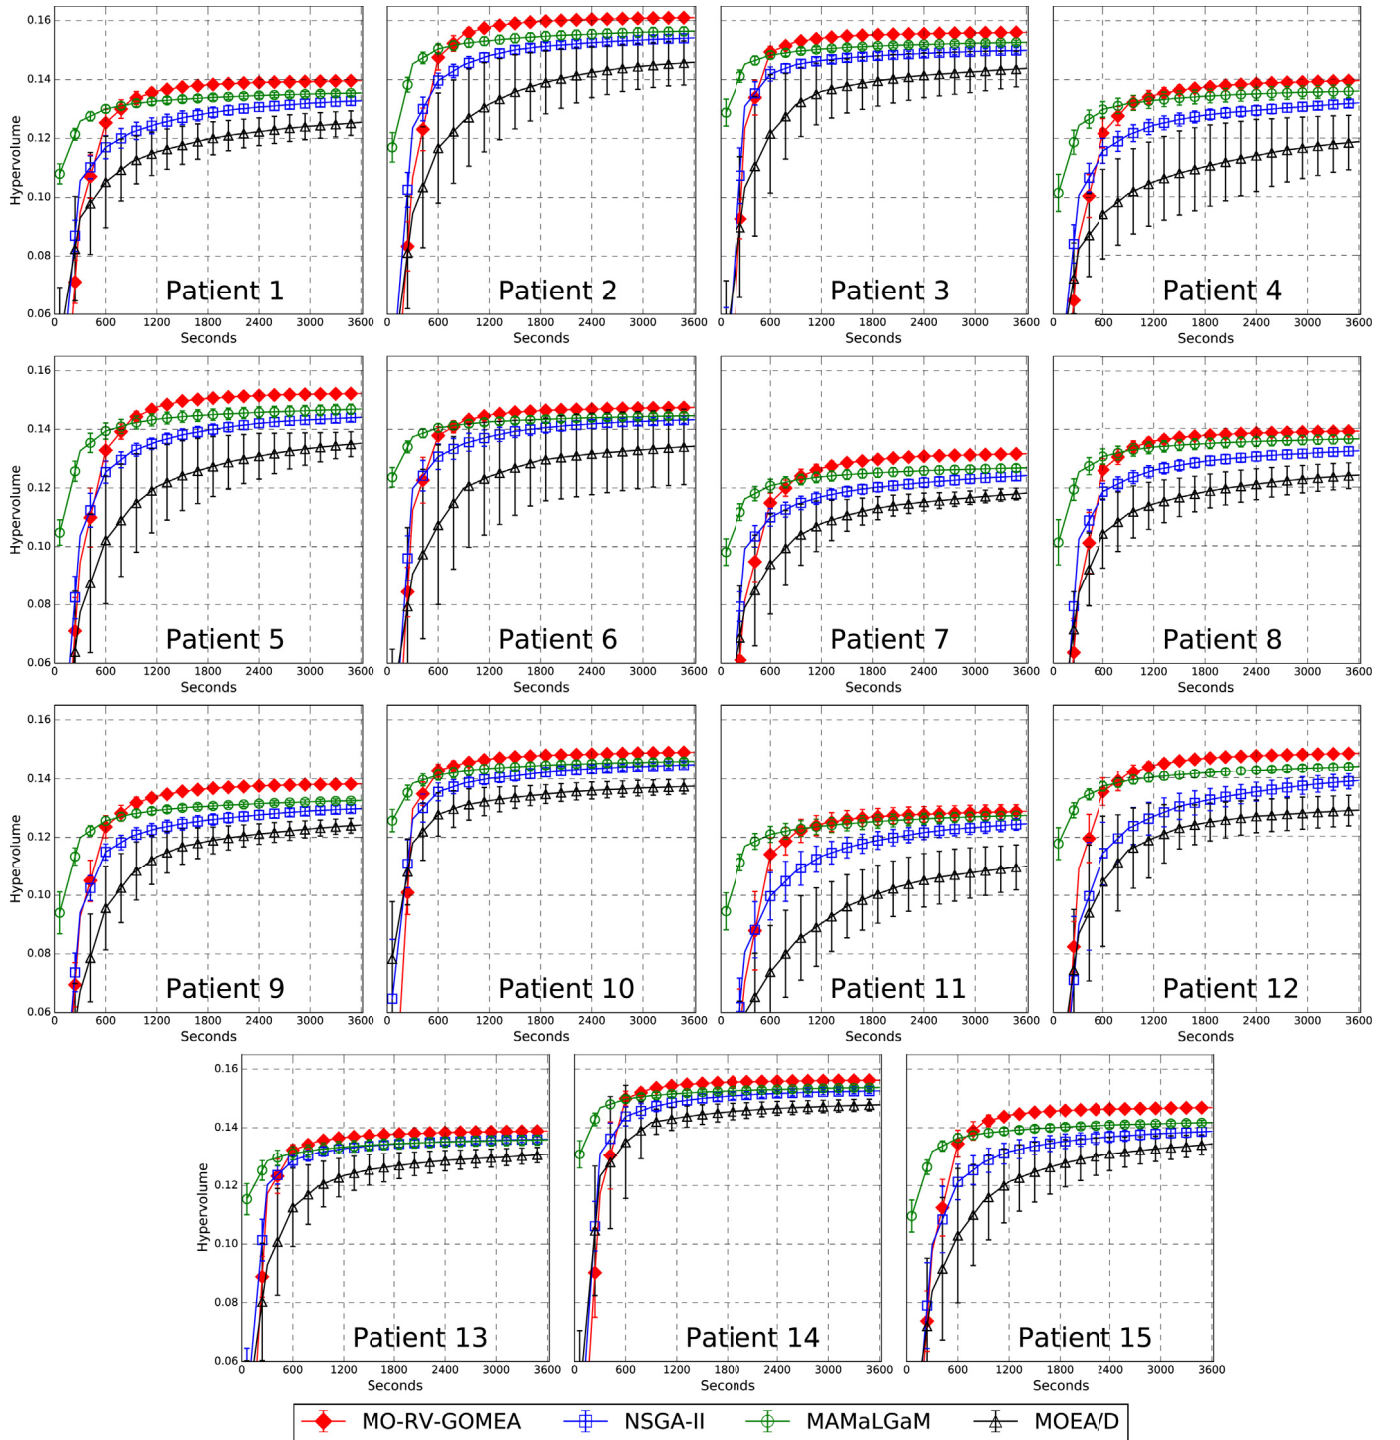


Fig. 4. The average hypervolume values of the Pareto fronts of optimization algorithms over time. Error bars represent the standard deviation.

tial evaluations indeed substantially accelerate MO-RV-GOMEA. While the hypervolume development graphs of MO-RV-GOMEA variants with partial evaluations become relatively flat after about 20 min, indicating convergence, the graphs of the variants without partial evaluations are nowhere near as flat after 1 h, suggesting that the optimizers are still in the middle of the search. Note that besides this extra required computing time, there is no difference between the black-box and gray-box versions of MO-RV-GOMEA. In the black-box case, the additional evaluations, especially for small linkage sets, are actually harmful in terms of computational efficiency. It would be more efficient in that case to use only large linkage sets or even a single linkage set with all

variables. In that case, MO-RV-GOMEA is quite close to MAMaLGaM in its behavior [15]. Regardless of the employed linkage model tested here, MO-RV-GOMEA with partial evaluations always outperforms the corresponding variant without partial evaluations, and the differences between the final hypervolume values obtained by each variant after 1 h are found to be statistically significant ($p < 0.001$). As the problem structure of the BT treatment planning problem allows partial evaluations to be carried out, for all experiments of MO-RV-GOMEA in the remainder of this article, only the variants with partial evaluations are considered.

Fig. 3 also shows that MO-RV-GOMEA with the LT performs slightly

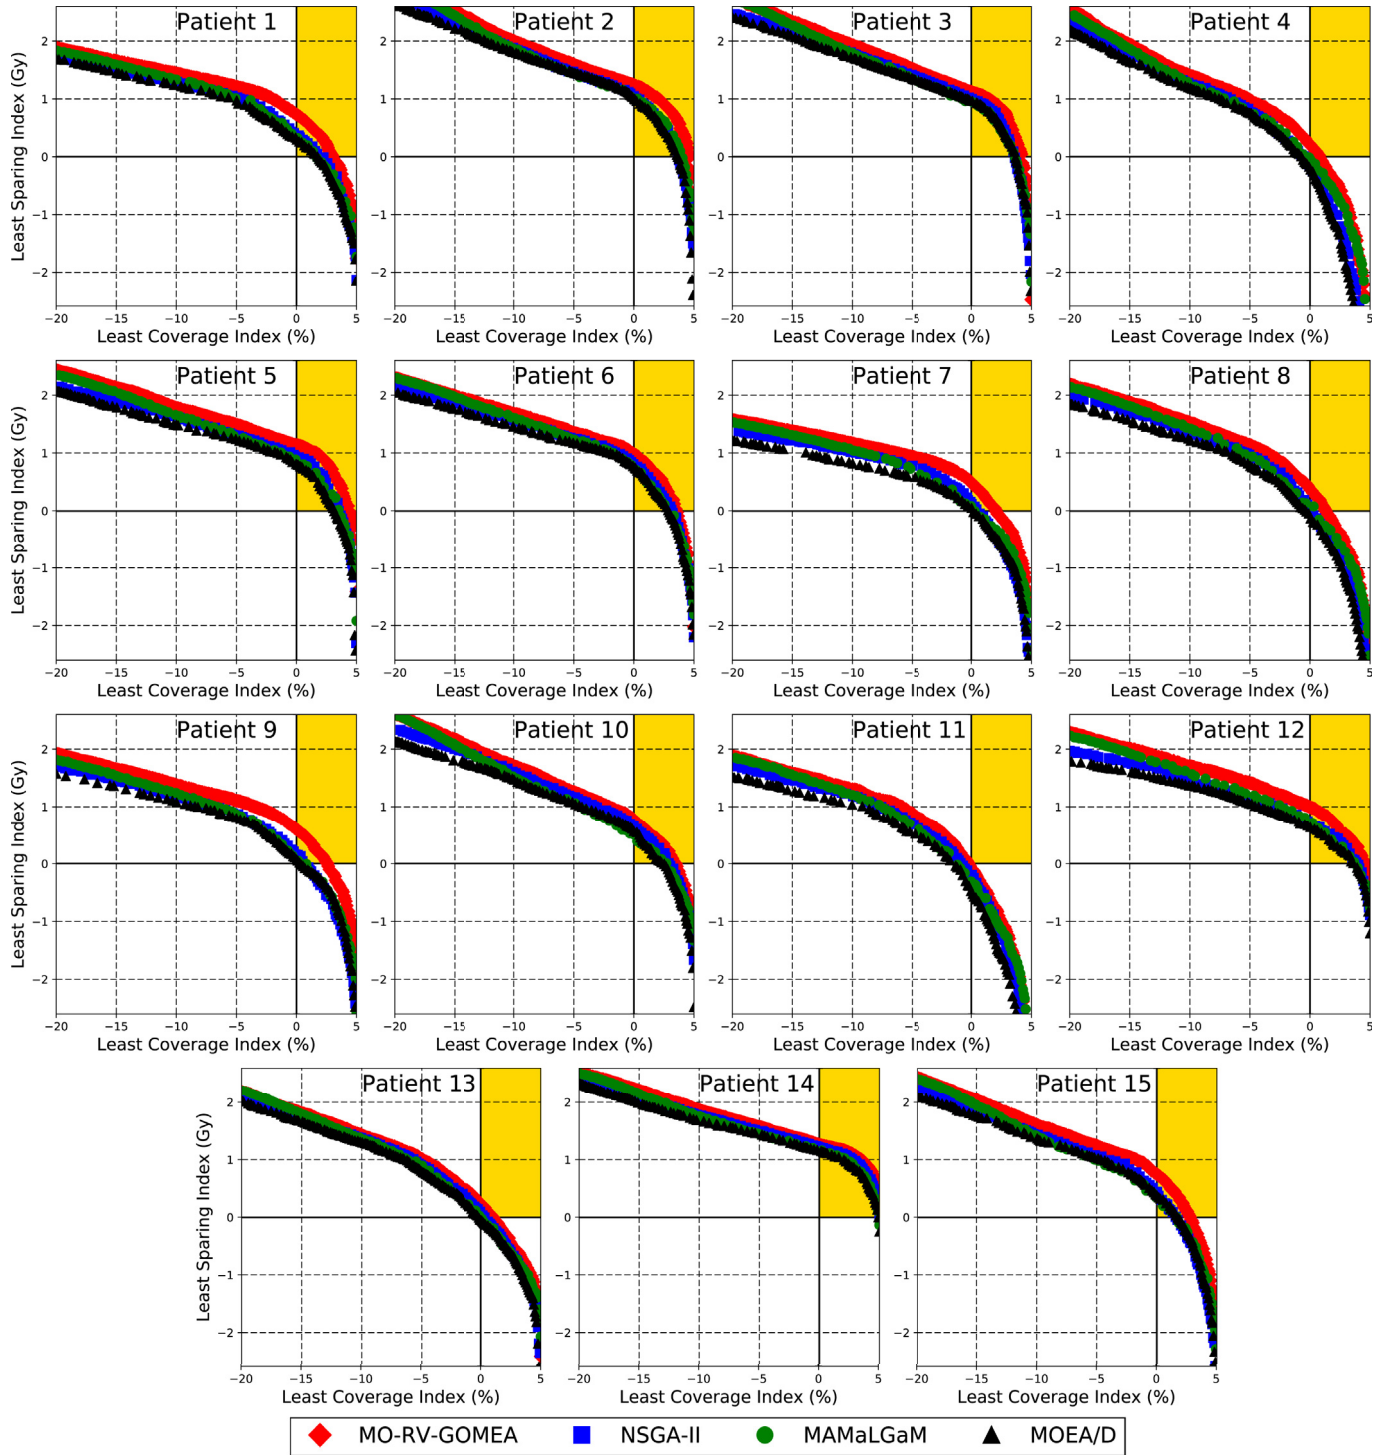


Fig. 5. Pareto fronts combined from 30 independent optimization runs of each algorithm after running for 1 h.

better than the variant with the UF model. These differences are found to be statistically significant ($p = 0.009$ for patient 2, and $p < 0.001$ for the other two patients). This confirms the suggestion that there exist certain useful dependencies between dwell times that can be exploited by the variation operator. Furthermore, MO-RV-GOMEA with the fixed LT is found to perform slightly better than the variant with the LT learned in each generation, where the differences are again found to be statistically significant ($p < 0.001$ for all 3 cases). This is reasonable because the LT which is constructed *a priori* based on geometry information of dwell positions (i.e., the case of gray-box optimization

with problem-specific knowledge) likely captures the linkages among dwell times better than the LT which needs to be learned from the working population (i.e., the case of black-box optimization). Moreover, since such LTs can be constructed offline, the linkage learning procedure in every generation is not needed anymore, thus enhancing the performance of MO-RV-GOMEA in terms of number of evaluations per second. In all the following experiments in this article, we therefore employ only the MO-RV-GOMEA variant with the fixed LT and partial evaluations.

6.2.2. Hypervolume results of MOEAs

Fig. 4 shows the hypervolume development graphs of 4 MOEAs over the running time (in seconds) averaged over 30 independent runs of each MOEA for all 15 patient cases. In all cases, MOEA/D is found to be the worst optimizer among the tested algorithms. We note that this might be because the default setting of the neighborhood size for solution variation (i.e., 20) is not the best for tackling our problem formulation. However, parameter tuning for MOEA/D is not within the scope of this work. MAMaLGaM almost always slightly outperforms NSGA-II (statistically significant with $p < 0.001$), except for patient case 13, for which the final hypervolume results of both algorithms are not statistically significantly different ($p = 0.93$). This suggests that the generating of offspring solutions by sampling the multivariate Gaussian distribution estimated over the solution space as done in MAMaLGaM is more effective than the solution recombination with the univariate SBX operator as done in NSGA-II. In all cases, MO-RV-GOMEA is the optimization algorithm with the best performance among the MOEAs that are considered in this article. The performance gaps between MO-RV-GOMEA and MAMaLGaM are found to be statistically significantly different ($p < 0.001$). Regarding the BT treatment planning problem, the explicit hierarchical linkage information encoded in the LT employed by MO-RV-GOMEA and its exploitation by the GOM operator, including the different improvement acceptance regimes in extreme clusters, is more beneficial to optimization than the single-level joint dependency information contained in the multivariate probability distributions of promising solutions in different clusters employed by MAMaLGaM. Furthermore, the hypervolume slopes of NSGA-II and MAMaLGaM nearly flatten out after 1 h, indicating convergence, but their obtained hypervolume values are still statistically significantly smaller than those of MO-RV-GOMEA.

Finally, there is a marked faster increase of the hypervolume as obtained by MAMaLGaM than by MO-RV-GOMEA in the initial phase (of approximately the first 10 min). In principle, if the LT model of MO-RV-GOMEA contains only the root node, the performance of MO-RV-GOMEA is similar to that of MAMaLGaM because only the Gaussian distributions concerning all problem variables jointly are considered just as in MAMaLGaM. However, due to the processing of many linkage sets of smaller sizes in solution variation, MO-RV-GOMEA has a certain overhead, causing MO-RV-GOMEA to appear slower than MAMaLGaM in the hypervolume development. Such exploitation of hierarchical linkage information rewards in the later phase for MO-RV-GOMEA in terms of obtaining Pareto fronts of better quality.

6.2.3. Pareto front results

Fig. 5 shows the Pareto fronts of non-dominated plans obtained by NSGA-II, MAMaLGaM, MOEA/D, and MO-RV-GOMEA for all 15 patient

cases. Each presented Pareto front is combined from the 30 approximation sets obtained by 30 independent optimization runs of each MOEA after 1 h by discarding dominated solutions and only keeping over-all non-dominated solutions. Note that we plot LCI objective values in terms of percentage (i.e., %) and LSI objective values in absolute terms (i.e., Gy) because it is easier for planners to interpret coverage values (LCI) in terms of how many percentages (%) of target volumes and organ sparing values (LSI) in terms of how much dose (Gy) are above/below the thresholds. Our formulations of the two objectives LCI and LSI (see Section 3) imply that treatment plans that satisfy all clinical requirements (in Table 1) should Pareto dominate the point (0, 0). Graphically, clinically-acceptable treatment plans are mapped to the top-right corner of the graphs above (0, 0), which is colored in gold. Pareto fronts shown in Fig. 5 indicate that treatment plans that satisfy all clinical requirements are achievable in all cases. The Pareto fronts obtained by MO-RV-GOMEA are always better than the ones obtained by NSGA-II, MOEA/D, and MAMaLGaM. Especially for cases 4 and 11, only the Pareto fronts obtained by MO-RV-GOMEA reach the clinically-acceptable corner. It can be clearly seen that in this *golden* corner, the solutions of MO-RV-GOMEA (Pareto-)dominate the solutions of all the other MOEAs, suggesting that exploiting the linkage information between problem variables (i.e., dwell times) benefits the optimization algorithm in reaching treatment plans of higher quality.

Table 2 shows the statistics of the hypervolume value of the approximation sets obtained after 1 h of 30 runs of NSGA-II, MOEA/D, MAMaLGaM, and MO-RV-GOMEA. MO-RV-GOMEA outperforms other MOEAs in all patient cases considered in this work and the results are found to be statistically significant.

The planning process is time-limited in practice, and should not exceed 1 h. Otherwise, patients would suffer too much from the inconvenience due to the inserted catheters when waiting too long for a plan to be made. Therefore, instead of spending the whole time budget on running optimization algorithms, it would be more beneficial to investigate the Pareto front results obtained after a shorter running time. Fig. 6 shows the Pareto fronts combined from 30 independent optimization runs of each MOEA after running for 10 min (note that given the efficiency of the CPU we used in this study (i.e., AMD processors), the efficiency of modern CPUs (e.g., Intel processors) would allow these results to be obtained in less than 5 min). In most cases (except for patient case 4), MO-RV-GOMEA obtains Pareto fronts of higher-quality treatment plans faster than both NSGA-II and MOEA/D. The plans found by MO-RV-GOMEA dominate the plans of NSGA-II and MOEA/D in the golden corner. At the timepoint of 10 min, the combined fronts obtained by MO-RV-GOMEA are slightly better than those of MAMaLGaM in some cases, and vice versa in other cases. For many cases (except for the cases 4, 7, 8, and 11), it can be seen that MO-RV-

Table 2

Means and standard deviations (in brackets) of the hypervolume value of the approximation sets obtained after 1 h of 30 runs of NSGA-II, MOEA/D, MAMaLGaM, and MO-RV-GOMEA. The best mean value for each case is presented in bold. Statistical significance is indicated with the * symbol.

Patient ID	NSGA-II	MOEA/D	MAMaLGaM	MO-RV-GOMEA
1	0.132830 (0.001459)	0.125463 (0.004070)	0.135470 (0.000837)	0.139617 (0.000737)*
2	0.154127 (0.000889)	0.145853 (0.007444)	0.156447 (0.001528)	0.161010 (0.000522)*
3	0.149913 (0.000933)	0.143773 (0.005809)	0.152653 (0.000711)	0.156027 (0.000629)*
4	0.132010 (0.001274)	0.118783 (0.009071)	0.136023 (0.002176)	0.139640 (0.001128)*
5	0.144123 (0.001196)	0.135247 (0.004102)	0.146940 (0.001750)	0.152233 (0.000778)*
6	0.143277 (0.001054)	0.134271 (0.012791)	0.144657 (0.000955)	0.147487 (0.000540)*
7	0.124253 (0.001318)	0.118257 (0.001966)	0.126880 (0.001285)	0.131733 (0.000783)*
8	0.132407 (0.001486)	0.124153 (0.004292)	0.136477 (0.001544)	0.139137 (0.001250)*
9	0.129827 (0.001179)	0.124240 (0.002442)	0.132663 (0.000683)	0.138337 (0.000742)*
10	0.144620 (0.000929)	0.137540 (0.002564)	0.145813 (0.001024)	0.148950 (0.001188)*
11	0.124663 (0.001936)	0.110032 (0.007660)	0.127560 (0.001414)	0.128910 (0.002337)*
12	0.139123 (0.002493)	0.128910 (0.005193)	0.144083 (0.001157)	0.148597 (0.000857)*
13	0.136023 (0.000871)	0.130753 (0.002573)	0.135807 (0.001712)	0.138853 (0.000950)*
14	0.152620 (0.000829)	0.147890 (0.001978)	0.153780 (0.000737)	0.156240 (0.000387)*
15	0.138803 (0.001601)	0.134413 (0.003331)	0.141740 (0.000812)	0.146903 (0.000451)*

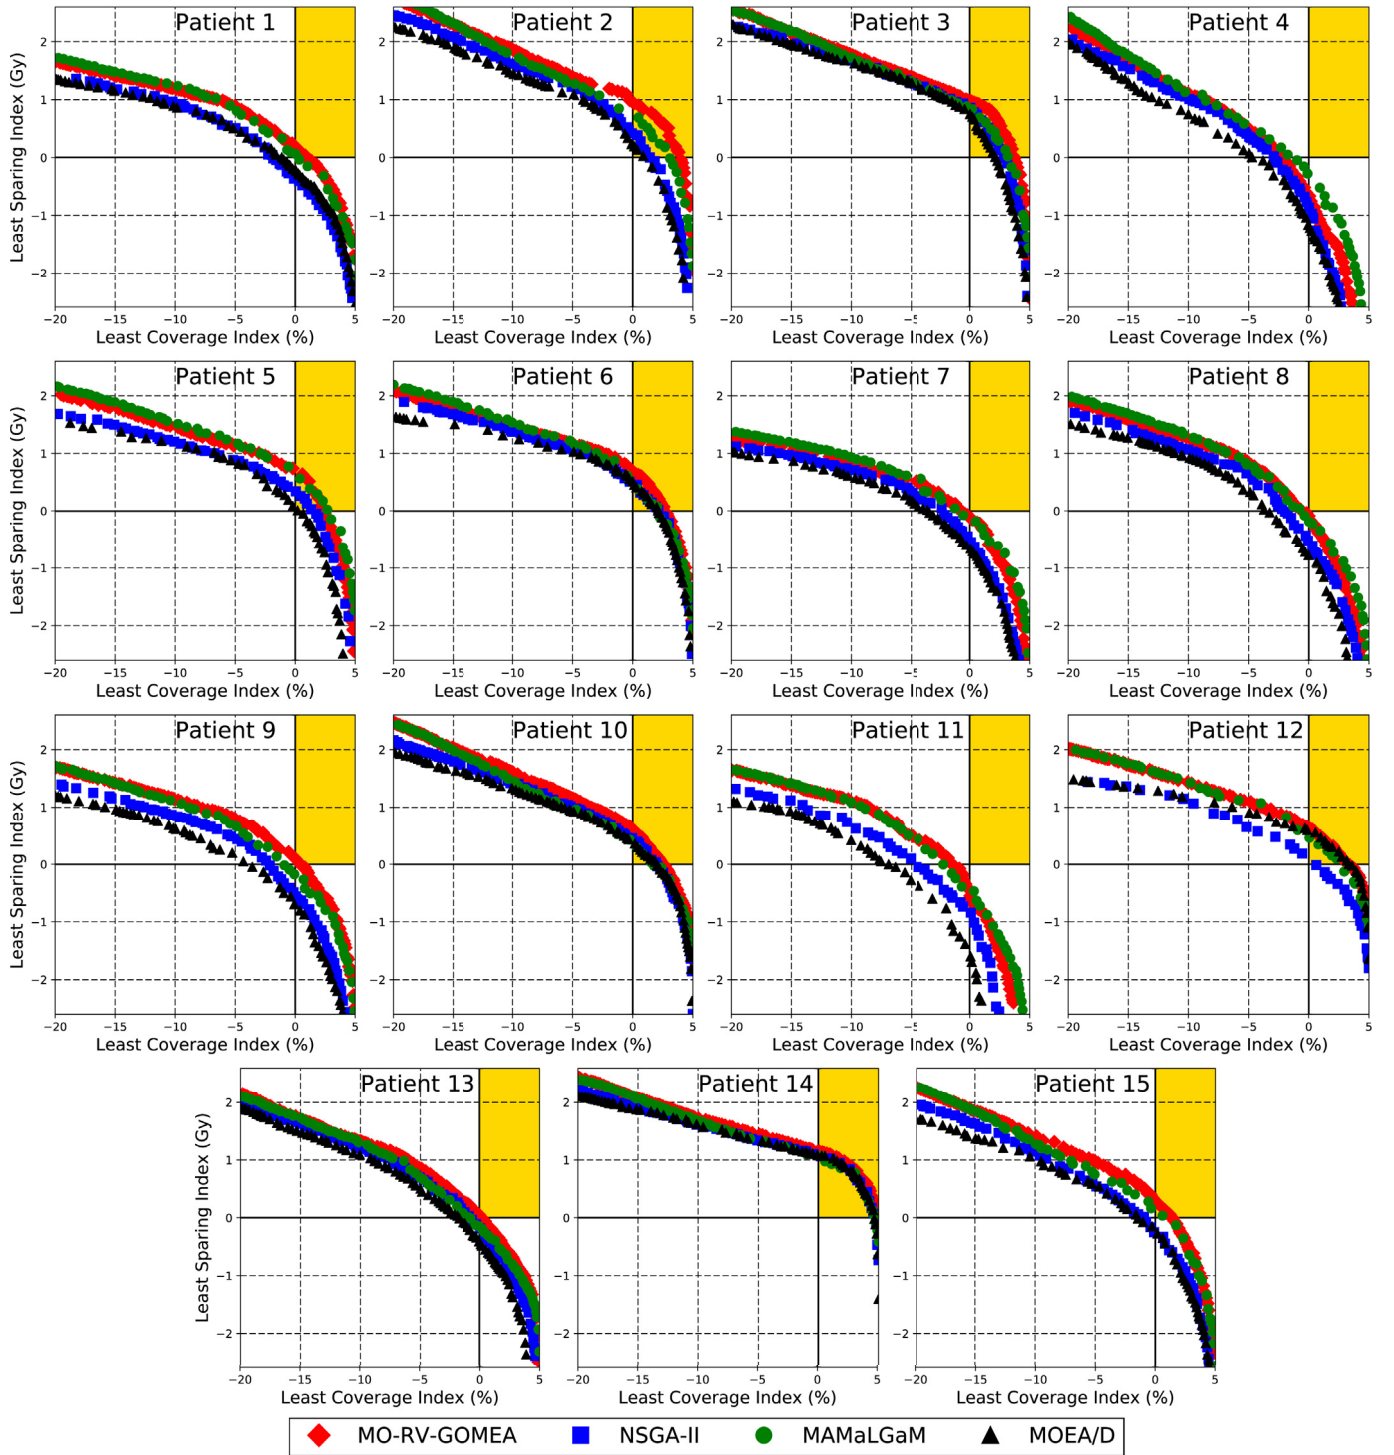


Fig. 6. Pareto fronts combined from 30 independent optimization runs of each algorithm after running for 10 min.

GOMEA and MAMaLGaM (also except for the cases 9 and 13) manage to reach the golden corner within 10 min. Also, it can be seen in Fig. 6 that the Pareto fronts obtained by MO-RV-GOMEA contain many more non-dominated solutions than all the other MOEAs, making them much smoother than those of other MOEAs. This is due to the GOM operator as used in MO-RV-GOMEA that performs solution variations by making series of local changes on existing solutions, thus substantially increasing the probability of obtaining many more non-dominated solutions, resulting in Pareto fronts of much finer-grained trade-offs. This is also the key reason why MO-RV-GOMEA is slower than MAMaLGaM in the

initial phase, which is compensated for by the Pareto fronts being of higher quality in the later phase (see Fig. 4). Note that, if planners are satisfied with the approximation set obtained along the run, instead of running for 1 h, MO-RV-GOMEA can be terminated earlier and the obtained treatment plans can be used for further fine-tuning if necessary.

Table 3 shows the statistics of the hypervolume value of the approximation sets obtained after 10 min of 30 runs of NSGA-II, MOEA/D, MAMaLGaM, and MO-RV-GOMEA. At the timepoint of 10 min, for 9 cases (i.e., cases 1, 2, 4, 5, 6, 7, 8, 9, and 11) approximation sets

Table 3

Means and standard deviations (in brackets) of the hypervolume value of the approximation sets obtained after 10 min of 30 runs of NSGA-II, MOEA/D, MAMaLGaM, and MO-RV-GOMEA. The best mean value for each case is presented in bold. Statistical significance is indicated with the * symbol.

Patient ID	NSGA-II	MOEA/D	MAMaLGaM	MO-RV-GOMEA
1	0.117026 (0.003911)	0.105189 (0.015535)	0.130043 (0.001323)*	0.125323 (0.004591)
2	0.139672 (0.002529)	0.116762 (0.018681)	0.150410 (0.001765)*	0.147510 (0.004258)
3	0.141886 (0.002380)	0.121563 (0.020112)	0.148293 (0.001226)	0.149350 (0.001870)*
4	0.115622 (0.004238)	0.094063 (0.015273)	0.129770 (0.002671)*	0.121663 (0.004976)
5	0.125469 (0.003793)	0.102237 (0.021891)	0.139440 (0.002771)*	0.132886 (0.003793)
6	0.130736 (0.004367)	0.107361 (0.027323)	0.140517 (0.001297)*	0.137896 (0.003035)
7	0.110006 (0.002995)	0.093991 (0.017202)	0.120817 (0.002232)*	0.114923 (0.003385)
8	0.118473 (0.002811)	0.104159 (0.011790)	0.130467 (0.002593)*	0.125929 (0.003773)
9	0.115002 (0.002735)	0.095677 (0.014362)	0.125670 (0.001476)*	0.123630 (0.003827)
10	0.135623 (0.003064)	0.127846 (0.007444)	0.141347 (0.001720)	0.142683 (0.002115)
11	0.099698 (0.008109)	0.073793 (0.015856)	0.121020 (0.002723)*	0.114173 (0.005679)
12	0.114146 (0.010589)	0.104691 (0.022149)	0.137163 (0.001638)	0.135273 (0.004709)
13	0.128640 (0.002153)	0.112556 (0.013397)	0.131167 (0.002277)	0.132290 (0.001915)
14	0.143879 (0.003240)	0.135038 (0.019455)	0.149813 (0.001107)	0.149963 (0.002525)
15	0.121262 (0.006081)	0.102923 (0.022943)	0.136447 (0.001511)	0.134479 (0.004638)

obtained by MAMaLGaM have better hypervolume values than those obtained by MO-RV-GOMEA and the results are found to be statistically significant. However, the final hypervolume results after 1 h in Table 2 and the hypervolume development graphs over running time in Fig. 4 indicate that given longer running time (e.g., more than 10 min), MO-RV-GOMEA catches up and outperforms MAMaLGaM in all cases. As explained earlier in Section 6.2.2, the initial slower performance of MO-RV-GOMEA compared to MAMaLGaM is due to the overhead incurred by the processing of many more linkage sets of smaller sizes during solution variation of MO-RV-GOMEA (see Section 4.4).

We note that it is still beneficial to also obtain parts of the Pareto-optimal front outside the golden corner. First, there might exist no treatment plan that satisfies the whole clinical protocol due to the quality of the catheter implants or the geometry of the involved organs, which is difficult to recognize before performing the optimization process (see Section 3). Second, depending on the specific situation of the patient and the preferences of the treating physician, the desired trade-off could be outside the golden corner. In practice, treatment planners can approve a plan located outside the golden corner. It is therefore important to provide the planners with well-spread Pareto fronts of non-dominated treatment plans so that the planners have more insight into the available trade-offs for each patient.

6.2.4. Variance of Pareto front results

Figs. 5 and 6 show the fronts of non-dominated solutions combined from 30 independent runs of each algorithm. In practice, however, typically only one optimization run would be performed. Hence, these results are overly optimistic and it is therefore important to investigate the consistency of each algorithm in terms of the variance in its obtained Pareto fronts. To this end, Fig. 7 shows, for the first three patient cases, all 30 fronts of non-dominated solutions obtained by each optimization algorithm after running for 10 min (blue-green color) and 1 h (red color). A complete overview of the results of all patient cases can be found in the [supplementary material](#). The objective space is divided into a grid and the results are represented as 2D histograms of density, such that the darker the color of a cell is, the higher the frequency that a solution can be found at that position, i.e., it conveys how frequent treatment plans with a certain quality can be obtained. Intuitively, thinner and darker curves of Pareto fronts indicate less variance, and thus higher algorithmic consistency and reliability.

MOEA/D is shown to exhibit the most variance in its results and the performance is thus the least robust. For all algorithms, a higher number of non-dominated solutions are found in the region associated with optimizing solely the LSI objective. This extreme region contains treatment plans that would result in under-irradiation, which are normally

easy to obtain by simply making dwell times as small as possible. The two MOEAs NSGA-II and MO-RV-GOMEA obtain Pareto fronts which are more evenly-distributed among different regions. Fig. 7 clearly confirms that MO-RV-GOMEA has superior consistency in its performance compared to other algorithms. In all cases, it can be seen that MO-RV-GOMEA obtains more non-dominated solutions with higher density than other MOEAs at both timepoints of 10 min and 1 h. Planners then only need to run MO-RV-GOMEA once instead of having to repeat optimization multiple times to choose the best results as in the cases of other MOEAs. Moreover, further speed enhancements of MO-RV-GOMEA are possible. The most straightforward of which are using a more modern CPU and a multi-resolution scheme for the number of dose calculation points such that this is increased over time, introducing more precision in dose calculations as the algorithm nears the Pareto-optimal front.

6.3. Discussions on decision making strategies

Solving our multi-objective BT treatment planning problem formulation results in an approximation set of diverse treatment plans that exhibit possible trade-offs between the coverage of target volumes and the sparing of OARs. Decision making strategies can then be applied to select a single treatment plan that is deemed the most preferable to be carried out. One such strategy is to employ a utility function to calculate the utility values of all solutions in the resulting approximation set and the solution with the best utility value can then be selected. A simple utility function is to calculate the Euclidean distance in the objective space from each trade-off solution to the *utopian* point (i.e., the point that represents both optimal target coverage and optimal OARs sparing) [13]. A more sophisticated utility function specific to BT that can be used is the Conformal Index (COIN), which relates to how much target volumes and how much healthy tissues outside target volumes are covered by a certain dose level. It further relates to how much OAR volumes are irradiated at a dose level higher than their clinical criteria [10]. However, such utility functions have not been widely used in practice and it has been pointed out that utility functions alone are not sufficient for insightful decision making in BT treatment planning because valuable trade-off information cannot be fully captured by utility functions [10].

A more well-informed decision making strategy is the direct navigation of the Pareto front formed by the solutions in the final approximation set. Since our problem formulation has just 2 objectives, Pareto fronts can be easily visualized and treatment planners can straightforwardly locate the desired trade-off as well as inspect how one objective needs to be compromised in order to improve the other objective. Direct navigation on Pareto fronts is more informative because each patient

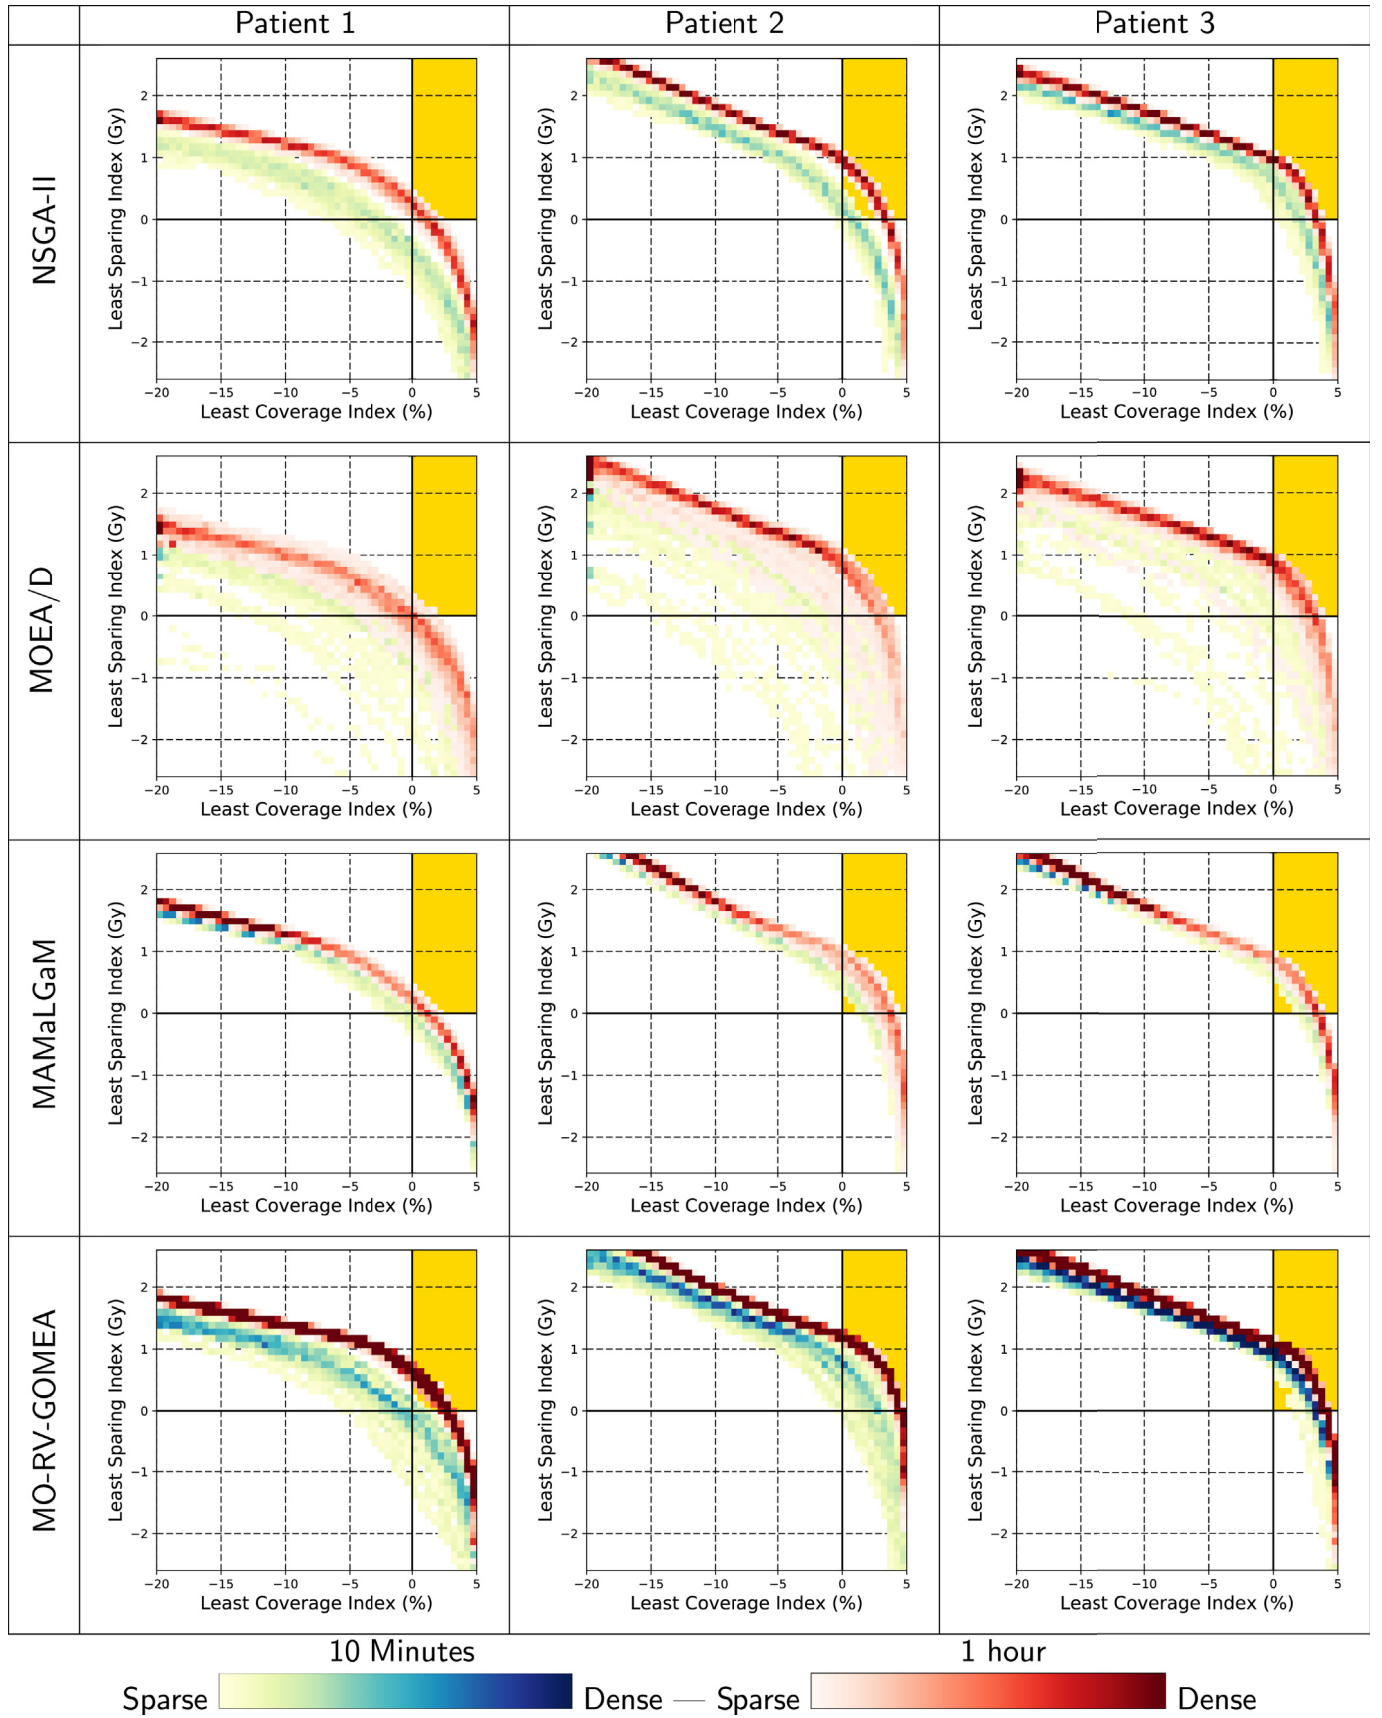


Fig. 7. Histogram of solution distributions of 30 optimization runs of each algorithm after running for 10 min (blue-green color) and 1 h (red color). Darker colors indicate high density. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

case requires a different trade-off between target coverage and organ sparing (e.g., depending on the age, health status, or organ geometry) and treating physicians, based on their experience, can have different subjective preferences regarding what a suitable treatment plan should be (e.g., a plan with good target coverage and sufficient organ sparing versus a plan with sufficient target coverage and good organ sparing). Such information cannot be straightforwardly converted into a utility function. As the final approximation set can contain many treatment plans, inspecting the whole Pareto front may take a lot of assessment time. A more efficient alternative is to present treatment planners with a smaller set of potential treatment plans (e.g., 5 or 10 plans) that are representative of the final Pareto front (e.g., a subset of evenly-distributed plans). Planners can then select a suitable plan from this subset of treatment plans to be carried out. Final decision making in BT treatment planning, however, requires further input from radiation oncologists and is left for future work.

7. Conclusions

BT treatment planning is an inherently multi-objective problem with multiple conflicting criteria of DV indices. Most optimization tools in real-world clinically available software for BT, however, treat the problem in a single-objective manner and do not optimize directly the DV indices used in the clinic for treatment plan evaluation. We argue that the BT treatment planning problem with its original DV criteria should be handled in the true multi-objective manner so that treatment plans satisfying the clinical protocol as much as possible can be obtained and that treatment planners can investigate possible trade-offs between DV indices before selecting a treatment plan that exhibits the desired trade-off for the patient case at hand. We have therefore studied and compared the use of MOEAs for real-world BT treatment planning. Moreover, we have proposed a novel bi-objective formulation for the BT treatment planning problem that closely matches intuition and typical clinical protocol. All treatment criteria in the clinical protocol were classified into two groups, namely Target Coverage and Organ Sparing, which were then used to define two corresponding optimization objectives, namely Least Coverage Index and Least Sparing Index, and other related constraints. The obtained Pareto fronts can be straightforwardly visualized and are easy to interpret and evaluate against the clinical protocol. In addition to a selection of well-known MOEAs (for black-box optimization), we considered the recently-introduced MO-RV-GOMEA, an MOEA that exploits hierarchical linkages between problem variables and is especially suited for gray-box optimization (a typical real-world scenario) by employing partial evaluations when existing solutions are altered in only a few variables. Combining the exploitation of linkage knowledge based on geometry information of the catheter implant and the performing of partial evaluations, MO-RV-GOMEA was found to be capable of outperforming other MOEAs, such as NSGA-II, MOEA/D, and MAMaLGaM, in obtaining BT treatment plans of higher quality with much higher consistency in less runtime. Experimental results on a data set of 15 prostate cancer patient cases confirmed that MO-RV-GOMEA is a promising MOEA for real-world applications, such as BT treatment planning, that is worth further research and wider applications.

Conflicts of interest

Dr. T. Alderliesten, Dr. A. Bel, and Dr. P.A.N. Bosman are involved in projects supported by Elekta.

Acknowledgement

This work is part of the research programme IPPSI-TA with project number 628.006.003, which is financed by the Netherlands Organisation for Scientific Research (NWO) and Elekta (628.006.003). The authors would like to thank Dr. Niek van Wieringen (Department

of Radiation Oncology, Academic Medical Center, Amsterdam, The Netherlands) for his support in providing the patient data, Dr. Rob van der Laarse (Quality Radiation Therapy B.V., Zeist, The Netherlands) for his advice on radiation dose calculation, and Ernst Kooreman for his contribution in this work.

Appendix A. Supplementary data

Supplementary data related to this article can be found at <https://doi.org/10.1016/j.swevo.2017.12.003>.

References

- [1] J.L.M. Venselaar, D. Baltas, P. Hoskin, A.S. Meigooni, Introduction and innovations in brachytherapy, in: *Comprehensive Brachytherapy, Imaging in Medical Diagnosis and Therapy*, Taylor & Francis, 2012, pp. 3–8, <https://doi.org/10.1201/b13075-3>.
- [2] A.M. Dinkla, R. van der Laarse, E. Kaljouw, B.R. Pieters, K. Koedooder, N. van Wieringen, A. Bel, A comparison of inverse optimization algorithms for HDR/PDR prostate brachytherapy treatment planning, *Brachytherapy* 14 (2) (2015) 279–288, <https://doi.org/10.1016/j.brachy.2014.09.006>.
- [3] E. Lessard, J. Pouliot, 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) (2001) 773–779, <https://doi.org/10.1118/1.1368127>.
- [4] A. Karabis, S. Giannouli, D. Baltas, HIPO: a hybrid inverse treatment planning optimization algorithm in HDR brachytherapy, *Radiother. Oncol.* 76 (2005) S29, [https://doi.org/10.1016/S0167-8140\(05\)81018-7](https://doi.org/10.1016/S0167-8140(05)81018-7).
- [5] R. Alterovitz, E. Lessard, J. Pouliot, I.-C.J. Hsu, J.F. O'Brien, K. Goldberg, Optimization of HDR brachytherapy dose distributions using linear programming with penalty costs, *Med. Phys.* 33 (11) (2006) 4012–4019.
- [6] T. Siau, A. Cunha, A. Atamtürk, I.-C. Hsu, J. Pouliot, K. Goldberg, IPiP: a new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices, *Med. Phys.* 38 (7) (2011) 4045–4051, <https://doi.org/10.1118/1.3598437>.
- [7] B.L. Gorissen, D. den Hertog, A.L. Hoffmann, Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR brachytherapy, *Phys. Med. Biol.* 58 (4) (2013) 1041–1057, <https://doi.org/10.1088/0031-9155/58/4/1041>.
- [8] L. De Boeck, J. Beliën, W. Egyed, Dose optimization in high-dose-rate brachytherapy: a literature review of quantitative models from 1990 to 2010, *Oper. Res. Health Care* 3 (2) (2014) 80–90, <https://doi.org/10.1016/j.orhc.2013.12.004>.
- [9] R. van der Laarse, P.A.N. Bosman, Dose optimization, in: D.Y. Song, B.R. Pieters, K. Tanderup (Eds.), *Emerging Technologies in Brachytherapy*, CRC Press, Boca Raton, Florida, 2017, pp. 79–98, Ch. 6.
- [10] M. Lahanas, D. Baltas, N. Zamboglou, A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy, *Phys. Med. Biol.* 48 (3) (2003) 399.
- [11] K. Deb, *Multi-objective Optimization Using Evolutionary Algorithms*, John Wiley & Sons, Inc., 2001.
- [12] M. Lahanas, D. Baltas, N. Zamboglou, Anatomy-based three-dimensional dose optimization in brachytherapy using multiobjective genetic algorithms, *Med. Phys.* 26 (9) (1999) 1904–1918, <https://doi.org/10.1118/1.598697>.
- [13] M. Lahanas, N. Milickovic, D. Baltas, N. Zamboglou, Application of multiobjective evolutionary algorithms for dose optimization problems in brachytherapy, in: E. Zitzler, L. Thiele, K. Deb, C.A. Coello Coello, D. Corne (Eds.), *Evolutionary Multi-criterion Optimization: First International Conference, EMO 2001 Zurich, Switzerland, March 7–9, 2001 Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 574–587, https://doi.org/10.1007/3-540-44719-9_40.
- [14] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Trans. Evol. Comput.* 6 (2) (2002) 182–197, <https://doi.org/10.1109/4235.996017>.
- [15] A. Bouter, N.H. Luong, C. Witteveen, T. Alderliesten, P.A.N. Bosman, The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm, in: *GECCO'17: Proceedings of the Genetic and Evolutionary Computation Conference, ACM, New York, NY, USA, 2017*, pp. 537–544, <https://doi.org/10.1145/3071178.3071274>.
- [16] D. Thierens, P.A.N. Bosman, Optimal mixing evolutionary algorithms, in: *GECCO'11: Proceedings of the Genetic and Evolutionary Computation Conference, ACM, New York, NY, USA, 2011*, pp. 617–624, <https://doi.org/10.1145/2001576.2001661>.
- [17] A. Bouter, T. Alderliesten, C. Witteveen, P.A.N. Bosman, Exploiting linkage information in real-valued optimization with the real-valued gene-pool optimal mixing evolutionary algorithm, in: *GECCO'17: Proceedings of the Genetic and Evolutionary Computation Conference, ACM, New York, NY, USA, 2017*, pp. 705–712, <https://doi.org/10.1145/3071178.3071272>.
- [18] N.H. Luong, A. Bouter, M.C. van der Meer, Y. Niatsetski, C. Witteveen, A. Bel, T. Alderliesten, P.A.N. Bosman, Efficient, effective, and insightful tackling of the high-dose-rate brachytherapy treatment planning problem for prostate cancer using evolutionary multi-objective optimization algorithms, in: *Genetic and Evolutionary Computation Conference, GECCO Companion'17, Berlin, Germany*,

- July 15–19, 2017, ACM, New York, NY, USA, 2017, pp. 1372–1379, <https://doi.org/10.1145/3067695.3082491>.
- [19] P.A.N. Bosman, The anticipated mean shift and cluster registration in mixture-based edas for multi-objective optimization, in: GECCO'10: Proceedings of the Genetic and Evolutionary Computation Conference, ACM, New York, NY, USA, 2010, pp. 351–358, <https://doi.org/10.1145/1830483.1830549>.
- [20] Q. Zhang, H. Li, MOEA/D: a multiobjective evolutionary algorithm based on decomposition, *IEEE Trans. Evol. Comput.* 11 (6) (2007) 712–731, <https://doi.org/10.1109/TEVC.2007.892759>.
- [21] P.J. Hoskin, A. Colombo, A. Henry, P. Niehoff, T.P. Hellebust, F.A. Siebert, G. Kovacs, GEC/ESTRO recommendations on high dose rate afterloading brachytherapy for localised prostate cancer: an update, *Radiother. Oncol.* 107 (3) (2013) 325–332, <https://doi.org/10.1016/j.radonc.2013.05.002>.
- [22] R. Nath, L.L. Anderson, G. Luxton, K.A. Weaver, J.F. Williamson, A.S. Meigooni, Dosimetry of interstitial brachytherapy sources: recommendations of the AAPM radiation therapy committee task group no. 43, *Med. Phys.* 22 (2) (1995) 209–234, <https://doi.org/10.1118/1.597458>.
- [23] M.J. Rivard, B.M. Coursey, L.A. DeWerd, W.F. Hanson, M. Saiful Huq, G.S. Ibbott, M.G. Mitch, R. Nath, J.F. Williamson, Update of AAPM task group no. 43 report: a revised AAPM protocol for brachytherapy dose calculations, *Med. Phys.* 31 (3) (2004) 633–674, <https://doi.org/10.1118/1.1646040>.
- [24] M.J. Rivard, W.M. Butler, L.A. DeWerd, M.S. Huq, G.S. Ibbott, A.S. Meigooni, C.S. Melhus, M.G. Mitch, R. Nath, J.F. Williamson, Supplement to the 2004 update of the AAPM task group no. 43 report, *Med. Phys.* 34 (6 Part1) (2007) 2187–2205, <https://doi.org/10.1118/1.2736790>.
- [25] T.M. Deist, B.L. Gorissen, High-dose-rate prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing, *Phys. Med. Biol.* 61 (3) (2016) 1155–1170, <https://doi.org/10.1088/0031-9155/61/3/1155>.
- [26] K. Deb, An efficient constraint handling method for genetic algorithms, *Comput. Methods Appl. Mech. Eng.* 186 (2–4) (2000) 311–338, [https://doi.org/10.1016/S0045-7825\(99\)00389-8](https://doi.org/10.1016/S0045-7825(99)00389-8).
- [27] N.H. Luong, P.A.N. Bosman, Elitist archiving for multi-objective evolutionary algorithms: to adapt or not to adapt, in: *Parallel Problem Solving from Nature - PPSN XII - 12th International Conference*, Taormina, Italy, September 1–5, 2012, Proceedings, Part II, 2012, pp. 72–81, https://doi.org/10.1007/978-3-642-32964-7_8.
- [28] J. Knowles, D. Corne, Properties of an adaptive archiving algorithm for storing nondominated vectors, *IEEE Trans. Evol. Comput.* 7 (2) (2003) 100–116, <https://doi.org/10.1109/TEVC.2003.810755>.
- [29] D. Thierens, D.E. Goldberg, Mixing in genetic algorithms, in: *5th International Conference on Genetic Algorithms*, ICGA 1993, 1993, pp. 38–47.
- [30] M. Pelikan, K. Sastry, D.E. Goldberg, Multiobjective hBOA, clustering, and scalability, in: *GECCO'05: Proceedings of the Genetic and Evolutionary Computation Conference*, ACM Press, New York, New York, USA, 2005, pp. 663–670, <https://doi.org/10.1145/1068009.1068122>.
- [31] N.H. Luong, H. La Poutré, P.A.N. Bosman, Multi-objective gene-pool optimal mixing evolutionary algorithms, in: *GECCO'14: Proceedings of the Genetic and Evolutionary Computation Conference*, ACM, New York, NY, USA, 2014, pp. 357–364, <https://doi.org/10.1145/2576768.2598261>.
- [32] A. Kraskov, H. Stögbauer, P. Grassberger, Estimating mutual information, *Phys. Rev. E* 69 (2004):066138 <https://doi.org/10.1103/PhysRevE.69.066138>.
- [33] I. Gronau, S. Moran, Optimal implementations of UPGMA and other common clustering algorithms, *Inf. Process. Lett.* 104 (6) (2007) 205–210.
- [34] G.R. Harik, F.G. Lobo, A parameter-less genetic algorithm, in: *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-volume 1*, Morgan Kaufmann Publishers Inc, 1999, pp. 258–265.
- [35] E. Zitzler, L. Thiele, Multiobjective optimization using evolutionary algorithms — a comparative case study, in: A.E. Eiben, T. Bäck, M. Schoenauer, H.P. Schwefel (Eds.), *Parallel Problem Solving from Nature PPSN V: 5th International Conference*, Amsterdam, The Netherlands September 27–30, 1998 Proceedings, Springer Berlin Heidelberg, 1998, pp. 292–301, <https://doi.org/10.1007/BFb0056872>.