Model-based evolutionary algorithms for finding diverse high-quality solutions

with an application in brachytherapy for prostate cancer

S.C. Maree

Cover art: *Hill-Valley Clustering*, by S.C. Maree
Layout design: S.C. Maree
Printed by: Ipskamp Printing
ISBN: 978-94-6421-234-1

The research described in this thesis was primarily carried out at the department of Radiation Oncology in the Amsterdam UMC, University of Amsterdam, Amsterdam, the Netherlands.

This work is part of the research programme IPPSI-TA with project number 628.006.003 which is financed by the Dutch Research Council (NWO) and Elekta Brachytherapy, Veenendaal, the Netherlands, and carried out in collaboration with the Centrum Wiskunde & Informatica, Amsterdam, The Netherlands. Printing of this thesis was made possible with support from Elekta.

Model-based evolutionary algorithms for finding diverse high-quality solutions

with an application in brachytherapy for prostate cancer

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. K.I.J. Maex
ten overstaan van een door het College voor Promoties ingestelde
commissie, in het openbaar te verdedigen in de Aula der Universiteit
op woensdag 17 maart 2021, te 14.00 uur

door Stefanus Cornelis Maree
geboren te Bergschenhoek

Promotiecommissie

Promotores:

Prof. dr. P.A.N. Bosman	Technische Universiteit Delft
Prof. dr. C.R.N. Rasch	AMC-UvA

Copromotores:

Dr. T. Alderliesten	Leids Universitair Medisch Centrum
Dr. A. Bel	AMC-UvA

Overige leden:

Prof. dr. L.J.A. Stalpers	AMC-UvA
Prof. dr. M.B. van Herk	AMC-UvA
Prof. dr. T.H.W. Bäck	Universiteit Leiden
Dr. C. Koedooder	AMC-UvA
Dr. A. Auger	Inria

Faculteit der Geneeskunde

"Hoe harder de strijd, des te mooier de overwinning."

Contents

Preface

This thesis describes the scientific journey I started in the beginning of 2016. To help you understand its scope, it is good to know that it was part of a larger research project. My work was mainly conducted at the Academic Medical Center of the Amsterdam University Medical Centers. Together with industry partner Elekta, they provided us with a challenging problem in treatment planning of high-dose-rate brachytherapy for prostate cancer. My dear colleagues and close collaborators at Centrum Wiskunde & Informatica, the national research institute for mathematics and computer science in the Netherlands, proposed an algorithmic approach to solve this problem. This approach is based on their, and by now also my, extensive experiences with model-based evolutionary algorithms. This thesis covers both the fundamental design of model-based evolutionary algorithms for finding diverse high-quality solutions, and their application in the treatment of prostate cancer with brachytherapy. This led to a cheerful moment in the spring of 2020, when the first prostate cancer patient was treated according to a plan constructed via our approach. That being said, I hope you enjoy reading this thesis as much as I enjoyed its research.

Stef Maree
November 2020

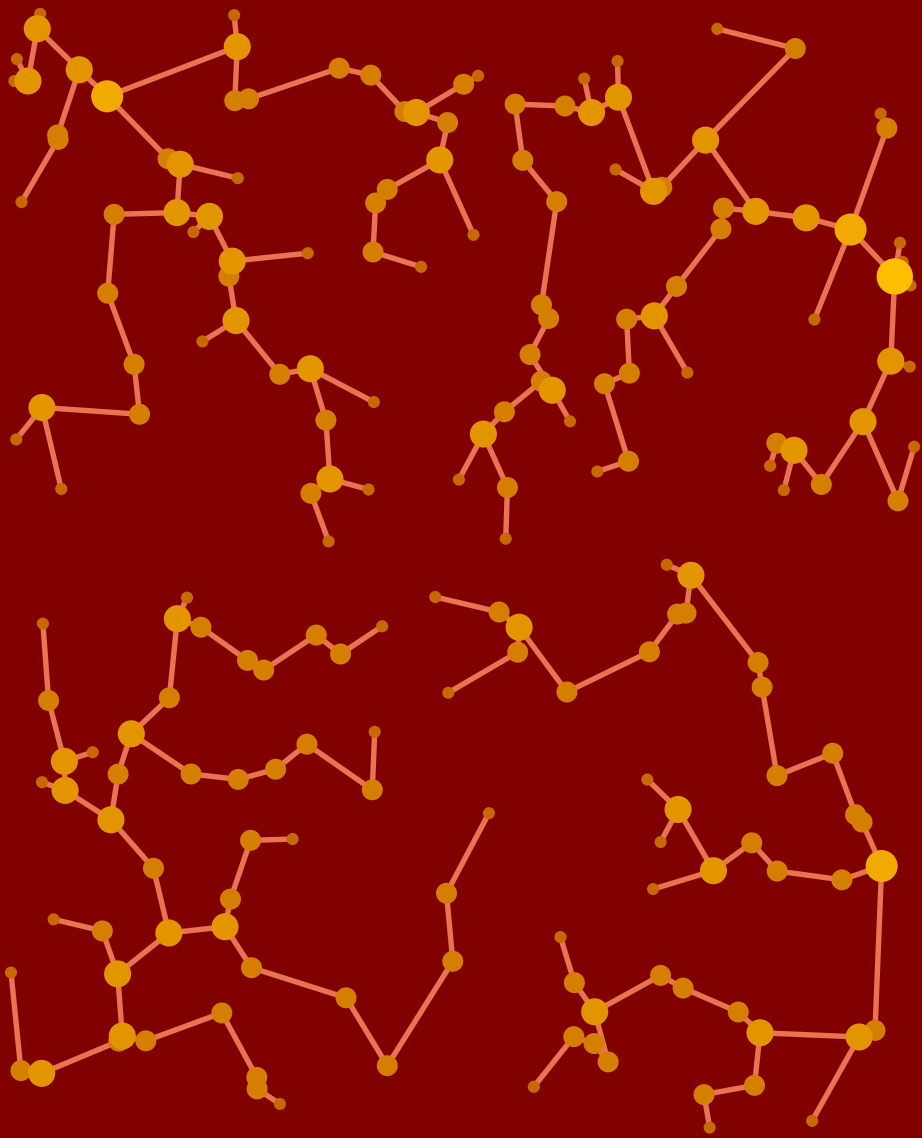

Summary

(Nederlandse samenvatting op pagina 231)

Model-based evolutionary algorithms for finding diverse high-quality solutions

with an application in brachytherapy for prostate cancer

Our society increasingly relies on the leveraging of data and algorithms to solve high-impact optimization problems. These problems are often fairly complex, making it hard to gain a sufficient understanding of their internal structure in order to design a problem-specific optimization approach or to apply exact optimization algorithms. To still be able to solve these problems, a *metaheuristic* can be employed. Metaheuristics are search algorithms that can be applied without assuming any problem specific knowledge. They are however, as the name implies, heuristics, meaning that there are typically no guarantees of finding the optimal solution. To search efficiently, metaheuristics have a certain search bias, meaning that they attempt to exploit certain problem features. When the problem at hand does not exhibit features that can be exploited, the search can be very inefficient, or result in only low-quality solutions. *Model-based metaheuristics* mitigate this risk by exploiting problem features through learnable models. The use of such models makes it explicit which problem features can be exploited. By doing so, high-quality solutions can be obtained with some certainty for the class of problems that properly aligns with the model being used.

Metaheuristics are a valuable tool for real-world optimization, in which it is often more important that a high-quality solution is obtained quickly than that the obtained solution is provably optimal. In a real-world scenario, it can however be difficult to capture all desirable aspects of a solution into a quantitative objective function. This is particularly true when the desired solution is a trade-off between conflicting aspects, such as the price-quality trade-off in economy; the bias-variance trade-off in machine learning; or the trade-off between an effective treatment and side effects in medicine. It is for such problems often unknown beforehand how these trade-offs should manifest in a desired solution. It is thus not uncommon that the

solution resulting from optimization is not as desired. To overcome this, we focus in this thesis on the design, development, and application of metaheuristics for finding not one solution, but a set of *diverse* high-quality solutions. By explicitly searching for diverse high-quality solutions, and comparing them afterwards, the normally implicit trade-offs of an optimization problem can be made explicit. This insight can help decision makers to select the most desirable solution for their problem. *Evolutionary algorithms* are sampling-based metaheuristics that are naturally well suited for this task of finding multiple high-quality solutions, as they already maintain a population of solutions to guide the search process. What then remains is to promote a form of diversity within this population during optimization. We achieve this via two different approaches: *multi-objective optimization*, and *multimodal optimization*.

In this thesis, we apply model-based evolutionary algorithms to an optimization problem that arises in the treatment of prostate cancer with high-dose-rate *brachytherapy*. Brachytherapy is a form of internal radiation therapy, in which a high dose of radiation is used to kill tumor cells. In doing so, radiation dose to nearby healthy tissue can however not be avoided. During the planning of the treatment it is determined how the radiation dose should be delivered such that the tumor is irradiated as much as possible, while surrounding tissue is spared as much as possible. These conflicting aims make treatment planning inherently a multi-objective optimization problem, for which a number of conflicting planning criteria are used in the clinic. Clinically-available methods for brachytherapy treatment planning use a weighted-sum approach to combine simplified planning criteria into a single-objective optimization problem. These methods must however be fine-tuned manually to obtain a treatment plan with desirable trade-offs between the non-simplified criteria. To overcome this unintuitive and time-consuming process, we use a recently introduced bi-objective planning model. This model is based directly on the clinical planning criteria and can be optimized efficiently with the Multi-Objective Gene-pool Optimal Mixing Evolutionary Algorithm (MO-GOMEA). By doing so, a set of diverse treatment plans is obtained, in which each plan has a different trade-off between radiation dose to the tumor and dose to nearby healthy tissue. A physician can then select the most desirable plan from this set. To aid in the selection, the set of plans can be visualized as an insightful trade-off curve, in which it can directly be seen which plans satisfy all planning criteria.

In Chapter 2, we show that the bi-objective planning model can easily be configured based on the planning criteria used in our clinic and that the optimization of this model represents our clinical practice well. An observer study among experienced physicians demonstrates that plans obtained via bi-

objective treatment planning can be considered clinically acceptable. These plans are retrospectively even preferred over the clinically used plan in 98% of the cases. The observers furthermore highly appreciate the possibility to compare high-quality plans with diverse trade-offs and consider the results to be insightful.

In Chapter 3, we apply a multi-objective evolutionary algorithm to automatically tune the parameters of two clinically available methods for brachytherapy treatment planning. Tuning is aimed at maximizing the objectives of the bi-objective planning model. We show that by doing so, treatment plans with good objective values can be obtained. However, direct optimization of the bi-objective planning model with MO-GOMEA results in treatment plans with better objective values. Per-patient automatic parameter tuning is furthermore too time-consuming for clinical practice. We therefore constructed standard parameter values, i.e., *class solutions*, from the automatically tuned plans. These class solutions are a good starting point for manual parameter tuning, but cannot entirely overcome it.

In Chapter 4, we develop a new approach to efficient and effective multimodal optimization. The aim of multimodal optimization is to obtain all global optima of an optimization problem. This is typically achieved by exploring multiple modes, that is, the high-fitness regions in the search space, hence the name multimodal optimization. In a black-box scenario, it is essential that no assumptions are made on the number of modes, or on their shape or size. We introduce hill-valley clustering that can be used to identify modes. Hill-valley clustering uses the simple hill-valley test to determine whether two solutions belong to the same hill (mode) by checking if there is a valley in between. The Hill-Valley Evolutionary Algorithm (HillVallEA) is a two-phase method that combines hill-valley clustering with an evolutionary algorithm. HillVallEA is state of the art in multimodal optimization and is twofold winner of the multimodal optimization competition held annually at the Genetic and Evolutionary Computation Conference.

HillVallEA is extended to multi-objective optimization in Chapter 5. In multimodal multi-objective optimization, the aim is to obtain all globally optimal solutions, i.e., all solutions in the *Pareto set*. This adds a layer of complexity to multi-objective evolutionary algorithms for multimodal optimization, as a balance needs to be found between diversity within a solution set (i.e., different trade-offs), and among solution sets (i.e., different modes). We show that the multi-objective HillVallEA outperforms other multi-objective optimization algorithms in multimodal optimization on a set of multimodal benchmark problems. Furthermore, and perhaps most importantly, we show that it is capable of obtaining, maintaining, and improving solution sets in multiple modes simultaneously.

Most well-known evolutionary algorithms for multi-objective optimization are

based on Pareto dominance, in which the fitness of a solution is determined by the number of solutions in the population it dominates. A fundamental limitation of this approach is however that these methods stagnate when the entire population is non-dominated. Indicator-based multi-objective optimization can overcome this by formulating the problem as a high-dimensional single-objective optimization problem. In Chapter 6, we show how the hypervolume indicator can be adapted in order to take dominated solutions into account, resulting in the Uncrowded Hypervolume (UHV) indicator. To be able to efficiently solve the high-dimensional single-objective problem, we exploit the known structure of the UHV with the single-objective GOMEA. We show that the resulting UHV-GOMEA can converge to a subset of the Pareto set, while maintaining a population of solutions, thereby being able to outperform other methods on multimodal multi-objective problems.

Multi-objective optimization provides a decision maker with a set of diverse high-quality solutions, from which the most desirable solution can be selected. Traversing the corresponding trade-off curve provides a natural ordering of these solutions. This order does however not necessarily map to a smooth trajectory in decision space. This forces the decision maker to inspect the decision variables of all solutions individually, which can make the selection of the most desirable solution time consuming and unintuitive. In Chapter 7, we use the UHV-based problem formulation to explicitly search for solution sets that are smoothly navigable. For this, solutions sets are parameterized as smooth Bézier curves in the decision space. We solve the Bézier problem formulation with GOMEA, which we refer to as BezEA. BezEA shows to be competitive, while smooth navigability is guaranteed. We furthermore show that BezEA can efficiently solve the bi-objective brachytherapy treatment planning problem by exploiting that plans on a Bézier curve can be efficiently evaluated. Obtained differences in objective values are small, suggesting that BezEA is a good alternative to MO-GOMEA for bi-objective treatment planning when navigational smoothness is desired.

Concluding, the contribution of this thesis is twofold. First, we validate bi-objective treatment planning for high-dose-rate prostate brachytherapy. Second, we contribute to the general knowledge on algorithms for finding diverse high-quality solutions. We develop the simple yet effective hill-valley clustering method, which can be employed for efficient and effective multimodal optimization. In addition, we advance the field of multi-objective optimization by demonstrating how efficient hypervolume-based multi-objective optimization with convergence to optimality can be accomplished. Finally, we show how this approach can be used to obtain smoothly navigable solution sets, which can further assist the decision maker in selecting the most desirable solution for the problem at hand.

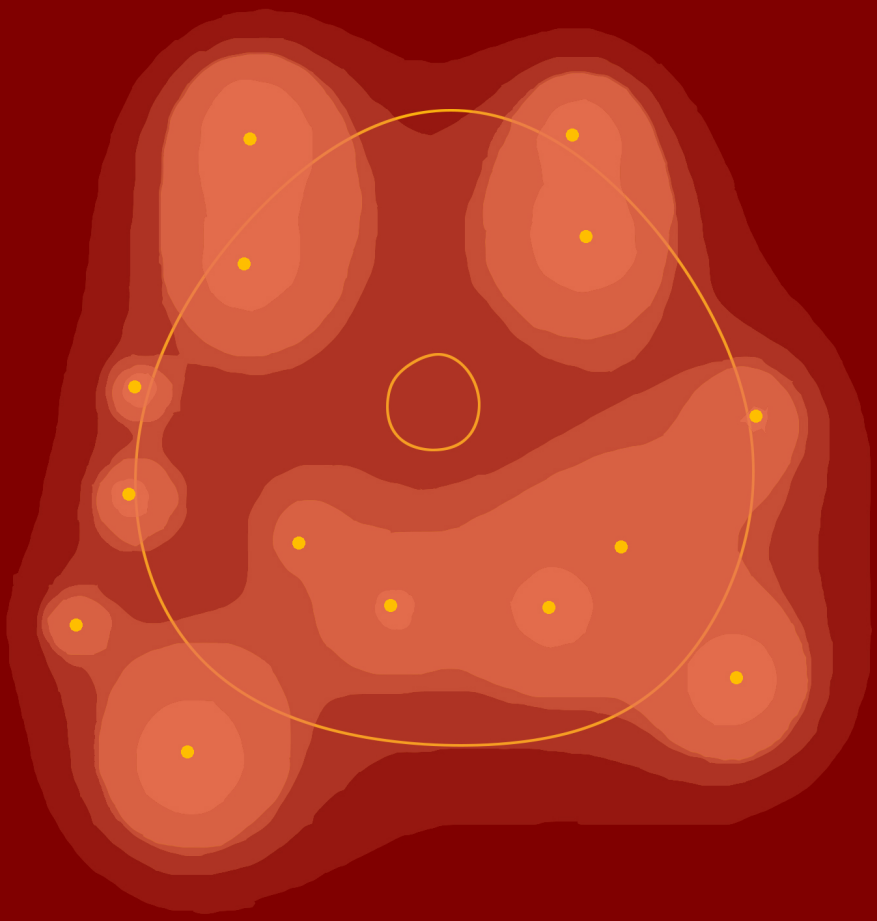

1 Introduction

1.1 Black-box optimization

Our society increasingly relies on the leveraging of data and algorithms to solve high-impact optimization problems. These problems are often fairly complex, making it hard to gain a sufficient understanding of their internal structure in order to design a problem-specific optimization approach or to apply exact optimization algorithms [39]. To still be able to solve these problems, a *black-box* perspective can be taken, in which it is assumed that beforehand nothing is known about the internal structure of the problem. In other words, the problem is viewed as a black box, for which some input parameters can be set, and an output is computed, but whatever happens inside this box is unknown. In such a scenario, a *metaheuristic* can be employed. Metaheuristics are search algorithms that can be applied without assuming any problem-specific knowledge [129]. They are however, as the name implies, heuristics, meaning that there are typically no guarantees of finding the optimal solution. The search strategy employed by a metaheuristic is usually designed to efficiently explore the *search space* of all possible solutions. To ensure efficiency, metaheuristics have a certain search bias, meaning that they attempt to exploit certain problem features. A drawback is that when the problem at hand does not exhibit features that can be exploited, the search can be very inefficient, or result in only low-quality solutions. A way to mitigate this risk

is via the use of *model-based metaheuristics*. Model-based metaheuristics exploit problem features that can be captured through learnable models. The use of such models makes it explicit which problem features can be exploited. By doing so, high-quality solutions can be obtained with some certainty for the class of problems that properly aligns with the model being used. Moreover, using a learnable model typically enlarges the class of problems that can be efficiently handled as the model can be adapted to the problem during optimization.

Metaheuristics are a valuable tool for real-world optimization, in which it is often more important that a high-quality solution is obtained quickly than that the obtained solution is provably optimal. Additionally, in a *grey-box* scenario, when parts of the internal structure of the problem are known, metaheuristics provide a good framework in which the additional information can be exploited to further improve the search [88, 129, 182].

Even though metaheuristics can be applied in a black-box scenario without the need for a deep understanding of the problem, a quantitative objective function that measures the quality of a solution is still required. For real-world optimization problems, it is not always straightforward to formulate this objective function such that it properly quantifies all important aspects of the problem at hand. There are often multiple aspects that need to be taken into account, and these aspects can be subjective, hard to quantify, or computationally expensive to evaluate. It is in a black-box scenario often unknown beforehand which aspects will be prevalent in high-quality solutions. It is thus not uncommon that afterwards, when the obtained solution is inspected by a domain expert, it is not as desired [182].

A straightforward approach to obtain more-desirable solutions is to adapt the objective function to better reflect practical desires, and repeat the optimization process with the newly constructed objective. This can however be difficult and time consuming, and the list of aspects that needs to be taken into account can be endless. An alternative approach is to present a set of *diverse* high-quality solutions to the domain expert, of which each of these solutions is of high-quality, but with different underlying aspects. This set of diverse solutions can then be used to gain insight in the problem structure and allows the expert to use so-far unformulated objectives or external information to selecting a desired solution. Such unformulated objectives typically relate to an intuition of what solutions should look like, or what qualities they should have, but are hard to formulate exactly. Once there are however alternatives to choose from, such qualities may be intuitively very clear to a domain expert.

If it is known upfront that multiple solutions should be obtained, optimization approaches can be employed that are more efficient than simply restarting and

repeating algorithms that aim for a single solution. In this thesis, we design, study, and apply two techniques to find a set of diverse high-quality solutions: *multi-objective optimization* and *multimodal optimization*. Both of these techniques serve a different practical purpose, as we will describe below. However, let us first discuss the class of algorithms we employ to do so: model-based evolutionary algorithms.

1.2 Model-based evolutionary algorithms

The optimization problems we consider in this thesis are numerical optimization problems, of which the to-be-optimized decision variables are real-valued. Whenever the derivatives of the real-valued objective function are available, it is almost always beneficial to use them in the optimization process to guide the search [6]. Well-known iterative methods that exploit derivative information are for example Adam gradient descent [110] and the Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm (L-BFGS) [43]. However, not in all cases are the derivatives of the objective function available. In a black-box scenario, the formulation of the objective function is assumed to be unknown, and the derivatives are thereby also unknown. Even when the precise formulation of the objective function is known, it can be computationally expensive to compute its derivatives, they could be of no use because the objective function exhibits noise, or they could simply not exist. In that case, we enter the field of *derivative-free optimization*. In derivative-free optimization, the derivatives of the objective function are not used to guide the optimization process. Derivative-free optimization and black-box optimization for numerical problems are very similar concepts, but have a different historical background. As a mathematical study, derivative-free optimization focuses on methods that can be mathematically analysed to prove convergence or to determine stopping criteria. Metaheuristics for black-box optimization are historically designed to be fast, but are generally unsupported by any rigorous convergence analysis, which is why they are named *heuristics*. We focus on mainly metaheuristics here, and make no distinction between derivative-free algorithms and black-box algorithms.

Metaheuristics can roughly be categorized into point-based methods and set-based methods. Point-based methods modify a single candidate solution to sequentially improve upon it, potentially making use of previously evaluated solutions. Point-based methods include naive approaches such as random search [178], exhaustive search and grid search [6], and the more advanced approaches such as Bayesian optimization [156], simulated annealing [170], tabu search [74], and the Nelder-Mead method [163].

Set-based methods on the other hand maintain and improve multiple solutions simultaneously by using population characteristics to guide the search. Set-based methods include genetic and evolutionary algorithms [11], and swarm intelligence [106]. The distinction between these methods can however be subtle or not clear at all. The main difference between these methods can be found in their origin, and the subsequent terminology. Evolutionary algorithms can be roughly described as inspired by Darwin's theory of evolution, and borrow a lot of terminology from the field of biology. Evolutionary algorithms maintain a set or *population* of solutions that are simultaneously improved over the course of multiple iterations or *generations*. Solutions that have a better objective value or *fitness* value (these terms are used interchangeably in this thesis) are more likely to produce *offspring* solutions, in a survival-of-the-fittest fashion. As we continue a research line on evolutionary algorithms in this thesis, we will use terminology from this field also to describe set-based methods that arose from a different background.

A commonly used approach to generate new solutions in numerical optimization with evolutionary algorithms is via sampling [33, 86, 206]. A probability distribution is used to model high-fitness regions in decision space, or to model a probable direction of improvement. Evolutionary algorithms that use such a probabilistic model to guide the search can be classified as *model-based evolutionary algorithms* or *estimation-of-distribution algorithms* [90, 121]. The distribution is adapted over time by combining information from high-fitness solutions, and is used to subsequently sample new solutions from. The Gaussian distribution is often used as search distribution, mainly due to its simplicity [31, 86], but alternatives such as beta distributions [206] or Bayesian networks [166] have been proposed as well. Successful model-based evolutionary algorithms are the Covariance Matrix Adaptation Evolution Strategies (CMA-ES) [86], Natural Evolution Strategies (NES) [206], the Adapted Maximum-likelihood Gaussian Model Iterated Density-estimation Evolutionary Algorithm (AMaLGaM) [31], and the related Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) [37]. In this thesis, we make extensive use of AMaLGaM, and the related GOMEA. These algorithms were chosen due to their good performance [31, 37] and a robust design that allows for making various adaptations without disrupting the internal mechanisms, and thereby performance, too much.

1.2.1 AMaLGaM

Let $f : \mathcal{X} \to \mathbb{R}$ be a to-be-minimized objective function, where $\mathcal{X} \subseteq \mathbb{R}^n$ is the n-dimensional search space, and $\mathbf{x} \in \mathbb{R}^n$ a solution vector. AMaLGaM maintains a population $\mathcal{P}^t = \{\mathbf{x}_1^t, \mathbf{x}_2^t, \ldots, \mathbf{x}_N^t\}$ of N solutions $\mathbf{x}_i^t$ that is updated as the number

of generations $t \geq 1$ increases. The population is generally initialized using uniform sampling, or using problem-specific simple heuristics. Then, selection is performed, which can be represented by assigning a weight $w_i^t \in [0, 1]$ to each solution $\mathbf{x}_i^t$ such that $\sum_{i=1}^{N} w_i^t = 1$. In AMaLGaM, truncation selection is used, where the $\lfloor 0.35 \cdot N \rfloor$ solutions with best fitness $f(\mathbf{x}_i^t)$ all get equal weight, and the non-selected solutions get zero weight. Maximum likelihood estimation [81, 89] is used to estimate the mean vector $\boldsymbol{\mu}^t$ and covariance matrix C^t from the best solutions in the population,

$$\boldsymbol{\mu}^t = \sum_{i=1}^{N} w_i^t \cdot \mathbf{x}_i^t, \quad \text{and,} \quad C^t = \sum_{i=1}^{N} w_i^t \cdot (\mathbf{x}_i^t - \boldsymbol{\mu}^t)(\mathbf{x}_i^t - \boldsymbol{\mu}^t)^T.$$

New solutions $\mathbf{x}_i^{t+1}$ are then sampled from the Gaussian distribution given by,

$$\mathbf{x}_i^{t+1} \sim \mathcal{N}(\boldsymbol{\mu}^t + b_i^t \cdot \boldsymbol{\mu}_{\text{shift}}^t, c^t \cdot C^t).$$

Here, two additional terms have been added. The first addition is the *anticipated mean shift* $\boldsymbol{\mu}_{\text{shift}}^t$ [31], for which $b_i^t \geq 0$ specifies whether it is applied to that solution (when $b_i^t > 0$), and the magnitude of the shift. The second addition is the distribution multiplier $c^t \geq 0$, which prevents the algorithm from premature convergence by enlarging the search distribution if deemed necessary using a mechanism called the *standard deviation ratio* [31]. The best solution in the current generation, the *elite*, is always maintained, which is a process called *elitism*. Adaptations such as elitism can violate (implicit) model assumptions, as the population is no longer solely sampled from the currently estimated distribution [22]. AMaLGaM was found to be particularly robust to these types of adaptations, making it a very practical algorithm to build further extensions upon.

There are two versions of AMaLGaM that we equip in this thesis. In AMaLGaM-full, a Gaussian distribution with a full-covariance matrix is estimated. The advantage of using a full-covariance matrix is that it can capture dependencies between decision variables, allowing AMaLGaM-full to solve non-separable problems [31]. A downside of this full-covariance distribution is that a Cholesky decomposition is required to sample from it, which has a computational complexity of $\mathcal{O}(n^3)$. Additionally, a large population size N is required in order to accurately estimate all $\frac{1}{2}n(n-1)$ (co)variances [137]. This becomes infeasible or at least very computationally expensive for high-dimensional problems.

As an alternative, AMaLGaM-univariate was introduced. In this version of AMaLGaM, only the n variances on the diagonal of the covariance matrix are estimated, and all covariances are set to zero. The number of model parameters that needs to be estimated with a univariate Gaussian distribution scales linear

in the problem dimensionality n. Additionally, AMaLGaM-univariate has a generational complexity of $\mathcal{O}(n)$, making it a good method for solving high-dimensional optimization problems with weak dependencies, but problems with strong dependencies can no longer be efficiently solved [31].

1.2.2 GOMEA

For many problems, instead of having to perform the entire objective function evaluation, it is possible to quickly update the objective value of a solution when only a few decision variables are changed, which is known as a *partial evaluation*. The real-valued GOMEA is a type of evolutionary algorithm that is excellently suited to exploit partial evaluations [37].

In GOMEA, a *linkage model* is used to specify which subsets of decision variables should be considered dependent. For each decision variable subset, a full-covariance Gaussian search distribution is maintained, based on the same mechanisms as used in AMaLGaM. In GOMEA, new values for only the decision variables specified in that subset are sampled and evaluated. Then, an intermediate selection step is added: if these newly sampled values are found to worsen the solution, the old values are restored. GOMEA can thus be interpreted as interleaved instances of AMaLGaM-full, where each instance optimizes a subset of the decision variables.

By equipping GOMEA with a linkage model that consists of a single subset that contains all decision variables, i.e., $L_{\text{full}} = \{\{1, 2 \ldots, n\}\}$, it is essentially equivalent to AMaLGaM-full. Alternatively, by equipping GOMEA with a univariate linkage model $L_{\text{uni}} = \{\{1\}, \{2\}, \ldots, \{n\}\}$ in which each decision variable is its own subset, one may be inclined to think it essentially performs a search similar as AMaLGaM-univariate. However, the intermediate selection step added to GOMEA affects the search behavior. Intermediate selection enforces an even stronger independent processing of decision variables in different subsets. To balance the advantages of these two extreme examples of a linkage model, the *linkage tree* model can be employed [194]. For a problem with $n = 4$ decision variables, an example of a linkage tree is,

$$L_{\text{tree}} = \big\{\{1\}, \{2\}, \{3\}, \{4\}, \{2, 3\}, \{1, 2, 3\}, \{1, 2, 3, 4\}\big\}.$$

The linkage tree is constructed by starting from the univariate linkage model, and iteratively merging two subsets. After each merge, the newly constructed subset is added to the linkage model. By construction, the linkage tree contains $2n - 1$ subsets of decision variables that are of different size, and of which some overlap. The merge order can be determined during optimization by exploiting observed

dependencies between decision variables in the population, or by using expert knowledge or other external information that indicates which decision variables are likely to be dependent [37]. Additional constraints can also be put on the linkage model. For example, the linkage tree can be truncated by discarding large subsets, as it is very uncommon that a problem has strong dependencies between all decision variables. The combination of the ability to exploit partial evaluations and the ability to model dependencies between decision variables without the need to estimate an $n \times n$-covariance matrix makes GOMEA state-of-the-art in high-dimensional *grey-box* optimization [37].

1.3 Multi-objective optimization

We so far considered the scenario in which there is a single objective function $f(\mathbf{x})$ that needs to be optimized. However, in many real-world scenarios, there are multiple important aspects of a problem that need to be taken into account in the optimization process. When these aspects are conflicting, it can be difficult to intuitively and properly combine them in a single objective function. One approach to construct a single objective function is by using a weighted sum approach. In this approach, each of the m objectives $f_i(\mathbf{x})$ is given a weight $w_i \in \mathbb{R}$, and a single objective function is constructed by considering their weighted average, i.e., $f(\mathbf{x}) = \sum_{i=1}^{m} w_i f_i(\mathbf{x})$. This approach, in which multiple objectives are combined into a single scalar function is called *scalarization*. It is however often unclear how to set the weights before optimization, as the desired trade-off between the objectives is beforehand often unknown. To overcome this, the optimization process can be repeated multiple times with different weights, but this is a rather inefficient approach as a very similar optimization process is repeated over and over, and it can be difficult to set the weights such that the resulting solution has desirable objectives values. It is moreover not for all problems the case that all solutions with an optimal trade-off are actually the optimum of a weighted average scalarization.

Alternatively, these aspects can be formulated as separate objective functions, resulting in a *multi-objective optimization problem* or *multi-criteria decision making problem* [52]. In multi-objective optimization, the objective function is a vector function $f : \mathbb{R}^n \to \mathbb{R}^m$, given by $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}) \ \cdots \ f_m(\mathbf{x})]$. When these objective functions are conflicting, there no longer is a single optimal solution to the problem, but a set of multiple solutions that all have a different *trade-off* between the objectives. Optimality of multi-objective problems can then be defined formally in terms of *Pareto efficiency* [113]. A solution is said to *dominate* another solution if it

is better in at least one objective, and not worse in any of the other objectives. We say that a solution is *non-dominated* with respect to a given set of solutions if it is not dominated by any solution in the given set. A solution that is *Pareto optimal* cannot be improved in any objective without worsening it in any of the other objectives, and is thus non-dominated with respect to the entire solution space. The set of all Pareto-optimal solutions is called the *Pareto set*. The corresponding set of objective values forms the *Pareto front*.

As we are working with metaheuristics, we implicitly assume that we are aiming for good approximations of the optimum, instead of the optimum itself. In case of multi-objective optimization, such an approximation is a set of solutions, rather than a single solution. Such sets are commonly referred to as *approximation sets*. The objective values of the solutions in the approximation set form the *approximation front*. The aim of multi-objective optimization is then to find an approximation front that approximates the Pareto front. Interestingly, this aim is in itself two-sided. On the one hand, the approximation front should be as close as possible to the Pareto front (proximity). On the other hand, the solutions should be a good representation of the entire Pareto front (diversity) [33].

In case of bi-objective optimization (two objectives), the approximation front can be easily visualized as a *trade-off curve*, as shown in Figure 1.1. This trade-off curve can be used by a decision maker or domain expert to gain insight in the problem-specific trade-offs and it can aid in the selection of a desired solution [113].

1.3.1 Multi-objective optimization algorithms

Population-based methods are of specific interest for multi-objective optimization as these methods already maintain a set of multiple solutions that is used to explore the decision space. However, evolutionary algorithms for single-objective optimization cannot be used directly for multi-objective optimization. First, the selection mechanism needs to be adapted, as there is no longer a strict ordering of solutions based on quality, since solutions can be mutually non-dominated. A commonly used selection mechanism for multi-objective optimization is based on the number of solutions that dominate a given solution. This is used for example in the well-known Non-dominated Sorting Genetic Algorithm (NSGA-II) [53]. The second change that is required is that single-objective evolutionary algorithms aim for a single optimal solution, where multi-objective evolutionary algorithms (MOEAs) need to obtain a set of solutions with different trade-offs. To do this effectively, the sampling distribution needs to be adapted. In this thesis, we will mainly consider the multi-objective AMaLGaM (MAMaLGaM) [181] and the multi-objective GOMEA (MO-GOMEA) [38]. These methods use a number of

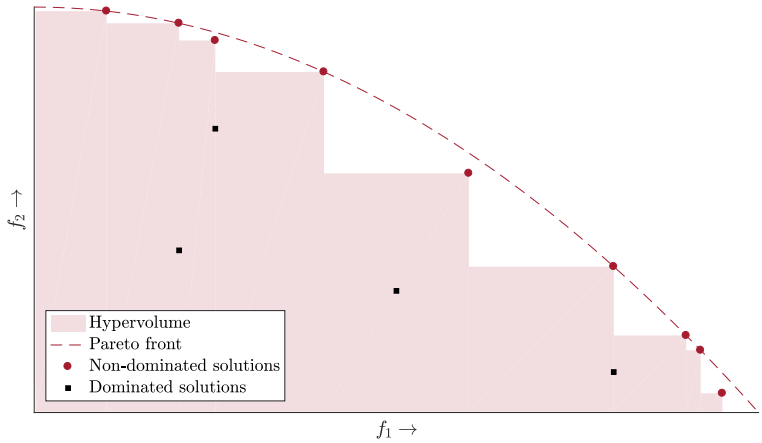

Figure 1.1: Objective space visualization of a bi-objective problem with to-be-maximized objective function $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}) \ ; \ f_2(\mathbf{x})]$. The dashed line shows the Pareto front, the theoretical optimum of this problem, which is here a continuous line segment in objective space, thereby containing an infinite number of Pareto-optimal solutions, which is common for real-valued optimization problems. Red circles represent a number of non-dominated solutions (obtained by an algorithm). The set of all obtained non-dominated solutions is called the *approximation set*, as it approximates the Pareto front. Some dominated solutions are illustrated by black squares. The shaded area represents the hypervolume [216] of the trade-off curve (i.e., the approximation front), which is a performance indicator for approximation sets.

overlapping Gaussian distributions, instead of a single one, to model the decision space. For this, the population is clustered into overlapping clusters, and based on the solutions in each cluster, the parameters of the Gaussian sample distributions are estimated, using mechanisms very similar to those of AMaLGaM.

1.4 Real-world optimization in brachytherapy for prostate cancer

We use MOEAs to solve a real-world optimization problem that arises in the treatment of prostate cancer. The prostate is a gland about the size of a walnut and is part of the male reproductive system [14]. It is located between the bladder and the penis, just in front of the rectum. The urethra runs through the center of the prostate, from the bladder to the penis, allowing urine to flow out of the body.

Cancer is the development of abnormal cells that divide uncontrollably and have the ability to infiltrate and destroy normal body tissue [198]. Cancer often has the

ability to spread throughout the body, thereby invading key organs and interfering with body functions that are necessary to live.

Prostate cancer is the most common cancer among men in Europe, and the third cause of cancer death for men [66]. In the last 25 years, the incidence of prostate cancer has increased significantly [41]. This can be attributed to the aging population, but also to an earlier detection of prostate cancer by the use of the prostate-specific antigen (PSA) blood test [44, 191]. Prostate cancer can be categorized in three categories of increasing risk [157]. For low-risk prostate cancer, watchful waiting, is a viable treatment option, as the tumor progression rate of prostate cancer is typically low [15, 111]. Intermediate or high-risk prostate cancer can be treated with radical prostatectomy, in which the entire prostate gland is removed, or with radiation therapy. In radiation therapy, or radiotherapy, a certain dose of ionizing radiation is delivered to the tumor which causes DNA damage that kills tumor cells or slows their growth. Prostate cancer cells, like many other cancerous cell types, are more sensitive to radiation than healthy cells [149]. This offers a so-called therapeutic window, which allows for the killing of cancer cells, while surrounding healthy tissue can be spared. It is however impossible to irradiate the tumor while fully preventing radiation dose to be delivered to surrounding healthy tissue. This makes the treatment of prostate cancer with radiotherapy inherently a multi-objective optimization problem.

Radiotherapy can either be applied externally in the form of external beam radiotherapy (EBRT), in which a linear accelerator is used to generate a beam of high-energy photons or protons, or internally in the form of brachytherapy, in which sources of ionizing radiation are permanently or temporarily implanted within the tumor. We distinguish two types of brachytherapy. Low-dose-rate brachytherapy is applied in the treatment of low-risk prostate cancer and uses small radioactive iodine or palladium seeds that are permanently implanted within the prostate [79, 168, 172]. For the treatment of intermediate and high-risk prostate cancer, high-dose-rate (HDR) brachytherapy is commonly applied, potentially as boost after EBRT [97–99, 117, 215]. In HDR prostate brachytherapy, on which we focus in this thesis, a set of catheters is implanted in or close to the tumor. Through the catheters, a small radioactive source (Iridium-192) is guided by an afterloader device. The longer the source dwells at certain positions in these catheters, the more the surrounding tissue is irradiated. By tuning these *dwell times*, the distribution of radiation dose can be sculpted in order to deliver an effective treatment in which the tumor is eradicated while healthy tissue is spared as much as possible. The delivery of the radiation dose itself typically takes a few minutes. After the treatment, the radioactive source is retracted into a shielded safe in the afterloader.

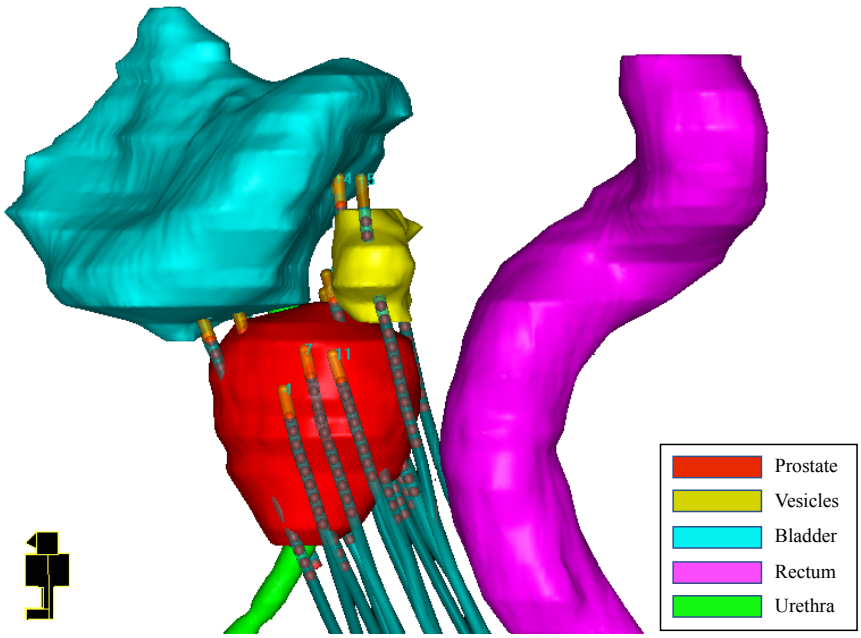

Figure 1.2: Three-dimensional reconstruction of the delineated targets and organs at risk, based on the delineated MRI scans (See Figure 1.3). Catheter tips are denoted in yellow, and dwell positions as red spheres within the catheters. Patient orientation as in the lower left corner.

1.4.1 Brachytherapy treatment planning

Treatment planning is the process of determining how the radiation dose should be delivered, by optimizing the dwell times. In the Amsterdam University Medical Centers, location Academic Medical Center, treatment planning is performed based on Magnetic Resonance Imaging (MRI) scans of the patient that are made after implantation of the catheters. On these scans, a physician, together with a Radiation Therapy Technologist (RTT) delineates: the target volumes, which are the entire prostate gland and possibly the base of the seminal vesicles; the catheters; and the Organs At Risk (OARs), which are for this application the urethra, bladder, and rectum; as shown in Figure 1.2. Based on the delineations, the dose distribution can be simulated [180]. In order to determine the quality of a proposed dose distribution or *treatment plan*, it can be visualized as a heatmap that is projected on top of an MRI scan, as shown in Figure 1.3. Besides a visual inspection of the dose distribution, a number of quality indicators is used to quantify different aspects of a dose distribution, which we refer to as Dose-Volume Indices (DVIs) [149]. DVIs are specified for a given region of interest, which can be either a target volume or

an OAR. Multiple studies have suggested relations between certain DVIs and local tumor control [92, 148] or adverse effects [26, 100, 190]. Based on these relations, a treatment protocol is recommended by the GEC-ESTRO. The GEC-ESTRO is a group of experts and practitioners of brachytherapy and is an amalgamation of the The Groupe Européen de Curiethérapie[1] (GEC) and the European SocieTy for Radiotherapy & Oncology (ESTRO). The treatment protocol used in our clinic is based on recommendations by the GEC-ESTRO, and describes the minimally desired dose to the target volumes and the maximally desired dose to the OARs. This results in a total of ten DVI-based *planning criteria*. As these DVIs measure different aspects of a treatment plan, their values give a good indicator of the quality of a treatment plan. However, a visual inspection is always performed to verify that there are no undesired aspects of the dose distribution that are not captured by the DVIs.

In clinical practice, planning time is limited, as schedules are usually tight, and the patient is waiting with the catheters implanted during treatment planning [77, 95, 189]. A longer treatment planning time implies that fewer patients can be treated, but also implies more discomfort for the patient, and a higher risk of catheter displacement, which might result in an incorrectly delivered treatment [70, 109, 116]. It is therefore important that good treatment plans are obtained within reasonable time.

1.4.2 Treatment planning methods

In order to obtain a desirable dose distribution, different treatment planning methods have been developed. Graphical optimization is a manual optimization tool that allows for a drag-and-drop of the iso-dose lines of the visualized dose distribution [159]. It is a valuable tool to quickly adapt the dose distribution locally, but to shape the entire dose distribution by graphical optimization can be very time consuming [60, 82, 115, 159]. Alternatively, *inverse planning methods* have been introduced. Different planning methods will be discussed in Chapters 2 and 3, but we address two inverse planning methods here that are clinically available in Oncentra Brachy (Elekta Brachytherapy, Veenendaal, the Netherlands): Inverse Planning Simulated Annealing (IPSA) [124] and Hybrid Inverse Plan Optimization (HIPO) [120]. These methods do not optimize the DVI-based planning criteria directly, but use a simplified dose-penalty model that allows for fast computations. The planner, either a physician or RTT, needs to specify a set of to-be-minimized dose requirements on different regions (targets and

[1] In French, brachytherapy is called curiethérapie, after Marie and Pierre Curie, who did pioneering research on radioactivity.

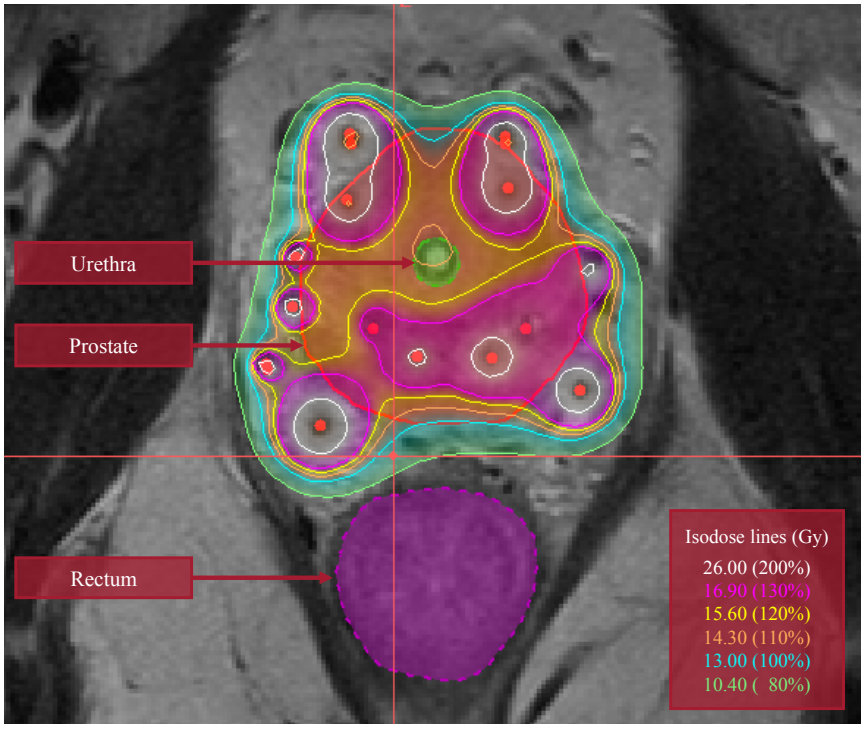

Figure 1.3: Simulated dose distribution, projected on an MRI scan of a prostate cancer patient treated with HDR brachytherapy. Iso-dose lines are shown in the color corresponding to the dosage shown in the legend, in gray (Gy) as well as in percentage of the planning-aim dose. Catheters are positioned perpendicular to the scan, and dwell positions are denoted by red dots.

OARs), and a set of corresponding penalty weights that indicates the importance of each requirement. A scalarization approach is then applied to combine all these requirements into a single objective function by considering their weighted average. In IPSA, this objective function is then optimized with simulated annealing [170], a black-box metaheuristic. HIPO exploits that the dose distribution and the simplified dose-penalty model are differentiable with respect to the dwell times [119], and uses the gradient-based L-BFGS to optimize its objective function [43].

It is in practice however unclear how to set the penalty weights before optimization, as the precise trade-off between the dose requirements is unknown, and can be different for each patient. It therefore generally happens that, after optimization, the resulting treatment plan is not as desired. One approach to overcome this is to repeat the optimization process multiple times with different weights, until a desired treatment plan is obtained [60, 154]. As this process

is performed manually, it is labor-intensive and it can be difficult to obtain desirable plans [60]. Multiple treatment planning approaches have furthermore been developed that replace manual iterative weight tuning by more intuitive approaches, such as Enhanced Geometrical Optimization with Interactive Inverse Planning (EGO-IIP) [61], as further discussed in Chapters 2 and 3. Although these methods are more intuitive to use than setting and tuning weights of inverse planning methods, it remains hard and time-consuming to get a good intuition about the nature of the underlying trade-offs this way.

1.4.3 Bi-objective treatment planning

An alternative approach to treatment planning is to formulate it as a multi-objective optimization problem [119]. In this thesis, we specifically focus on the bi-objective planning model [131]. In the bi-objective planning model, the planning criteria are grouped into one coverage objective and one sparing objective, that are referred to as the Least Coverage Index (LCI) and Least Sparing Index (LSI). Both the LCI and LSI are constructed by combining the coverage or sparing criteria in a worst-case manner. Here, a worst-case scalarization means that the value of the LCI/LSI is determined by the DVI of the coverage/sparing criteria that is most violated. These two objectives capture the biggest trade-off between all planning criteria: the trade-off between target coverage (dose maximization) and sparing of OARs (dose minimization).

By solving this bi-objective planning model, a set of treatment plans is obtained in which all plans have a different high-quality trade-off in terms of the LCI and LSI. Because the model is limited to only two objectives, these plans can be easily visualized as a patient-specific trade-off curve, as shown in Figure 1.4. An advantage of the worst-case construction of the LCI and LSI is that the trade-off curve directly shows which plans satisfy all planning criteria, which is the case when $\text{LCI} > 0$ and $\text{LSI} > 0$ holds. This corresponds to the upper right quadrant in Figure 1.4 [131]. Being able to directly see which plans satisfy all planning criteria and how they trade-off more coverage and more sparing can be a very valuable and insightful tool that could be used by the planner to select a desirable treatment plan. Bi-objective treatment planning thereby has the potential to transform treatment planning from a time-consuming manual optimization process into an insightful decision making problem.

1.4.4 Solving the bi-objective planning model

The DVIs, on which the clinical planning criteria and the LCI and LSI are based, are not differentiable, rendering gradient-based methods inapplicable for

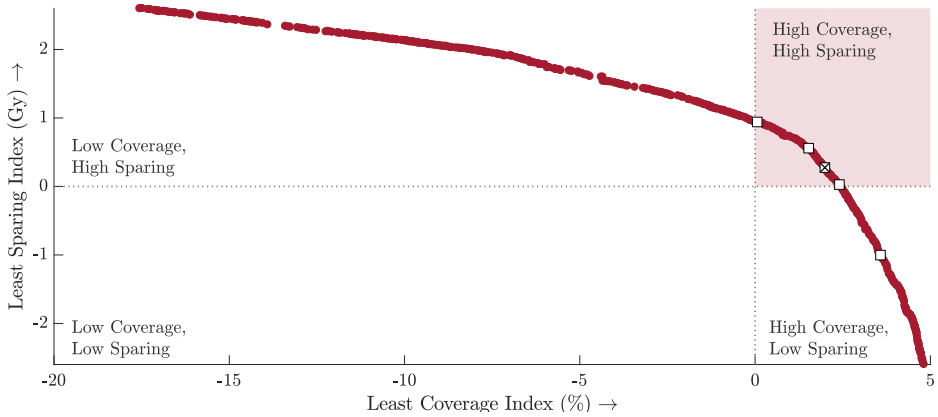

Figure 1.4: Example of a trade-off curve for a prostate cancer patient treated with brachytherapy. Each red dot represents a treatment plan. Plans in the upper right corner (shaded, with LCI > 0 and LSI > 0) satisfy all planning criteria. From this trade-off curve, five plans were selected with diverse (clinically-relevant) trade-offs, marked by the white squares. After a further inspection of the dose distributions of these five plans, the plan marked with ⊠ was selected and used as a basis to treat the patient with.

this application without simplifying or smoothing the underlying model [82]. This makes it particularly interesting to approach the bi-objective planning model with an MOEA. It was shown in [131] that MO-GOMEA was able to outperform a number of other MOEAs by obtaining the best trade-off curves in limited time. This was mainly achieved by using a problem-specific linkage tree in which dwell times corresponding to dwell positions that are geometrically nearby are considered dependent. Additionally, runtime could be significantly improved by exploiting partial evaluations, for which this problem is particularly well suited as dose contributions from different dwell positions can be updated independently. With these problem-specific adaptations to MO-GOMEA, computation time could be reduced to one hour. Although the experiments were performed on a somewhat older central processing unit (CPU), this is still too slow for clinical use. An extension, in which the resolution of the internal dose calculations is increased during optimization, reduced the required computation time to about five minutes [130]. To even further reduce computation time, the computation of the dose distribution and the DVIs as well as part of MO-GOMEA were implemented on a graphics processing unit (GPU). Using a GPU, computation time could be reduced to 30 seconds [36]. Additionally, optimization with high-resolution dose calculations could now be performed in three minutes. This makes bi-objective treatment planning sufficiently fast for clinical use.

1.4.5 Clinical acceptability

In practice, treatment planning starts with an automatically generated treatment plan, which is then manually fine-tuned by a planner. Even though the planning criteria function as guidance during treatment planning, the planner may have taken external factors into account, intentionally or even unintentionally. The planner has additional information about the patient and the disease, and may have a general intuition about desirable aspects of treatment plans. During treatment planning, the planner assesses plan quality not only by the DVI-based planning criteria, but also by a visual inspection of the dose distribution. It is thus not directly clear that optimizing the DVI-based planning criteria, as in the bi-objective planning model, is sufficient for obtaining clinically acceptable treatment plans. We address this in our first research question:

Research question 1:
Can clinically acceptable prostate brachytherapy treatment plans be obtained by optimizing the bi-objective planning model with MO-GOMEA?
(Chapter 2)

1.4.6 Tuning inverse planning methods

Inverse planning methods for brachytherapy treatment planning, such as IPSA or HIPO, which were discussed above, combine conflicting aims into a single objective by using a weighted-sum approach. The planner searches for a desirable treatment plan by adapting penalty weights of different target volumes and OARs. Even when there exist penalty weights that would result in a desirable plan, it is not straightforward to obtain these weights, especially under clinical time pressure [60]. A difficulty in comparing performance of planning methods is then that it is unclear whether the tuning had been done properly, and even if it were, the resulting plans can be mutually non-dominated, meaning that they are not better or worse than another, but just different [60].

The search for penalty weights that result in a desirable plan is an optimization problem itself. As such, the bi-objective treatment planning model can be used to automatically tune the penalty weights of an inverse planning method in order to maximize the DVIs on which the planning criteria are based. In this way, inverse planning methods can be aimed at obtaining a set of plans with diverse trade-offs. The performance of different inverse planning methods can then be compared insightfully in terms of their obtained maximally-achievable trade-off curves.

The clinically-available IPSA and HIPO do however not directly optimize the DVI-based planning criteria, but solve a simplified dose-penalty model. While this model can be adapted to a great extent by tuning penalty weights, it is unclear whether it is configurable in such a way that the same plan quality in terms of the DVI-based planning criteria can be obtained as by directly optimizing these planning criteria in the bi-objective planning model with MO-GOMEA. It is further unclear whether automatic tuning of the penalty weights of inverse planning methods can overcome manual treatment planning. We address this in the following research question:

Research question 2:
Can time-consuming manual treatment planning for prostate brachytherapy be overcome by automatic tuning of the penalty weights of clinically-available inverse planning methods? *(Chapter 3)*

1.5 Multimodal optimization

So far, we considered optimization algorithms for finding diverse solutions via multi-objective optimization, in which all solutions have a different trade-off between different conflicting objective functions. Another approach for finding diverse solutions is via *multimodal optimization*. In (single-objective) multimodal optimization, the aim is not just to obtain a single optimal solution of an optimization problem, but multiple globally optimal solutions, and potentially also high-quality locally optimal solutions, as visualized in Figure 1.5 [125].

Multimodal optimization can be a useful approach in real-world optimization when the formulation of the objective function is not straightforward or certain and may thus not be a perfect reflection of the ultimate goal. By aiming the search for a set of diverse (locally) optimal solutions with a good objective value, the domain expert can inspect the various alternatives and select the most desirable solution. Multimodal optimization is specifically interesting when the underlying problem is multimodal with a moderate number of similar-quality *modes* or *hills*. In that case, similar quality solutions with very different decision values exist, that are therefore likely to differ in aspects that are not captured by the objective function.

A naive approach to multimodal optimization is by restarting the (evolutionary) algorithm multiple times, each time initialized in a different region of the decision space. However, this approach may still have a high chance of ending up in the same

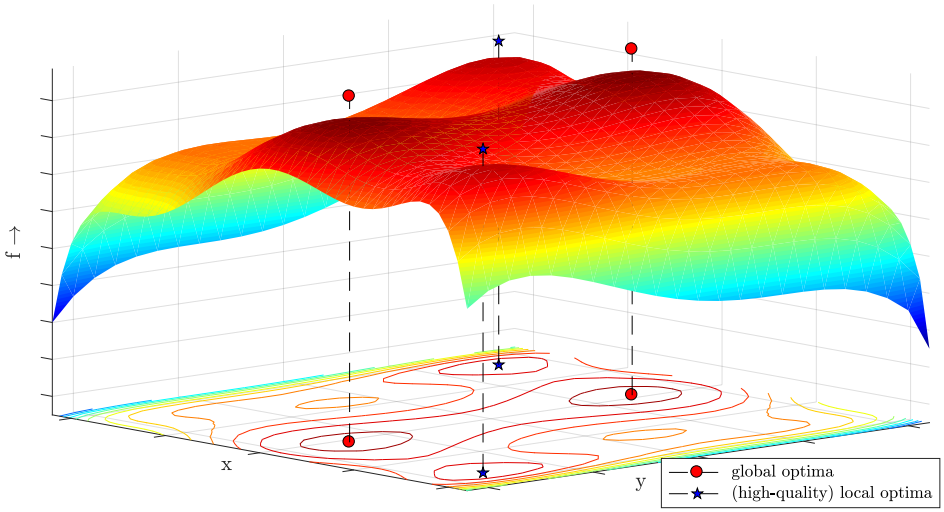

Figure 1.5: Illustration of (single-objective) multimodal optimization, in which the aim is to obtain all global optima and potentially also high-quality local optima.

optimum multiple times, especially when the size of the modes differ. Many restarts are then required before the smaller modes will be explored. *Niching* is the process of maintaining diversity in the population by spreading out the search over multiple high-fitness regions in the decision space, known as *niches*. Classically, niching was used to prevent metaheuristics from premature convergence [50, 75]. Currently, niching methods are commonly used to perform multimodal optimization [126].

A recent approach to multimodal optimization is the repelling-subpopulations method [3], in which multiple instances of a search algorithm are randomly initialized in the search space. To prevent multiple instances from converging to the same optimum, rejection sampling is used to push them away from each other, and also from previously obtained optima. As opposed to random initialization of search algorithms within the search space, two-phase methods aim in the first phase at locating different niches in the search space. In the second phase of these methods, core search algorithms are specifically initialized in these niches. The Nearest-better Evolutionary Algorithm (NEA2+) [174, 175] is one such two-phase method, which uses Nearest-Better Clustering (NBC) in the first phase to cluster an initial set of solutions. NBC is a fitness-informed clustering method that uses a distance measure based on the concept of a nearest-better solution. The idea behind NBC is that local optima can be detected by the observation that there are no *nearby* solutions with a better fitness. However, the distinction between *nearby*

and *far away* is rather difficult to make in a black-box scenario, as it depends on the fitness landscape, the number of solutions in the initial population, and the problem dimensionality. The effectiveness of this approach was consequently found to deteriorate for problems with larger problem dimensionality. This leads to the following research question:

Research question 3:
How can model-based evolutionary algorithms be adapted to perform efficient and effective (single-objective) multimodal optimization, without making assumptions on the number of modes, or on their shape or size?
(Chapter 4)

1.5.1 Multimodality in multi-objective optimization

The concept of multimodal optimization can also be applied to multi-objective problems. In case of multimodality in a multi-objective decision space, the set of Pareto-optimal solutions can contain multiple solutions that have exactly the same objective values, but differ in their decision values. MOEAs by default discard or ignore solutions that are of similar quality (or slightly worse), even if they have very different decision values [53]. This means that such MOEAs cannot obtain all Pareto-optimal solutions, making them unsuitable for *multimodal multi-objective optimization*. The aim of multimodal multi-objective optimization is to obtain all (locally) Pareto-optimal solutions, that is, all solutions in the Pareto set [107].

The challenges in multimodal multi-objective optimization are similar to those of single-objective multimodal optimization, in the sense that it is essential for effective black-box multimodal methods to make as few assumptions as possible on the number of modes, their shape, and their size. However, in contrast to single-objective optimization where the aim is to find a single solution, MOEAs already aim to obtain multiple solutions. This adds a layer of complexity to MOEAs for multimodal optimization, as a balance needs to be found between diversity in objective space and in decision space. Additionally, instead of detecting whether two solutions belong to a different mode, it is now required to detect whether two approximation sets belong to a different mode. Mechanisms that detect whether approximation sets reside in a single mode are therefore required, especially when the aim is to explicitly distinguish between multiple modes in a multimodal multi-objective decision space. Recently, a number of MOEAs for multimodal optimization (MMOEAs) have been introduced, by applying niching

techniques to existing multi-objective evolutionary algorithms. These MMOEAs are then aimed to approximate the Pareto set with a set of diverse high-quality solutions in decision space [127, 192, 211]. As we will see in Chapter 4, hill-valley clustering is a niching technique that performs very well for single-objective multimodal optimization. This makes it interesting to investigate how it can be extended to multi-objective optimization problems, which leads to the following research question:

Research question 4:
How can niching via hill-valley clustering be applied to multi-objective optimization problems to allow for efficient and effective multi-objective multimodal optimization, without making assumptions on the number of modes, or on their shape or size? *(Chapter 5)*

1.6 Hypervolume-based multi-objective optimization

Domination-based MOEAs are today arguably the most frequently used type of MOEA. These methods implicitly balance proximity to the Pareto set and diversity within the population [33]. This however tends to lead to stagnation in terms of convergence to the Pareto set when the majority of the population becomes non-dominated [18]. Hypervolume-based multi-objective optimization has shown promising results to overcome this [21, 204]. The hypervolume indicator measures the area under the trade-off curve (bounded by a reference point), as illustrated in Figure 1.1 [216]. The hypervolume indicator is rather expensive to compute for problems with more than two objectives, but it is particularly interesting as it is *Pareto-compliant* [216], meaning that the solutions in an approximation set with optimal hypervolume are Pareto-optimal. By making use of the hypervolume indicator, multi-objective optimization problems can be reformulated as a single-objective optimization problem in which the aim is to optimize the hypervolume of a fixed-size approximation set [204, 216]. However, the hypervolume indicator cannot be used directly in this application, as it does not take dominated solutions into account. As a result, this would lead to no guidance for an optimization procedure to push dominated solutions to the non-dominated front. For this reason, it was so far mainly applied as secondary selection criterion after non-dominated sorting, most notably in the $\mathcal{S}$-metric selection evolutionary multi-objective optimization algorithm (SMS-EMOA) [21]. In the recently introduced Sofomore framework

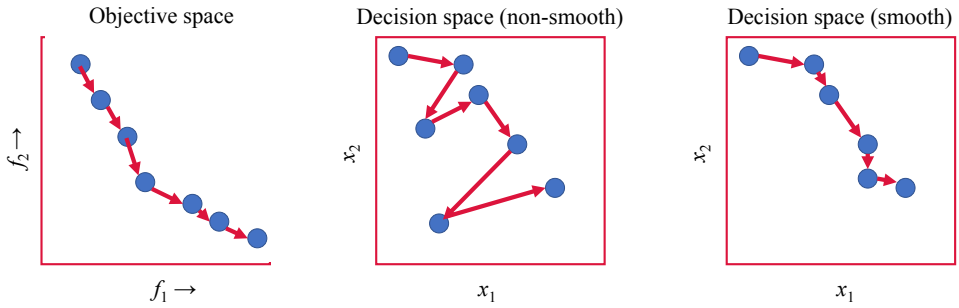

Figure 1.6: Illustration of how the trade-off curve of a set of solutions (blue circles) implies a navigational order for the decision maker in objective space in the left subfigure. The corresponding decision variables do not necessarily yield a smooth trajectory in decision space (middle and right subfigure).

[195], this is overcome by making use of the *uncrowded distance* [63], which we will discuss in detail in Chapter 6. To handle the high dimensionality of the resulting problem, it was proposed to formulate it as a set of multiple interleaved single-objective dynamic problems, that iteratively improve a single approximation set. However, as this implies that actually only a single approximation set is being updated and optimized iteratively, it loses many advantages of population-based multi-objective optimization, such as escaping local optima in multimodal problem landscapes. This gives rise to the following research question:

Research question 5:
How can evolutionary multi-objective optimization approaches be designed where the individuals in the population represent entire approximation sets and fitness is directly based on the hypervolume indicator or extensions thereof so as to ensure convergence to a subset of the Pareto set?
(Chapter 6)

1.6.1 Navigational smoothness of the approximation set

The aim of bi-objective optimization is to obtain an approximation set of high-quality solutions. A decision maker can navigate this set to select a desired solution, potentially using the visualized trade-off curve. The trade-off curve provides an intuitive navigational ordering of solutions to traverse, e.g., from left to right as illustrated in Figure 1.6. This ordering does however not necessarily map to a smooth trajectory through decision space. Especially when the problem at hand is multimodal, solutions with similar trade-offs can be very different in terms of their

decision values. If decision makers are interested in the values of the decision variables, they are then forced to inspect the decision values of each solution individually, which can be very time consuming. Additionally, when solutions seem very similar in terms of their objective values, but are then found to be very different in terms of their decision values, the selection of a solution with desirable properties is no longer intuitive, which potentially further complicates the selection.

Imposing a form of smoothness or continuity in terms of decision variables between solutions in the approximation set as a restriction upon the population of an MOEA is not straightforward and has, to the best of our knowledge, not yet been done. A reason for this is that control over approximation sets as a whole is needed. However, typical domination-based MOEAs use single-solution-based mechanics. Hypervolume-based methods can offer a solution here. In these methods, solution sets are directly parameterized, which allows for natural extensions or restrictions upon this parameterization. This gives rise to the following research question:

Research question 6:
How can smoothly navigable approximation sets be obtained efficiently and effectively via hypervolume-based bi-objective optimization, thereby making the selection of a desirable solution more intuitive for a decision maker? *(Chapter 7)*

Navigational smoothness in brachytherapy treatment planning

In brachytherapy treatment planning, it is time-consuming to inspect the dose distribution of a plan to verify its quality. Due to clinical time constraints, it is therefore not feasible to inspect many different plans. If the dose distribution of plans vary in a *smooth* and predictable manner when the trade-off curve is traversed, it might not be required to inspect all plans individually. This could make plan selection easier and more intuitive. Since there is a direct relation between the dwell times of a plan and the dose distribution, this smooth navigation can be imposed upon the decision variables (dwell times) of solutions. This results in the following final research question:

Research question 7:
Are smoothly navigable approximation sets for bi-objective treatment planning of prostate brachytherapy obtainable without too much cost in obtained trade-offs and computation time? *(Chapter 7)*

1.7 Outline of this thesis

This thesis can be roughly categorized in three partially-overlapping pillars: brachytherapy treatment planning, multi-objective optimization, and multimodal optimization, as schematically visualized in Figure 1.7. A reader interested in one particular pillar can follow one of the indicated paths.

In Chapters 2, 3, and 7, we focus on Research questions 1, 2, and 7, which are related to different aspects of brachytherapy treatment planning. In Chapters 4 and 5, we focus on Research questions 3 and 4, which are related to multimodal optimization. Further, (hypervolume-based) multi-objective optimization is covered in Chapters 6 and 7 in which we aim to answer Research questions 5 and 6. Finally, in Chapter 8, we discuss in general the results obtained in this thesis.

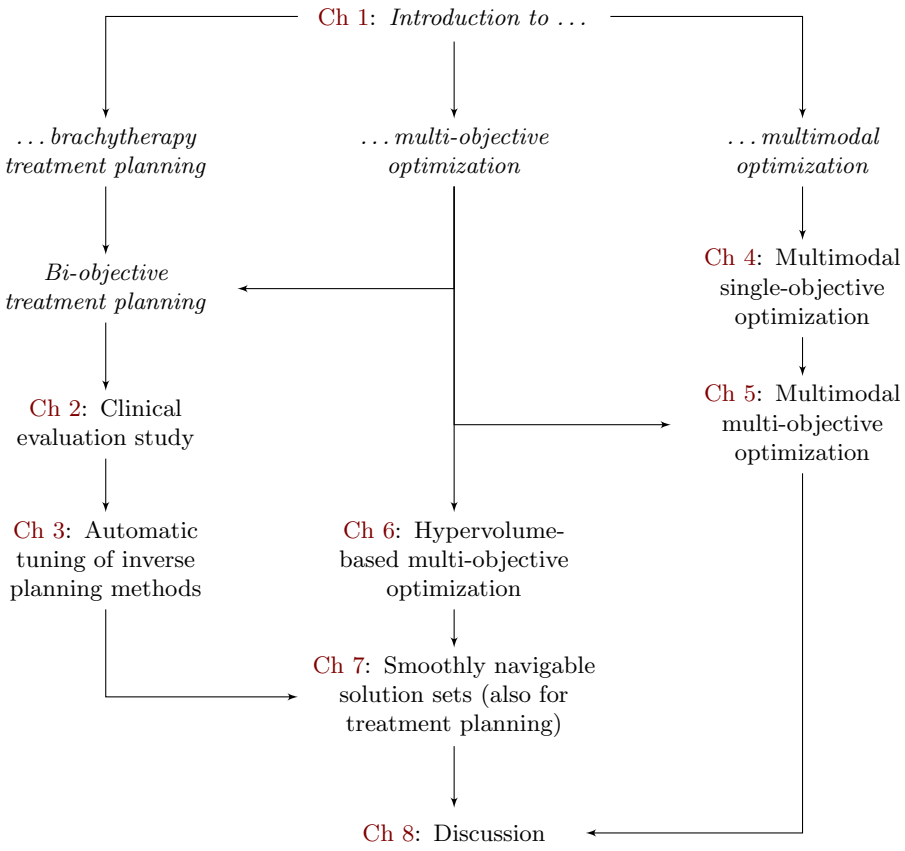

Figure 1.7: Outline of this thesis. Arrows indicate relations between chapters and topics.

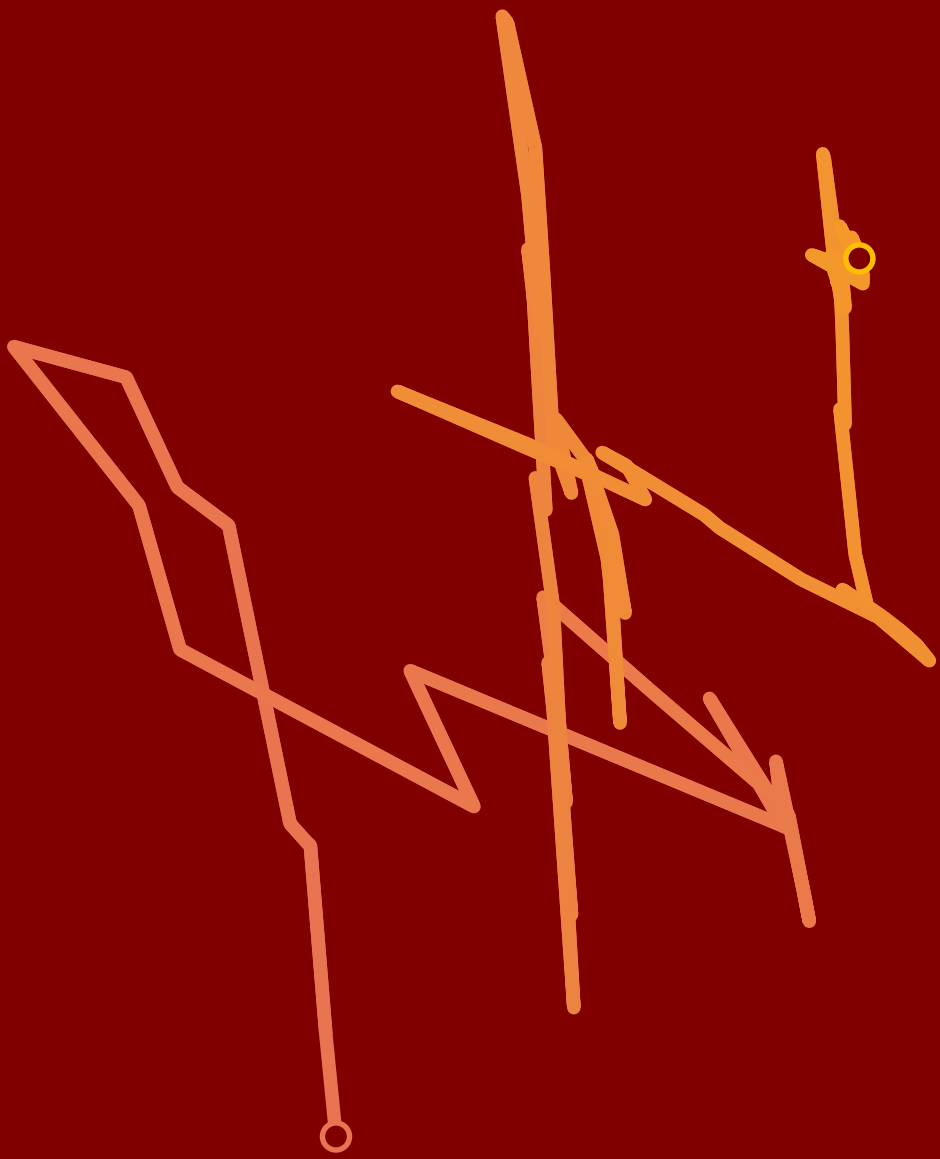

2

Clinical evaluation of bi-objective treatment planning for prostate brachytherapy

Research question 1:
Can clinically acceptable prostate brachytherapy treatment plans be obtained by optimizing the bi-objective planning model with MO-GOMEA?

This chapter is an adaptation of S.C. Maree, N.H. Luong, E.S. Kooreman, N. van Wieringen, A. Bel, K.A. Hinnen, G.H. Westerveld, B.R. Pieters, P.A.N. Bosman, T. Alderliesten. (2019) *Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy – a retrospective observer study.* Brachytherapy 18 (3), 396–403.

Abstract

Bi-objective treatment planning for high-dose-rate prostate brachytherapy is a novel treatment planning method in which two separate objectives are being used that represent target coverage and organ-at-risk sparing. In this chapter, we investigate the feasibility and plan quality of this method by means of a retrospective observer study. Current planning sessions are recorded to configure the bi-objective planning model and to assess its applicability to our clinical practice. Optimization software, MO-GOMEA, is then used to automatically generate a large set of plans with different trade-offs in the two objectives for 18 prostate cancer patients. Five plans per patient are selected from the set for comparison to the clinical plan, in terms of satisfaction of the planning criteria, and in a retrospective observer study. Three brachytherapists are asked to evaluate the plans and select a preferred one. Recordings demonstrate that the bi-objective planning model represents our clinical practice well. For 14/18 patients, automatically generated plans satisfy all planning criteria, compared to 4/18 clinical plans. In the observer study, in 53/54 cases, an automatically generated plan was preferred over the clinical plan. When asked for consensus among observers, this ratio was 17/18 patients. Observers highly appreciated the insight gained from comparing multiple plans with different trade-offs simultaneously. Automatically generated plans were considered equal or superior to the clinical plans. Additionally, presenting multiple high-quality plans provided novel insight into patient-specific trade-offs.

2.1 Introduction

Commercially-available treatment planning methods for high-dose-rate (HDR) prostate brachytherapy, such as Inverse Planning Simulated Annealing (IPSA) [4, 124] and Hybrid Inverse Plan Optimization (HIPO) [155], present a single plan to the planner. This plan can be adapted iteratively by changing dwell times, by drag-and-dropping of isodose lines [159], or by changing the parameters of the planning method. After each modification, the planner assesses the quality of the plan by a set of planning criteria that are based on the dose-volume indices (DVIs) as stated in the clinical protocol and by a visual inspection of the dose distribution. To be able to efficiently obtain high-quality plans using inverse planning methods, the underlying optimization model should closely match the planning criteria, while

computation time is still acceptable. The optimization model on which IPSA and HIPO are based – a dose-penalty model – is fast to compute, but does not always result in plans that adhere to the planning criteria, even if such plans do exist [77, 95]. Other methods [56, 82, 189] do model the treatment planning problem based on the DVIs of the planning criteria. However, similar to IPSA and HIPO, these methods combine multiple planning criteria by a weighting that the planner needs to set. Tuning the weights is a patient-dependent and non-trivial task, which makes treatment planning a difficult and time-consuming trial-and-error process [60, 145]. Obtaining acceptable plans by this procedure requires experience and is time-consuming, often taking more than 30 minutes [82, 115, 159].

To overcome the manual tuning of weights, treatment planning can be formulated as a multi-objective optimization problem, where each planning criterion is an objective [118]. Typically, there are 5 or more objectives, depending on the treatment site and clinical protocol [96]. The optimum of a multi-objective optimization model is not a single treatment plan, but a large set of plans, all with a different trade-off between the objectives. A single preferred plan then needs to be selected from this set, either manually or automatically. However, considering all planning criteria as separate objectives results in many objectives, which makes fast computing of all best trade-off plans computationally infeasible [132]. In practice, an interactive method can be used in which the planner steers optimization to get to a desirable plan [162, 185, 193]. Although more intuitive than setting weights, it can be hard to get intuition about the nature of underlying trade-offs this way. Alternatively, approaches exist to navigate the set of trade-off solutions without first computing them all, but approximations of the planning criteria (i.e., dose-penalty models rather than DVI models) and plan interpolations are then required for efficiency reasons, making it more difficult to obtain plans that are in line with the planning criteria [47, 77, 95, 158].

The novel treatment planning method that is used in the current study models the treatment planning problem as a bi-objective optimization problem, with only two objectives. One objective is based on target coverage and the other on organ-at-risk (OAR) sparing [135]. The objectives are directly based on the DVIs stated in the clinical protocol. As there are only two objectives, first computing a set of high-quality trade-off plans is computationally tractable and visualization of these plans as a trade-off curve is straightforward. This reduces treatment planning to a decision-making process of selecting the preferred plan from this trade-off curve.

In this chapter, we evaluate the use of bi-objective treatment planning for our clinical practice, and evaluate the quality of resulting plans in a retrospective observer study.

Table 2.1: HDR brachytherapy prostate protocol of a single planning-aim dose of 13 Gy. Volume (V) indices in percent of the planning-aim dose. Dose (D) indices are expressed in volume percentage or absolute volume (cm^3), and units gray (Gy).

	Volume	Use	Criteria	
Coverage criteria	Prostate	Target	$V_{100\%} > 95\%$	$D_{90\%} > 100\%$
	Vesicles	Target	$V_{80\%} > 95\%$	
Sparing criteria	Prostate	OAR	$V_{150\%} < 50\%$	$V_{200\%} < 20\%$
	Bladder	OAR	$D_{1cm3} < 86\%$	$D_{2cm3} < 74\%$
	Rectum	OAR	$D_{1cm3} < 78\%$	$D_{2cm3} < 74\%$
	Urethra	OAR	$D_{0.1cm3} < 110\%$	

2.2 Materials and methods

2.2.1 Patient and treatment characteristics

Between February 2015 and April 2017, 18 prostate cancer patients were treated in our clinic according to the protocol in Table 2.1, receiving single-dose HDR brachytherapy of 13 gray (Gy) a week after external beam radiation treatment with a dose schedule of 20 × 2.2 Gy. Median age at time of treatment was 68 (range: 58–84) years, Gleason Score was between 6 and 9 (ISUP grade group 1–5). The median urinary flow rate was 16.3 (range: 8.5–34.8) ml/s and the median prostate volume defined by Magnetic Resonance Imaging (MRI) after catheter placement was 31.9 (21.1–69.3) cm^3. A median of 16 (range: 14–20) catheters were implanted with a source step of 2.5 mm, totalling to a median of 413 (range: 250–668) dwell positions. Catheter implantation was performed using transrectal ultrasound under general or epidural anaesthesia according to a pre-plan, made in the operation theatre based on ultrasound imaging, in Oncentra Prostate (Elekta Brachytherapy, Veenendaal, the Netherlands) [169]. Visibility of the urethra was enhanced by a transurethral catheter with a bladder balloon.

After implantation, three orthogonal pelvic T2-weighted turbo spin echo MRIs (Ingenia 3T Philips Healthcare, Best, The Netherlands) with an in-plane resolution of 0.6 × 0.7 mm and 3.0 mm slice thickness with 0.3 mm gap were acquired and used for treatment planning. Imaging was taken with the patients lying on their back and legs flat, similar to the treatment position. These images were loaded into Oncentra Brachy (version 4.3–4.5, Elekta Brachytherapy, Veenendaal, the Netherlands) and used for catheter reconstruction and delineation of the regions of interest.

Initial plans of patients treated before mid-2015 were created using IPSA, initial plans after mid-2015 were created with HIPO. Both IPSA and HIPO were run

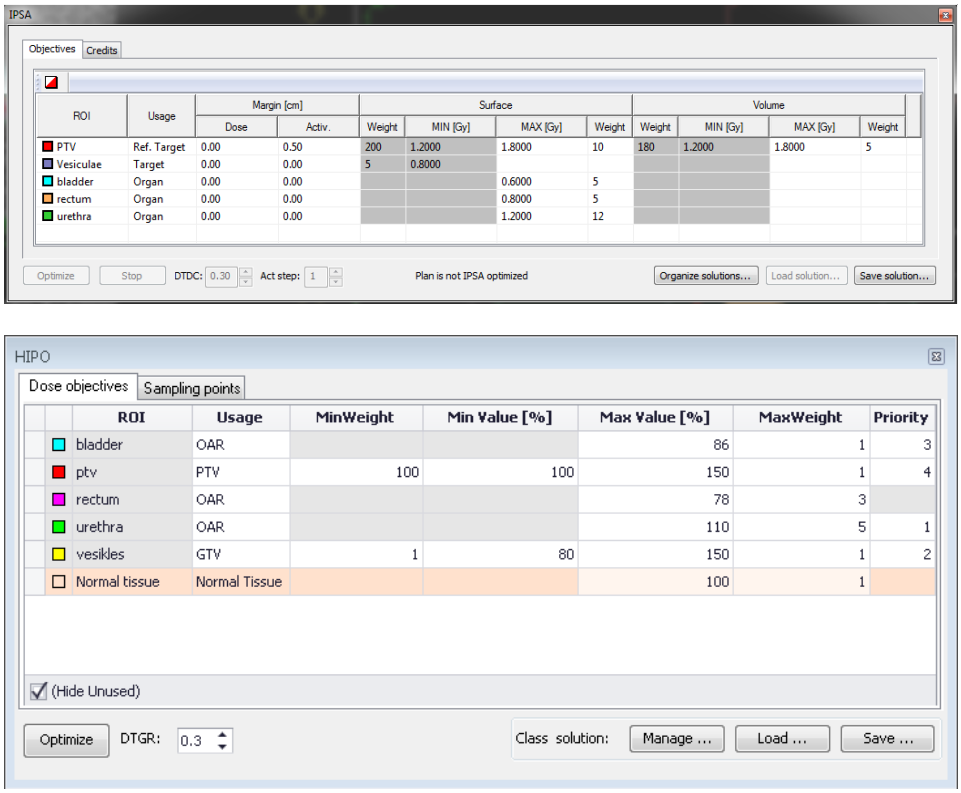

Figure 2.1: Class solutions of IPSA (top, used before mid-2015) and HIPO (bottom, used after mid-2015) for generation of an initial plan in clinical practice. After, manual (graphical) optimization was performed. These class solution may therefore be sub-optimal.

with a standard parameter set, i.e., a class solution, as provided in Figure 2.1. Plans were then manually fine-tuned using graphical optimization. Next, quality assurance checks were done by a medical physicist. Finally, the plans were assessed for clinical acceptability by a physician using the criteria in Table 2.1, and by visual inspection of the dose distribution.

To gain insight into the current planning process in our clinic and into the applicability of the bi-objective optimization model [135], clinical planning sessions were filmed, and changes in DVIs over time were recorded during the manual graphical optimization of five of the 18 patients. It was measured how many of the changes were dedicated to improving the DVIs, and how many focused on other aspects not explicitly mentioned in the clinical protocol. Based on this analysis, the bi-objective optimization model was configured [131, 146].

2.2.2 Configuring the bi-objective optimization model

The bi-objective optimization model groups the planning criteria (Table 2.1) into one coverage objective and one sparing objective. The Least Coverage Index (LCI) was constructed by combing the coverage criteria $V_{100\%}^{\text{prostate}}$ and $V_{80\%}^{\text{vesicles}}$ in a worst-case manner, as follows,

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

An LCI = 1.2% should be read as: *the worst covered target is covered 1.2% more than its aim and the other target has a higher coverage.* Thus, when maximizing the LCI, a certain level of coverage for both $V_{100\%}^{\text{prostate}}$ and $V_{80\%}^{\text{vesicles}}$ is guaranteed. The criterion $D_{90\%}^{\text{prostate}} > 100\%$ has been left out of the model, as it is automatically satisfied when $V_{100\%}^{\text{prostate}} > 90\%$. Moreover, explicitly maximizing it would lead to dose escalation, and it is currently unclear whether this is desirable [173].

The Least Sparing Index (LSI) was constructed in a similar worst-case approach from the sparing criteria, as follows,

$$\text{LSI} = 13 \text{ Gy} \times \min\left\{\begin{array}{l} 86 - D_{1\text{cm3}}^{\text{bladder}}, 74 - D_{2\text{cm3}}^{\text{bladder}}, \\ 78 - D_{1\text{cm3}}^{\text{rectum}}, 74 - D_{2\text{cm3}}^{\text{rectum}}, \\ 110 - D_{0.1\text{cm3}}^{\text{urethra}} \end{array}\right\}.$$

An LSI = 1.4 Gy should be read as: *the worst spared OAR is spared 1.4 Gy more than its criterion, and all other OARs are spared even more.* Thus, the LSI should also be maximized, and when LSI > 0 Gy holds, all sparing criteria are satisfied. Plans that satisfy all coverage and all sparing criteria are called *satisfactory* plans.

The criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ are the only sparing criteria based on volume indices. Directly adding them to the LSI would result in a comparison of indices with a different unit (percentage of volume compared to percentage of prescribed dose). Analysis of clinical plans, showed that the criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ were never violated, as shown in Table 2.3 at the end of this chapter. As the LSI uses a worst-case approach, and as these criteria were never violated, there would be no effect of leaving them out of the objectives. To guarantee that optimization generates plans that satisfy these two criteria, all plans that violate them are automatically discarded.

2.2.3 Automatic bi-objective treatment planning

For each patient, a large set of high-quality plans was automatically generated by optimizing plans under the bi-objective model. For this, patient DICOM RT-Struct and RT-Dose files were exported from Oncentra Brachy and processed by

our in-house developed TG-43 [180, 201] dose engine. As starting point for the optimization, all dwell positions were activated within a 5-mm margin of the targets (i.e., prostate and vesicles), excluding positions within a 1-mm margin of the urethra. Next, the dwell times associated with these dwell positions were initialized with a randomly chosen value between 0 and 1 second. The aim of treatment planning is then to optimize these dwell times.

The bi-objective model is non-convex, non-linear, and non-smooth, and to optimize plans according to it, a state-of-the-art multi-objective evolutionary algorithm, the Multi-Objective Gene-pool Optimal Mixing Evolutionary Algorithm (MO-GOMEA) [38, 131], was used. In MO-GOMEA the fact that the dose distribution can be quickly updated when only few dwell times change, is exploited. A run of MO-GOMEA was limited to one hour on a low-end processor. In a 1-hour run, MO-GOMEA produced a large set of 100–1000 treatment plans, all with different LCI/LSI trade-offs [131, 146].

As MO-GOMEA is a stochastic algorithm, it was run 30 times to assess its variance in final results, which was shown to be small [131]. From these 30 runs, a single set of plans was constructed by only retaining plans that exhibit the best trade-offs in LCI and LSI. Further research [34, 133] showed that using a Graphics Processing Unit (GPU) that is available in modern planning machines, the same results can be obtained within a matter of minutes.

For each patient, the resulting set of plans consisted of hundreds of high-quality plans. Presenting all these plans in an observer study was infeasible as software to quickly navigate through and compare that many plans is not yet available. We therefore manually selected five plans, as shown in Figure 2.2, for further analysis according to the following strategy: left and right of LCI = 0, above and below LSI = 0, and the plan in the middle. These plans are labelled from small to large LCI values as: high sparing, sparing, coverage/sparing, coverage, and high coverage. These five selected MO-GOMEA plans per patient were compared to clinical plans in terms of their (LCI, LSI)-values. Differences were tested for significance using a Wilcoxon matched-pair signed-rank test ($\alpha = 0.05$).

2.2.4 Observer study setup

A retrospective observer study was conducted with three physicians responsible in our clinic for HDR prostate brachytherapy. Each observer was individually presented with six plans per patient: the plan that was clinically used to treat the patient, and the five selected MO-GOMEA plans, without identifying which plan was which. The DVIs of the six plans were presented in a single overview,

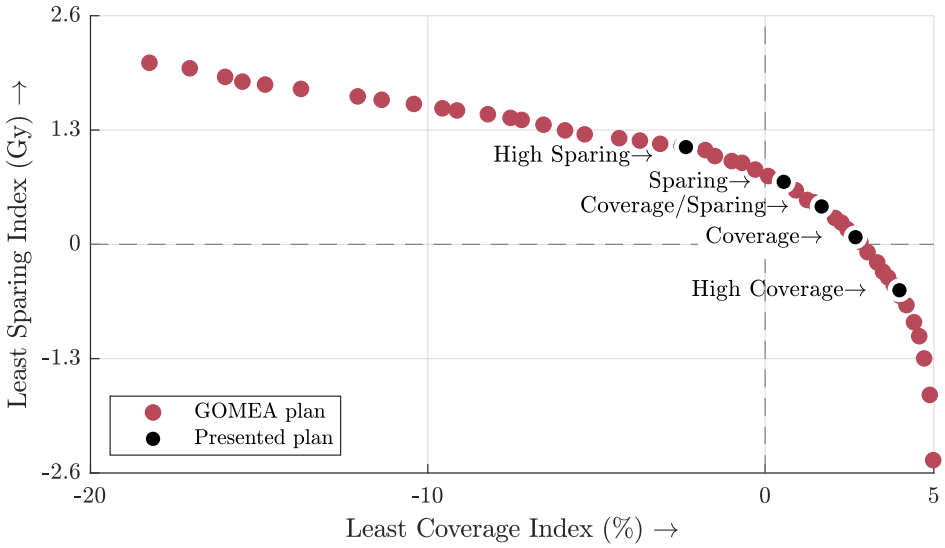

Figure 2.2: Example of a trade-off curve for one patient. Each dot represents a treatment plan obtained by MO-GOMEA. The axes correspond to the objectives of the bi-objective model. Both objectives should be maximized. Plans in the upper right corner, where the values for the Least Coverage Index and the Least Sparing Index are larger than zero, are satisfactory, i.e., satisfy all planning criteria. Large black dots illustrate the plans presented in the observer study.

similar to Oncentra Brachy. Then, the observer could inspect the corresponding dose distributions in Oncentra Brachy, one at a time.

Per patient, each observer was asked to answers the following questions:

- What plan do you prefer to treat the patient with?
- Which plans do you consider clinically acceptable?
- What plan would you dismiss?

After answering these questions, additional patient information was provided and the observer was asked if this changed his or her opinion on the preferred plan. The additional information comprised patient age and Gleason score for all patients. Urinary flow rate, available biopsy information, tumor location information, and potential tumor invasion of the seminal vesicles was available for the twelve most recent patients. Furthermore, for the eight most recent patients, an additional diffusion-weighted MRI with tumor location information was added as it had been available clinically after a recent adaptation of the clinical workflow.

Additional patient information was provided to mimic clinical practice as much as possible. IPSA and HIPO do not take this additional patient information into

account when initialized by a standard parameter set, but this information is known to the planner during manual optimization, and this information can then thus be incorporated into clinical plans. Similarly, MO-GOMEA does not take this additional information into account during optimization, but a planner might while selecting the desired MO-GOMEA plan. To investigate how this information changed decision making, we chose for this setup, where the questions are repeated after the first evaluation. It was recorded which aspects the observers assessed that were not mentioned in the clinical protocol, and how they approached decision-making.

This setup of presenting multiple MO-GOMEA plans was chosen because it provides additional insight in the patient-specific trade-off. This can be used during decision-making. The clinical plan was added to compare quality of MO-GOMEA plans and to assess clinical acceptance. This setup does however introduce a potential bias towards selecting a MO-GOMEA plan. We investigated statistical significance of the preferred plan by comparing against the null-hypothesis that all plans are equally likely to be selected, at $\alpha = 0.05$. Because the number of observers is too low to perform an observer variability study, p-values are reported per observer.

Consensus meeting

One week after the observer study, a consensus meeting was held where patients were discussed for which each observer selected a different preferred plan. These three plans were again blindly presented and observers were asked whether they could agree on a single preferred plan.

2.3 Results

2.3.1 Analysis of clinical planning sessions

The five filmed planning sessions lasted for a median of 33 minutes (range: 9–48), with an average of four drag-and-drop steps per minute, totaling to 525 modifications. These modifications were categorized based on the intended aim of corresponding change, as indicated verbally by the planner. The categories we distinguish are: Reducing dose to OARs; Increasing dose to the targets; 'Deactivation' is the (de)activation of additional dwell positions, that are either too close to an OAR, or more than 5 mm away from a target; 'Homogenizing' is spreading out the dose over multiple adjacent dwell positions, with the aim of making the dose distribution more homogeneous; Reducing hotspots is a particular

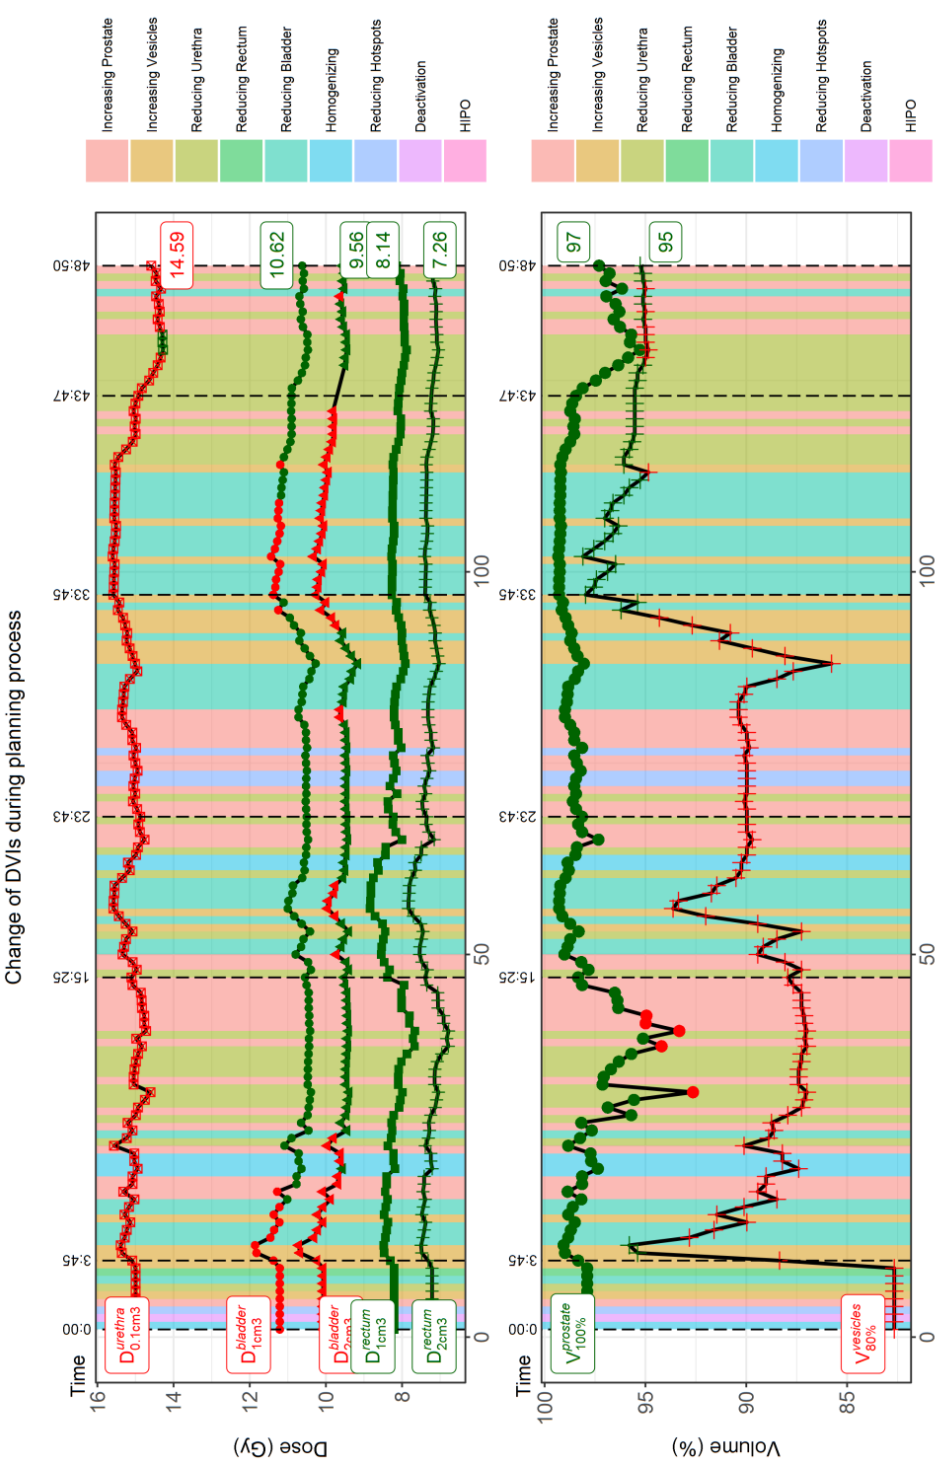

Figure 2.3: Transcript of the manual planning session for Patient 13. The number of changes are in chronological order, and the dashed vertical lines show some time stamps. The top figure shows the dose indices and the bottom figure shows the volume indices. The markers show whether the index satisfies the corresponding planning criterion (green) or not (red). The final value of each index is shown on the right. The background is colored based intended aim of the corresponding change. Transcripts for Patients 14–17 are shown at the end of this chapter in Figure 2.7.

form of dose homogenizing, aimed at reducing size and number of contiguous regions of more than 200% of the planning-aim dose; 'HIPO' indicates (initial) plans that were generated with HIPO. The remainder of the changes is performed using manual graphical optimization. Figure 2.3 visualizes the transcript of the manual planning session for Patient 13. The transcripts of the other four patients can be found at the end of this chapter in Figure 2.7.

Modifications were performed on an iterative basis, focusing on two conflicting criteria: 66% of the modifications on prostate vs. urethra, 20% on vesicles vs. bladder, and 1% on prostate vs. rectum. 50% of the changes were made in order to improve the least spared/covered volume. Two aspects not included in the clinical protocol, that were assessed, were improving dose homogeneity and reducing hotspots, i.e., volumes with dose higher than 200% of the prescribed dose, both in number and in size. For most of the manual plan-optimization time, both LCI and LSI had a negative value. Over time, they were improved iteratively, which corresponds to a zig-zag pattern in the bi-objective representation, as can be seen in Figure 2.4.

2.3.2 Automatic bi-objective treatment planning compared to clinical plans

Figure 2.5 shows the trade-off curves obtained with MO-GOMEA, together with the clinical plan and the five plans that were selected from the front for comparison. An overview of the DVIs associated with the clinical plan and the five selected plans for each patient is given in Table 2.3 at the end of this chapter.

For all patients, the MO-GOMEA plans had simultaneously a better LCI and LSI than the clinical plan. For four patients (10, 11, 16, 17), the clinical plan satisfied the clinical protocol, while MO-GOMEA plans satisfied the clinical protocol for 14/18 of the patients. For four patients (4, 7, 14, 18), neither the plans optimized by our method nor the clinical plan was satisfactory, caused by an unfavourable implant geometry. For some patients, the LCI value of the clinical plan is small because of a low vesicle coverage. For all clinical plans, $V_{100\%}^{\text{prostate}} \geq 93.3\%$ holds.

From the set of five MO-GOMEA plans per patient, MO-GOMEA plans with a similar or better LSI as the clinical plan had an LCI that is 3.5% larger than the LCI of the clinical plan, averaged over all patients (range: −0.6–14.7%, SD 4.3%, $p < 0.05$). MO-GOMEA plans with a similar or better LCI as the clinical plan had an LSI that is 0.85 Gy larger than the LSI associated with the clinical plan, averaged over all patients (range: 0.1–2.0 Gy, SD 0.6 Gy, $p < 0.001$).

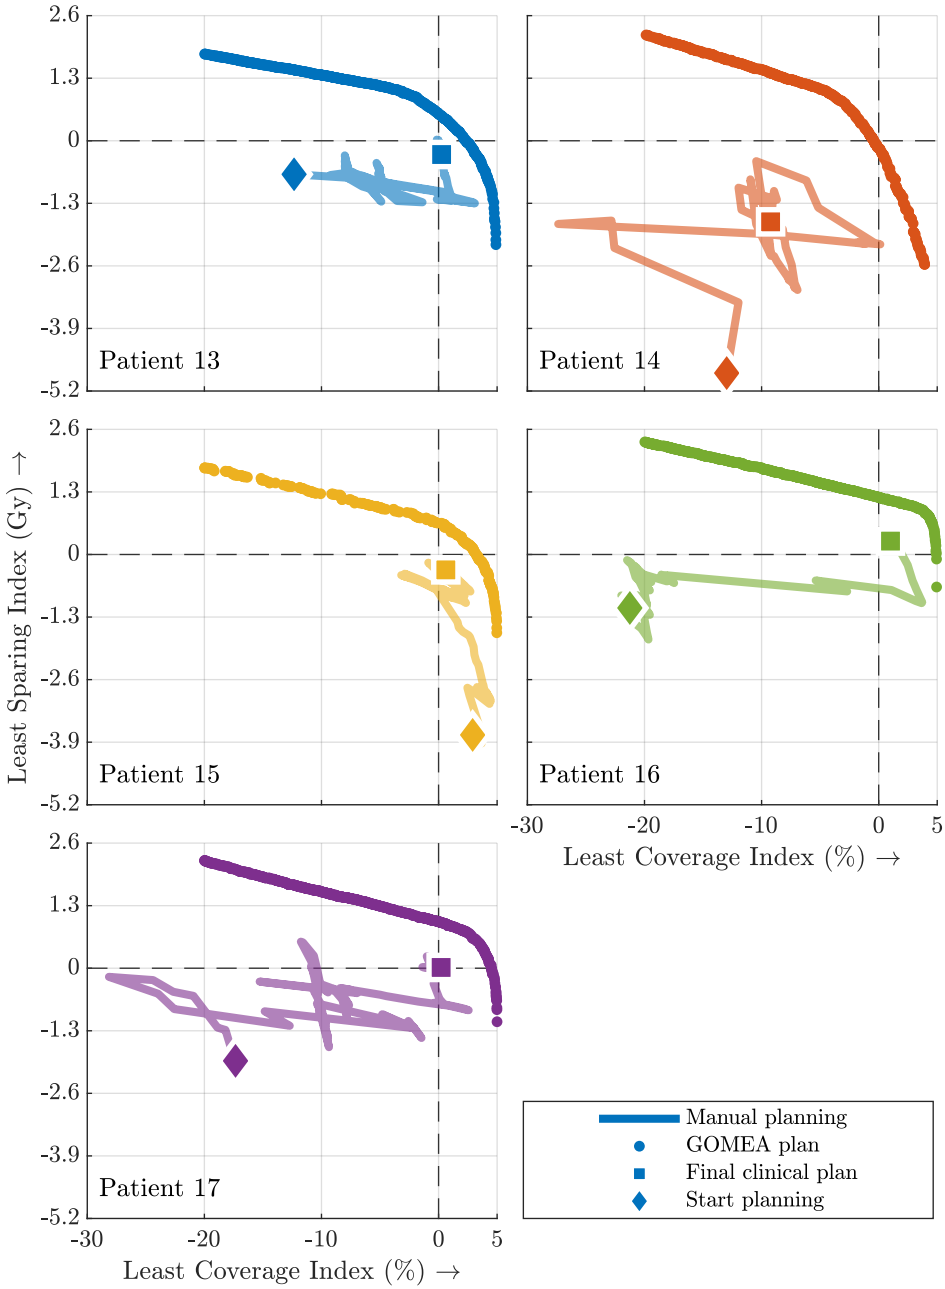

Figure 2.4: Manual planning process using graphical optimization visualized in the bi-objective representation for the five recordings. The trade-off curves obtained with MO-GOMEA are shown for comparison. A transcript of the planning process for Patient 13 is given in Figure 2.3.

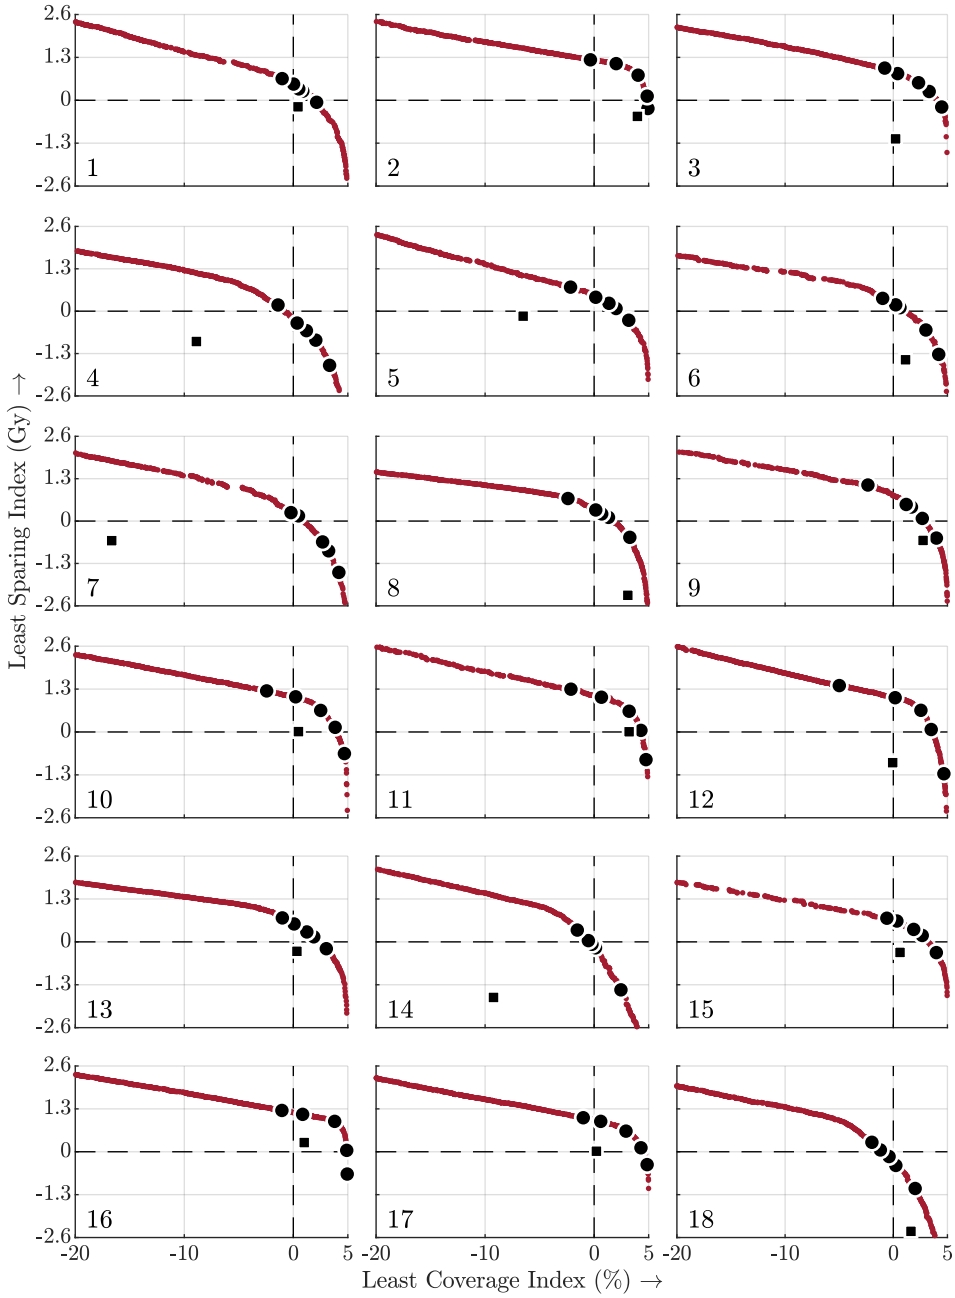

Figure 2.5: Trade-off curves for all 18 patients in small red dots. Black squares represent the clinical plans. Black circles are the MO-GOMEA plans presented in the observer study, representing from left to right on the trade-off curve the selected high-sparing, sparing, coverage/sparing, coverage, and high-coverage plan.

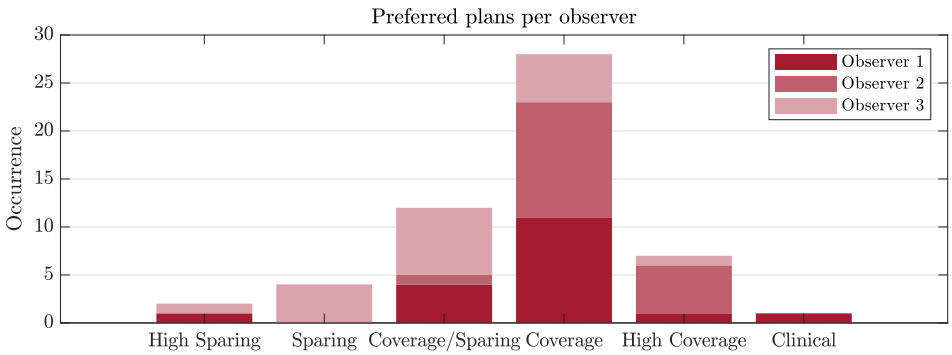

Figure 2.6: Distribution of the preferred plans after the clinical information was presented, per observer in the retrospective observer study with in total 54 cases.

2.3.3 Observer study

Table 2.2 shows the results of the observer study, before and after presenting additional clinical patient data, totaling to three observers $\times$ 18 patients = 54 cases. In all cases, the preferred plan was considered clinically acceptable. Furthermore, in all cases, one or more MO-GOMEA plans was considered clinically acceptable. In 53/54 (98%) cases, a MO-GOMEA plan was preferred. One observer selected a clinical plan once ($p = 0.173$), the two other observers never ($p = 0.038$). The coverage plan was preferred most often. Observers had different distributions of preferred plans, as shown in Figure 2.6. One observer checked all six dose distributions visually, while the other two first dismissed plans based on the DVIs, and only visually inspected the remaining one or two plans. The high-sparing plan was dismissed most often, for insufficient coverage. For five patients (28%) one or more observers dismissed the clinical plan.

The preferred plan was changed eight times after presenting clinical information, of which five times by one observer. Half the changes were to increase coverage due to a high Gleason score, the other half to decrease coverage due to a low Gleason score or bad urinary flow. Five of these eight changes were made for patients who had a diffusion-weighted MRI.

For most of the plans, prostate coverage versus urethra sparing was the dominating trade-off. In that case, observers generally looked for the plan with maximum prostate coverage (visually inspected, or based on the $V_{100\%}^{\text{prostate}}$) while satisfying the urethra sparing criterion. Although for some patients, prostate coverage of this plan was deemed insufficient, and a plan was chosen that violated the urethra sparing criterion. In the visual inspection of the dose distribution, observers focused on locations where the target was not covered. The whole

Table 2.2: Per patient each of the three observers indicated which of the plans should be dismissed, denoted by "–". Preferred plans that were chosen *before* clinical information was presented "∘", and that were changed after, are denoted by "×". Preferred plans that were chosen before clinical information was presented, and not changed after this information was available are denoted by "+". For patients 5, 8, and 11, the preferred plans after the consensus meeting are denoted by "*".

Patient	High Sparing	Sparing	Coverage / Sparing	Coverage	High Coverage	Clinical
1	– – –			++	+	
2	– – –		+	++		
3	– –		+	++	–	
4	– – –	∘ ∘		+×	×	
5	– – –	+		+*	+	
6	– –		++	+	–	
7	–		+++			– –
8	– – –	×		∘+	+*	
9	– – –		+	++		
10	– – –	×		++∘		
11	– – –		+	+		+*
12	– – –	+		++		
13	– –			++	+	–
14	+×			+∘	–	– –
15	–			+++		– –
16	– –		++	∘	×–	
17	– – –		∘	++×		
18			+	×	+∘	– – –

prostate was considered to be target volume, and an under dose can only be acceptable in parts where no tumor is expected. Observers also focused on hotspots and the activation of dwell positions in close proximity of the OARs. In a few cases, observers mentioned that they would like to try to manually improve MO-GOMEA plans by disabling dwell positions or by spreading out the dose more evenly over multiple dwell positions.

Overall, observers highly appreciated to see multiple plans simultaneously to get an impression of the achievable trade-offs. Decision making was said to be harder when none of the presented plans was satisfactory, i.e., satisfied all planning criteria.

Consensus Meeting

Three patients (5, 8, 11) were discussed in the consensus meeting (Table 2.2). Observers easily came to an agreement. For patient 11, the clinical plan was preferred over the other plans by one observer, and later in the consensus meeting

by all observers. For this patient, all three presented plans were satisfactory, and observers agreed all plans are clinically acceptable. The coverage/sparing plan was dismissed based on the DVIs. Finally, the coverage plan was dismissed as well because of a high value of the additionally computed $D^{\text{rectum}}_{0.01\text{cm3}} = 186\%$, compared to $D^{\text{rectum}}_{0.01\text{cm3}} = 85\%$ of the clinical plan.

2.4 Discussion

The novel bi-objective treatment planning method that was clinically evaluated for the first time in this chapter automatically generates a large set of high-quality treatment plans, from which the planner can then select the desired plan. This makes treatment planning an insightful decision-making process instead of a trial-and-error optimization process.

Recordings of the clinical planning sessions showed that the bi-objective model corresponded well to our clinical practice. Not all planning criteria in our clinical protocol were included in the bi-objective model. First, $D^{\text{prostate}}_{90\%} > 100\%$ was omitted, which is a criterion that should be satisfied, but should not be maximized further. This distinction is important to notice and it highlights the importance of considering each planning criterion carefully when configuring the model. Second, $V^{\text{prostate}}_{150\%} < 50\%$ and $V^{\text{prostate}}_{200\%} < 20\%$ could be excluded in our case, as these criteria have loose aims that were always achieved. Alternatively, if these criteria would need to be included anyway by addition to the LSI, because a different clinical protocol is used, or when extending this method to a different treatment site, they could be rewritten as $D^{\text{prostate}}_{50\%} < 150\%$ and $D^{\text{prostate}}_{20\%} < 200\%$ to avoid comparing indices with different units (e.g., percentage of volume with percentage of dose).

MO-GOMEA is by design able to handle more than two objectives. Additional criteria based on indices with different unit could therefore be added as third of fourth objective. Nevertheless, so far, for the brachytherapy treatment planning application we limited ourselves to the use of two objectives. The main downside of additional objectives is that it would make it harder to visualize the trade-off curve, which potentially complicates the decision making.

In the recordings, and during the visual inspection of the dose distributions in the observer study, observers focused not only on the planning criteria, but on other aspects as well: hotspots, the activation of dwell positions near the OARs, and the location of areas where the target was not covered. Ensuring target coverage locally can be incorporated in optimization by e.g., indicating which sub-volumes (e.g., quadrants) of the prostate should be fully covered because of tumor presence. The other aspects are not easily quantified. Earlier attempts have shown to potentially

deteriorate plan quality, and different attempts were shown to be inconsistent [13].

For some plans, observers mentioned they would like to try to manually improve the plan, mainly to satisfy criteria that are not explicitly in the clinical protocol. By not allowing further improvements, it was shown that MO-GOMEA plans are already clinically acceptable in their current form, without further tuning needed. Nevertheless, the potential to further improve the automatically generated plans by manual graphical optimization is of interest in future studies.

Several other planning approaches [77, 159] also model the treatment planning problem based on the (approximated) DVIs, but in a single-objective manner, where the trade-off is represented by weighting criteria. Figure 2.5 however shows that the trade-off is patient dependent and Figure 2.6 suggests that the preferred plan was in principle observer dependent, although observers could come to an agreement in the consensus meeting. These confirm the value and validity of the DVI-based multi-objective planning approach discussed in this chapter.

A limitation of the present study is that not for all patients the same set of clinical information was available to be presented in the observer study. The two-step approach in the observer study gives insight in how this affected decision making. The majority of changes of the preferred plan was made for patients who had a diffusion-weighted MRI available, which suggests that it is of additional value for decision making. It also suggests that more changes would have been made when this information was available for all patients. However, with the small total number of changes observers made, it is unknown if having the same information for all patients would have led to consistent changes in plan preference.

Future work will be to evaluate the applicability of bi-objective planning model in other institutes. Also, development of techniques that allow fast and insightful navigation of MO-GOMEA plans, instead of preselecting five plans, is of interest for the clinical implementation of bi-objective treatment planning.

2.5 Conclusion

To conclude, we retrospectively evaluated a novel bi-objective treatment planning method for HDR prostate brachytherapy for use in our clinic. The analysis of current clinical planning sessions showed that the bi-objective model was easily configured and represented our clinical practice well. The observer study demonstrated that the bi-objective method automatically generates plans of high-quality. For all patients and by all observers, resulting plans were preferred over the clinical plan in 98% of the cases. The ability to compare multiple high-quality plans was considered insightful and highly appreciated by the observers.

Table 2.3: (Part 1): DVIs of the presented plans in the observer study, see Figure 2.2 for a visualization. Presented plans are the Clinical (Cl), High Coverage (HC), Coverage (C), Coverage/Sparing (C/S), Sparing (S) and High Sparing (HS) plans. LCI and V-indices in percentage of volume, LSI and D-indices in percentage of the planning-aim dose (13 Gy).

Patient Plan	LCI	LSI	$V^{prostate}_{100\%}$	$V^{prostate}_{150\%}$	$V^{prostate}_{200\%}$	$D^{prostate}_{90\%}$	$V^{vesicles}_{80\%}$	$D^{bladder}_{1cm3}$	$D^{bladder}_{2cm3}$	D^{rectum}_{1cm3}	D^{rectum}_{2cm3}	$D^{urethra}_{0.1cm3}$
aim	>0.0	>0.0	>95	<50	<20	>100	>95	<86	<74	<78	<74	<110
1 Cl	0.0	-2.0	95	34	13	106	96	86	76	71	64	111
HC	1.6	-2.6	97	27	11	108	97	87	77	71	65	110
C	0.2	0.2	95	26	11	106	96	84	74	70	64	107
C/S	-0.2	0.7	95	26	11	105	95	84	73	70	64	106
S	-0.6	1.9	95	25	11	105	94	83	72	71	65	106
HS	-1.4	3.0	94	24	10	104	94	81	71	70	65	104
2 Cl	4.0	-4.0	99	26	10	112	100	80	70	51	44	114
HC	4.0	-2.0	99	24	8	110	101	88	75	64	56	111
C	3.9	0.4	99	21	8	107	101	86	72	56	48	109
C/S	3.0	5.6	98	19	7	104	101	80	67	59	50	104
S	0.8	8.4	96	18	7	103	101	77	64	54	46	102
HS	-1.8	9.3	93	16	6	102	101	76	64	55	47	101
3 Cl	0.0	-9.0	98	28	10	109	95	79	83	79	69	110
HC	3.4	-3.0	98	32	10	111	100	86	77	81	71	112
C	2.0	0.5	97	25	9	107	99	82	73	77	68	108
C/S	0.6	2.4	96	21	8	105	98	80	72	75	66	106
S	-1.4	4.8	94	20	8	103	96	78	69	72	64	104
HS	-2.1	5.6	93	19	8	102	95	77	68	71	62	103
4 Cl	-9.0	-7.0	99	37	18	109	86	89	81	68	60	115
HC	1.8	-14.5	98	37	16	113	97	97	88	89	77	123
C	-0.2	-8.0	96	33	14	109	95	90	82	78	68	116
C/S	-1.0	-5.5	95	31	13	107	94	87	80	79	69	114
S	-2.2	-3.7	94	31	12	105	93	86	78	75	65	113
HS	-5.0	0.6	93	27	11	103	90	81	73	74	64	108
5 Cl	-7.0	-1.0	96	36	16	106	88	86	75	71	61	111
HC	2.0	-1.3	97	30	11	107	99	87	75	77	68	110
C	0.8	1.0	96	28	11	105	98	85	73	76	67	108
C/S	0.2	2.2	95	27	10	104	98	83	72	76	67	107
S	-0.7	3.4	94	28	10	104	96	81	71	74	66	105
HS	-3.3	5.6	92	29	10	102	94	79	68	71	63	102
6 Cl	1.0	-11.0	98	35	14	109	96	80	68	80	71	121
HC	3.1	-9.7	98	34	13	113	99	89	78	88	78	120
C	2.0	-4.1	97	30	11	109	99	84	73	82	72	114
C/S	-0.6	1.4	94	27	10	104	96	80	69	77	68	109
S	-0.8	1.9	94	25	9	104	97	80	68	76	67	108
HS	-2.0	3.4	93	24	9	102	96	79	68	75	66	106
7 Cl	-17.0	-5.0	97	31	15	108	78	84	77	65	57	115
HC	2.1	-12.8	97	37	15	113	98	97	87	81	70	122
C	1.3	-6.9	96	31	13	109	96	90	81	74	65	116
C/S	0.5	-4.8	96	31	12	108	96	87	79	72	63	115
S	-2.0	1.3	94	27	12	104	93	79	73	72	64	109
HS	-2.9	1.5	93	26	11	105	92	78	72	72	63	107
8 Cl	3.0	-18.0	98	29	11	106	99	76	68	81	74	128
HC	2.3	-6.7	97	21	7	107	99	91	81	83	75	114
C	0.3	-1.2	95	18	6	104	97	85	75	78	71	109
C/S	-0.2	0.2	95	19	7	104	96	83	74	77	70	109
S	-0.8	1.7	94	17	6	104	96	79	72	76	69	108
HS	-3.5	3.8	91	16	6	101	94	80	70	74	67	105
9 Cl	3.0	-5.0	98	28	12	108	98	81	74	53	48	115
HC	2.3	-3.6	97	29	11	109	99	83	78	68	61	113
C	1.3	0.2	96	26	11	106	99	79	74	64	58	109
C/S	0.3	3.0	95	23	9	105	98	75	71	62	55	106
S	-0.3	4.0	95	22	9	104	97	75	70	60	54	106
HS	-3.6	8.6	91	21	9	101	93	70	65	60	54	101

Table 2.3: (Part 2): DVIs of the presented plans in the observer study, see Figure 2.2 for a visualization. Presented plans are the Clinical (Cl), High Coverage (HC), Coverage (C), Coverage/Sparing (C/S), Sparing (S) and High Sparing (HS) plans. LCI and V-indices in percentage of volume, LSI and D-indices in percentage of the planning-aim dose (13 Gy).

Patient	Plan	LCI	LSI	$V_{100\%}^{\text{prostate}}$	$V_{150\%}^{\text{prostate}}$	$V_{200\%}^{\text{prostate}}$	$D_{90\%}^{\text{prostate}}$	$V_{80\%}^{\text{vesicles}}$	$D_{1\text{cm}3}^{\text{bladder}}$	$D_{2\text{cm}3}^{\text{bladder}}$	$D_{1\text{cm}3}^{\text{rectum}}$	$D_{2\text{cm}3}^{\text{rectum}}$	$D_{0.1\text{cm}3}^{\text{urethra}}$
	aim	>0.0	>0.0	>95	<50	<20	>100	>95	<86	<74	<78	<74	<110
10	Cl	0.0	0.0	97	28	10	106	95	77	66	67	56	110
	HC	2.9	-5.7	98	24	8	106	98	90	80	74	63	114
	C	2.3	0.5	97	20	7	105	97	83	74	69	58	109
	C/S	0.1	4.7	96	19	6	104	95	80	69	69	58	105
	S	-1.3	8.1	94	16	6	102	94	75	65	68	57	102
	HS	-4.8	9.2	91	16	5	101	90	75	65	65	55	101
11	Cl	3.0	0.0	98	35	15	109	100	83	73	64	57	110
	HC	3.9	-7.1	100	35	16	109	99	93	79	67	59	116
	C	3.4	0.5	99	32	14	108	98	85	72	71	61	109
	C/S	2.9	4.4	98	29	13	105	98	82	69	67	59	105
	S	0.0	7.3	95	27	11	103	95	79	65	66	59	101
	HS	-3.2	10.0	93	25	11	101	92	76	64	63	55	100
12	Cl	0.0	-7.0	99	29	11	107	95	76	67	63	57	117
	HC	3.9	-9.7	99	28	10	109	100	91	80	72	65	120
	C	2.8	0.5	98	24	9	105	100	82	72	72	64	110
	C/S	1.7	4.8	97	25	11	104	98	79	69	72	65	105
	S	-0.5	7.5	95	26	10	103	95	77	67	70	62	102
	HS	-5.4	10.7	90	24	9	100	90	73	63	67	60	99
13	Cl	0.0	-2.0	97	20	9	105	95	82	74	63	56	112
	HC	1.5	-2.6	97	20	7	109	99	85	77	71	64	112
	C	0.4	0.4	95	17	6	106	98	81	74	68	61	109
	C/S	-0.3	1.5	95	18	6	105	97	80	72	68	61	108
	S	-1.3	3.4	94	16	6	104	96	78	71	68	61	106
	HS	-2.6	4.5	92	16	6	102	95	78	69	67	60	104
14	Cl	-9.0	-13.0	93	30	12	107	86	66	59	70	61	123
	HC	1.3	-11.1	96	34	14	112	99	94	85	78	72	120
	C	-0.9	-1.6	94	28	12	107	96	83	75	76	68	112
	C/S	-1.2	-0.4	94	26	11	106	96	81	74	78	70	110
	S	-1.6	0.1	93	27	11	106	98	80	73	74	67	110
	HS	-2.4	2.6	93	24	9	104	99	77	70	74	66	107
15	Cl	1.0	-2.0	96	29	13	105	100	85	76	76	66	112
	HC	2.0	-4.4	97	28	12	108	99	87	78	77	67	112
	C	0.7	-0.6	96	28	10	106	98	83	75	76	67	109
	C/S	0.0	0.7	95	25	9	104	97	82	73	75	66	107
	S	-1.1	3.1	94	24	9	103	96	79	71	73	64	105
	HS	-1.9	3.7	93	24	9	102	95	79	70	73	63	104
16	Cl	1.0	2.0	96	28	9	105	98	76	66	54	47	108
	HC	3.7	-6.0	99	29	10	112	99	92	79	65	58	114
	C	3.7	0.4	99	20	6	108	99	85	74	59	52	108
	C/S	2.6	6.9	98	20	7	104	98	77	67	63	55	103
	S	-0.3	8.3	95	19	6	103	95	76	66	62	54	101
	HS	-2.1	9.5	93	19	6	102	94	75	64	62	54	100
17	Cl	0.0	0.0	95	25	9	105	96	72	65	68	61	110
	HC	2.4	-3.1	97	25	8	109	100	85	76	78	70	113
	C	1.9	0.9	97	24	7	108	100	80	73	76	69	109
	C/S	0.5	4.2	95	19	6	105	99	76	68	74	67	105
	S	-1.7	6.6	93	19	5	102	97	74	66	71	64	103
	HS	-3.2	7.7	92	19	6	102	95	72	65	70	63	102
18	Cl	2.0	-19.0	97	39	15	108	99	101	93	70	63	128
	HC	2.0	-8.1	97	32	11	110	97	88	82	80	72	118
	C	-0.4	-2.8	95	30	11	108	95	83	77	77	69	113
	C/S	-1.4	-0.8	95	28	10	107	94	80	74	75	67	111
	S	-1.6	0.8	94	28	10	105	93	78	72	72	63	109
	HS	-4.1	2.7	93	26	10	104	91	77	71	73	65	107

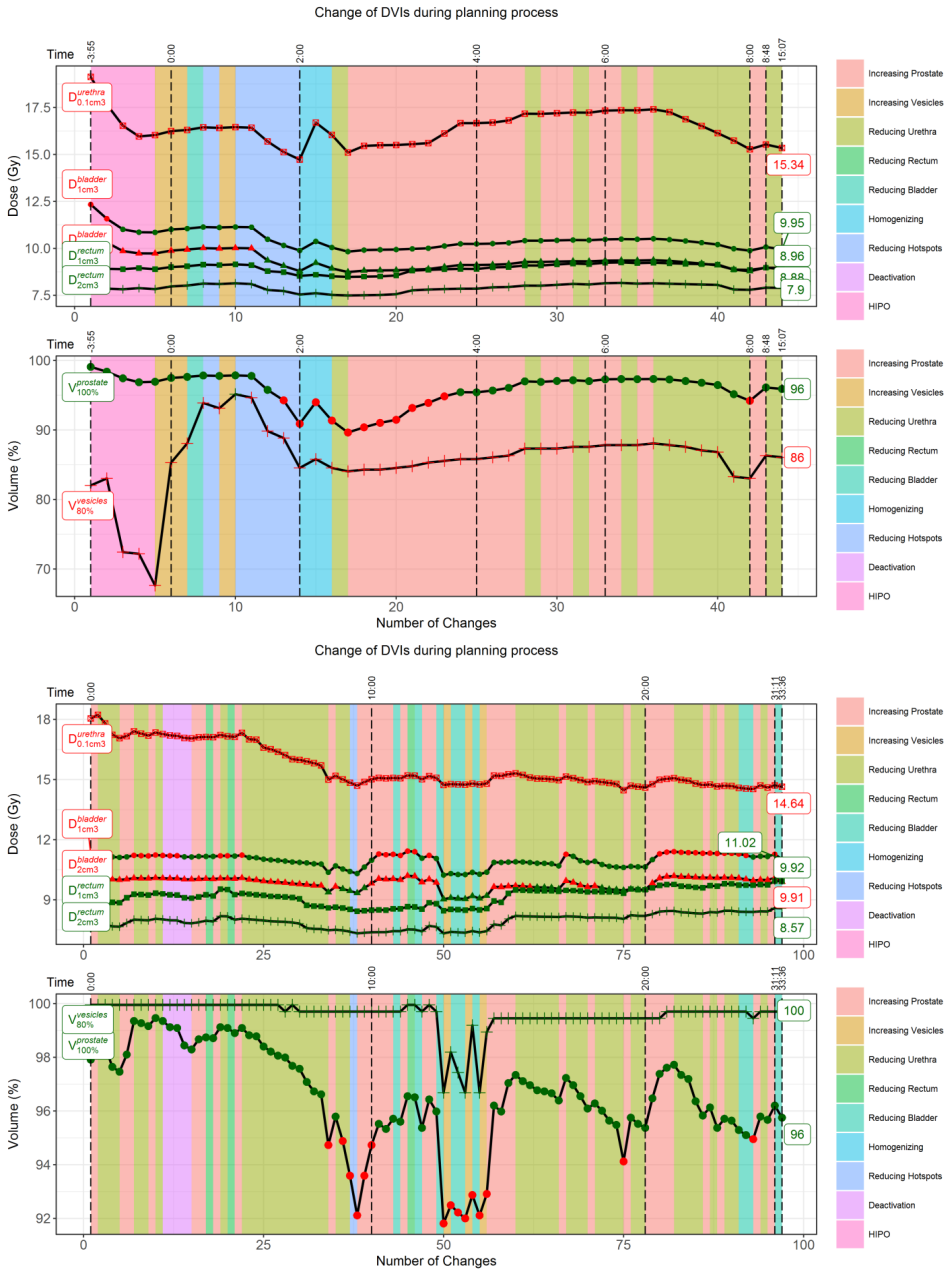

Figure 2.7: (Part 1): Transcript of the manual planning session for Patient 14 and 15. See Figure 2.3 for details. Note the change in both horizontal and vertical plot ranges.

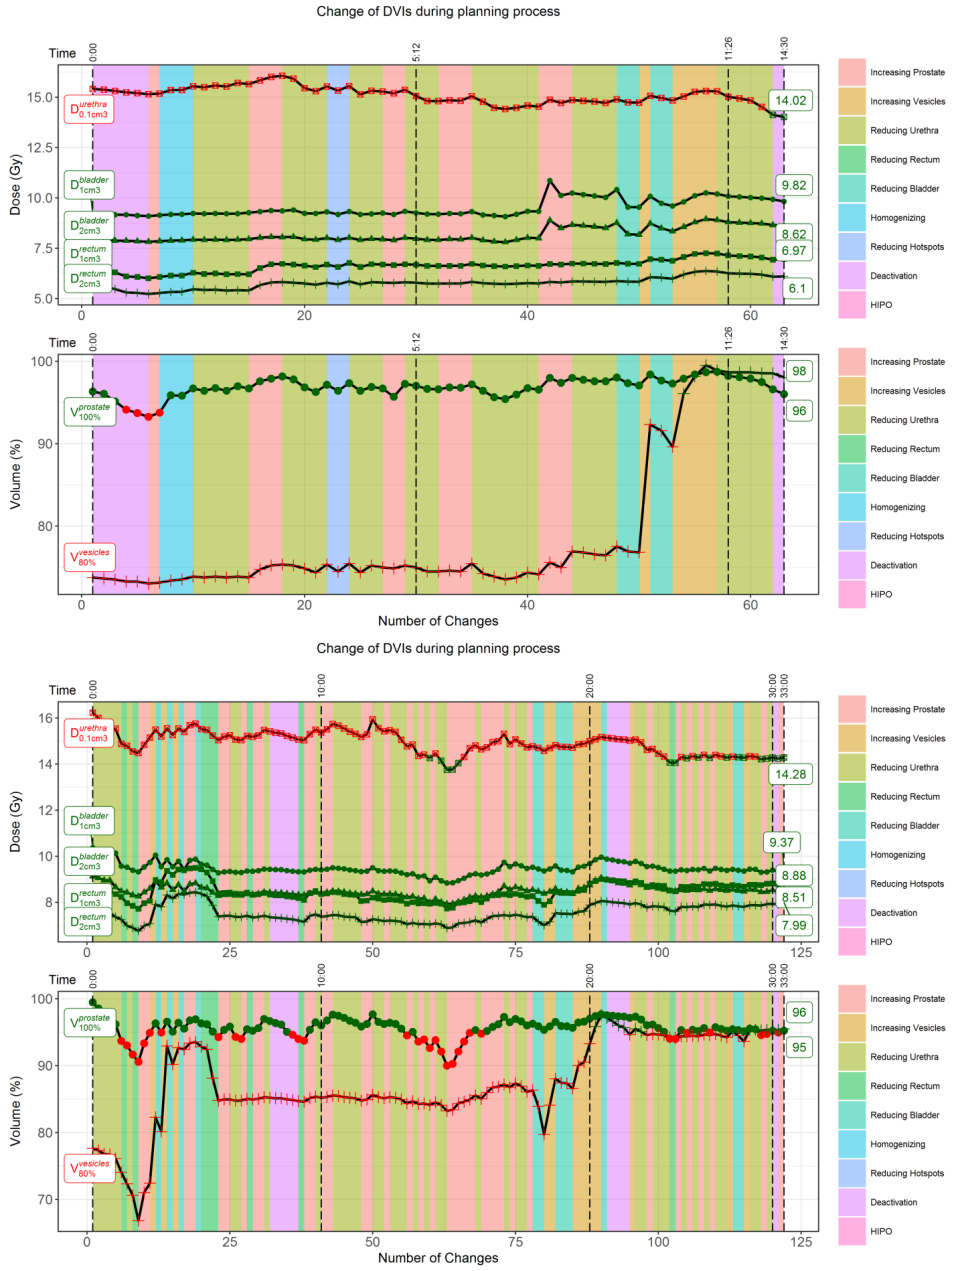

Figure 2.7: (Part 2): Transcript of the manual planning session for Patient 16 and 17. See Figure 2.3 for details. Note the change in both horizontal and vertical plot ranges.

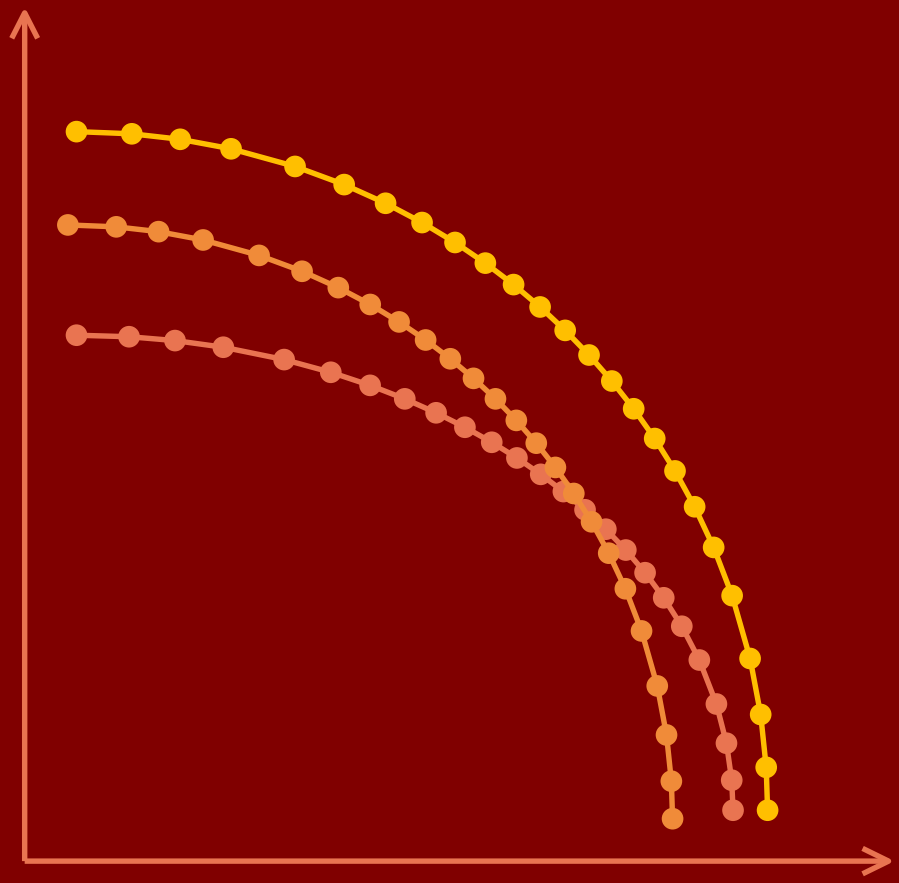

3

Automatic bi-objective parameter tuning of inverse planning methods for prostate brachytherapy

Research question 2:
Can time-consuming manual treatment planning for prostate brachytherapy be overcome by automatic tuning of the penalty weights of clinically-available inverse planning methods?

This chapter is an adaptation of S.C. Maree, P.A.N. Bosman, N. van Wieringen, Y. Niatsetski, B.R. Pieters, A. Bel, T. Alderliesten. (2020) *Automatic bi-objective parameter tuning for inverse planning of high-dose-rate prostate brachytherapy.* Physics in Medicine & Biology 65 (7), 075009.

Abstract

We aim to overcome the difficult and time-consuming manual parameter tuning that is currently required to obtain desirable treatment plans for high-dose-rate prostate brachytherapy. For this, we present an automatic bi-objective parameter-tuning approach for inverse planning methods. The treatment planning problem is modeled as a bi-objective optimization problem, in which DVI-based planning criteria related to target coverage are explicitly separated from organ-sparing criteria. In our parameter-tuning approach, the parameters of inverse planning methods are automatically tuned, aimed to maximize the two objectives of the bi-objective planning model. Automatic parameter tuning furthermore allows to construct standard parameter sets (class solutions) representing different trade-offs in a principled way, which can be directly used in current clinical practice. In this chapter, we consider the inverse planning methods IPSA and HIPO. Thirty-nine previously treated prostate cancer patients were included. We compare automatic parameter tuning, random parameter sampling, and optimizing the bi-objective planning model directly with MO-GOMEA. We show that for all patients, a set of plans with a wide range of trade-offs can be obtained using automatic parameter tuning. By tuning HIPO, better trade-offs can be obtained than by tuning IPSA. For most patients, automatic tuning of HIPO results in a trade-off curve similar to the trade-off curve obtained by optimizing the bi-objective planning model directly. Automatic parameter tuning was shown to improve obtained DVIs significantly compared to random parameter sampling. Finally, from the automatically-tuned plans, three class solutions were successfully constructed representing different trade-offs.

3.1 Introduction

High-dose-rate (HDR) brachytherapy is a treatment modality for prostate cancer with excellent therapeutic outcomes [10, 64]. In HDR prostate brachytherapy, catheters are implanted in the prostate. Within each catheter, a radioactive source can be temporarily stopped at multiple dwell positions to locally irradiate the surrounding tissue. Treatment planning is the process of determining the time this source stops at each dwell position, which is essential for delivering an effective treatment that optimizes the trade-off between tumor control and normal tissue

complications. The quality of a treatment plan is assessed by a visual inspection of the 3D dose distribution. In addition, and to quantify certain aspects, treatment planning criteria are often formulated in terms of dose-volume indices (DVIs) on the target volumes and organs at risk (OARs), derived from the dose-volume histogram [96, 183]. The purpose of treatment planning is to obtain a plan that has desirable trade-offs between the planning criteria. It was shown in Chapter 2 that achievable trade-offs are however patient specific, which makes treatment planning in clinical practice difficult and time consuming.

Inverse treatment planning methods aim to aid in the planning process by reducing planning time and effort [60]. In these methods, treatment planning is modeled as an optimization problem and optimization software is used to obtain high-quality plans according to the specified model. Two clinically widely-available planning methods are Inverse Planning Simulated Annealing (IPSA) [124] and Hybrid Inverse Plan Optimization (HIPO) [118]. In both IPSA and HIPO the treatment planning problem is modeled with a dose-penalty model. This model allows fast optimization, but resulting plans do not necessarily adhere to the planning criteria even when such plans exist [77, 94]. To overcome this, the planner needs to tune the parameters of the dose-penalty model and rerun the planning method, or adapt the plan manually using graphical optimization, either of which is difficult and time consuming as shown in Chapter 2.

Instead of using a dose-penalty model, treatment planning can also be approached by optimizing DVIs directly [159, 189]. One such approach is bi-objective planning, as discussed in Chapter 2. In the bi-objective planning model, the planning criteria in the clinical protocol are grouped into two separate objectives; one related to target coverage and one related to OAR sparing. Per objective, the grouped planning criteria are combined in a worst-case manner such that optimization searches for satisfactory plans, i.e., plans that satisfy all planning criteria. The result of optimizing the bi-objective planning model is not a single best plan, but a large set of plans that we call the *approximation set* or *trade-off set*, in which each plan has a different high-quality trade-off between the two objectives. Since the bi-objective planning model reduces all planning criteria to only two objectives, the trade-off set is easy to visualize as a trade-off curve, which shows maximally achievable plan quality in an insightful way. The trade-off curve can be used in clinical practice to intuitively select the preferable plan from the trade-off set. We showed in Chapter 2 that optimizing the bi-objective planning model with the Multi-objective Gene-pool Optimal Mixing Evolutionary Algorithm (MO-GOMEA) [37] resulted in plans that were considered clinically acceptable for all patients considered ($n = 18$). Furthermore, these automatically generated plans

were preferred over the clinically used plans in 98% of the cases, based on inspection of the 3D dose distribution and the achieved planning criteria.

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

We applied our automatic bi-objective parameter-tuning approach to the inverse planning methods IPSA and HIPO. We investigated how automatic parameter tuning differs from the maximally achievable plan quality obtained by directly optimizing dwell times with MO-GOMEA and from the clinically used treatment plans. We furthermore compared automatic bi-objective parameter tuning to random parameter sampling, i.e., a straightforward approach to generate a trade-off set, but without further optimizing the resulting plans [17].

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

3.2 Materials and methods

3.2.1 Patient data and treatment protocol

Between November 2014 and December 2018, 39 prostate cancer patients were treated in our medical center with a single-dose HDR brachytherapy of 13 gray (Gy) delivered a week after external beam radiation treatment with a dose schedule of 20×2.2 Gy. Catheter implantation was performed using transrectal ultrasound

under general or epidural anesthesia according to a pre-plan, made in the operating room in Oncentra Prostate (version 4.2.2, Elekta Brachytherapy, Veenendaal, the Netherlands) based on ultrasound imaging [169]. A transurethral catheter with a bladder balloon was used to enhance visibility of the urethra.

After catheter implantation, three orthogonal pelvic T2-weighted turbo spin echo Magnetic Resonance Imaging (MRI) scans (Ingenia 3.0T Philips Healthcare, Best, The Netherlands) with an in-plane resolution of 0.6×0.7 mm and 3.0 mm slice thickness with 0.3 mm gap were acquired and used for treatment planning. Imaging was done in supine position, similar to the treatment position. These images were loaded into Oncentra Brachy (version 4.3–4.5, Elekta Brachytherapy, Veenendaal, the Netherlands) and used for catheter reconstruction, delineation of the volumes of interest (prostate, vesicles, bladder, rectum, and urethra) and further treatment planning. The median delineated prostate volume on these images after catheter placement was 31.5 (range: 16.6–74.2) cm^3. Treatment planning criteria as used in our clinic are given in Table 2.1. All planning criteria are based on DVIs, and are adapted from the GEC-ESTRO HDR prostate guidelines [96]. As starting point for treatment planning, dwell positions were activated inside the target volumes plus a 5-mm margin, excluding dwell positions located within 1 mm of the urethra. A median of 17 (range: 14–20) catheters were implanted with a source-step of 2.5 mm, which resulted in a median of 230 (range: 121–353) active dwell positions.

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

3.2.2 Dose-volume calculations

We compared different approaches to generate treatment plans, either obtained using an inverse planning method, or using MO-GOMEA. To compute the dose distribution and DVIs corresponding to a list of dwell times, we used our in-house

developed TG-43 [180] dose engine, which was validated against Oncentra Brachy [202]. This dose-engine is based on randomly sampled dose-calculation points within each volume of interest. Using a larger number of dose-calculation points results in more precise estimates of the DVIs, but also requires more computation time and more computer memory. In this chapter, we chose to use a larger number of dose-calculation points than would be strictly necessary in clinical practice to show the maximally achievable performance of the methods. During optimization of the bi-objective planning model, 10^5 dose-calculation points were used for both the bladder and the prostate, $5 \cdot 10^4$ for the rectum, and $2 \cdot 10^4$ for both the vesicles and urethra, which were heuristically chosen based on the average size of each volume to provide sufficiently accurate computations of the DVIs [202].

To overcome a bias towards the dose-calculation points used in the optimization process, all obtained plans were afterwards re-evaluated based on newly sampled dose-calculation points; 10^5 for each of the five volumes of interest. This is the same number of dose-calculation points as used in the dose-volume histogram calculations of Oncentra Brachy in clinical practice in our medical center.

3.2.3 Dose-penalty model setup

The dose-penalty model, used by IPSA and HIPO, can be optimized efficiently, but does not directly relate to the planning criteria [77]. To set up the dose-penalty model, the planner specifies lower and/or upper dose limits per volume of interest, together with a corresponding importance weight for each limit. Then, when a dose-calculation point receives a dose that is lower than the lower dose limit or higher than the upper dose limit, a penalty is given. Finally, all penalties are multiplied by the corresponding importance weight, and the final score of a plan is the sum of all these weighted penalties [118, 124]. The optimization software within IPSA and HIPO then searches for a plan with the lowest total score.

We defined a lower dose limit on the target volumes (prostate and vesicles), and an upper dose limit on the OARs (bladder, rectum, and urethra). Setting an upper dose limit on the target volumes was found to be deteriorating plan quality in preliminary experiments on a limited number of patients, and these limits were thus not set. This is mainly due to our clinical planning criteria for the high-dose regions in the targets (i.e., $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$, see Table 2.1), which are relatively easy to satisfy, as was shown in Chapter 2.

This setup results in two parameters that need to be specified per volume of interest: a dose-limit value and its corresponding weight. Tuning both the dose limit and weight at the same time complicates parameter tuning, as, for example, decreasing the upper dose limit has a similar effect as increasing the corresponding

weight. Therefore, only one of the two is tuned, and the other is set based on the following reasoning. For the target volumes, the planning criteria are defined in terms of volume-indices, i.e., $V_{100\%}^{\text{prostate}} > 95\%$ and $V_{80\%}^{\text{vesicles}} > 95\%$. For such planning criteria, the maximum value can be achieved if the lower dose limit is set to the corresponding dose, i.e., 100% and 80% of the planning-aim dose respectively, and only the corresponding weight is tuned.

For the OARs, however, the planning criteria are defined in terms of dose-indices, e.g., $D_{2\text{cm3}}^{\text{bladder}} < 74\%$. When one would set the bladder upper dose limit to 74% of the planning-aim dose and the corresponding weight large, the best value that can be achieved is $D_{2\text{cm3}}^{\text{bladder}} = 74\%$, but never $D_{2\text{cm3}}^{\text{bladder}} < 74\%$, as the underlying dose-penalty model does not assign penalties below the upper dose limit. For planning criteria based on dose-indices, the upper dose limit is thus to be tuned, while we set the dose-limit based on the clinical protocol. There are however two planning criteria specified for the bladder and rectum in the clinical protocol, but only one upper dose limit can be specified per volume of interest in the dose-penalty model. Therefore, this dose limit needs to capture both criteria.

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

IPSA setup

IPSA is an inverse planning method that uses simulated annealing to optimize the underlying dose-penalty model [124]. We set up IPSA as shown in Table 3.1. Upper dose limits on OARs were only specified on the surface. As there are no dwell positions activated within the OARs, the highest dose will occur at the surface, and reducing surface dose will result in reduced total dose.

In contrast to the other methods discussed in this chapter handles IPSA dwell position activation internally. For this, the user must specify activation margins. We set these activation margins as currently specified in our clinical protocol (see Section 3.2.1), to minimize the difference between methods. However, internal

Table 3.1: Used parameter settings for the inverse planning methods IPSA and HIPO. For both methods, four parameters remain to be set by automatic parameter tuning: The lower dose limits L_b, L_r, and L_u for respectively the bladder, rectum, and urethra, and an overall weight W_{OARs} representing the importance of the organs at risk (OARs) with respect to the targets. Dose limits are given in percentage of the planning-aim dose of 13 Gy. Weights have no unit.

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

computations by IPSA might still result in a slightly different set of activated dwell positions. The number of dose-calculation points that IPSA uses internally cannot be controlled.

Finally, a Dwell Time Deviation Constraint (DTDC) can be specified for IPSA. The DTDC aims to increase dose homogeneity by limiting dwell-times differences within a catheter. Dose homogeneity is not a planning criterion in our clinical protocol, and there is thus no direct incentive to enable the DTDC, especially as it has been shown that it can deteriorate overall plan quality [13]. The DTDC is therefore disabled.

HIPO setup

In HIPO, the limited-memory Brodyen-Fletcher-Goldfarb-Shanno algorithm (L-BFGS) is used to optimize the underlying dose-penalty model [155]. The parameter setup used for HIPO is shown in Table 3.1. Dose limits are applied to both the surface and the volume depending on the sample point settings. According to our clinical practice, for each OAR, 500 surface points and 500 volume points were used. For each target, we used 1000 volume points, plus surface points with a density of 8 cm^{-2}. Compared to IPSA, HIPO requires two additional dose limits; an upper dose limit on normal tissue dose and an upper dose limit on the primary

target volume (prostate). To reduce the differences between the optimization aims in IPSA and HIPO, we effectively disabled these by setting the dose limit to 400% and the corresponding weight to 0.001, the lowest weight possible. Furthermore, HIPO does not accept overlapping volumes. Therefore, the intersection between the prostate and the urethra was set to only belong to the urethra. This is different from the interpretation by IPSA, Oncentra Brachy, and our dose engine, as there, intersections between volumes are considered part of both volumes. Similar to the DTDC of IPSA, the Dwell Time Gradient Restriction (DTGR) in HIPO was disabled.

3.2.4 The bi-objective planning model

In the bi-objective planning model, the planning criteria are grouped into one coverage objective and one sparing objective, that are referred to as the Least Coverage Index (LCI) and Least Sparing Index (LSI). Both LCI and LSI are constructed by combining the coverage or sparing criteria in Table 2.1 in a worst-case manner. For example, maximizing the LCI implies that the least-fulfilled criterion is improved. The criterion that is least fulfilled changes during the optimization process, thus, ultimately, all criteria are improved. This is similar to how planners were observed to typically perform manual optimization of clinical plans, by trying to improve the least-fulfilled criterion iteratively, as visualized in Figure 2.4. The LCI and LSI are given by,

$$\begin{aligned} \text{LCI} &= \min\left\{V_{100\%}^{\text{prostate}} - 95, V_{80\%}^{\text{vesicles}} - 95\right\}, \\ \text{LSI} &= 13\ \text{Gy} \times \min\left\{\begin{array}{l} 86 - D_{1\text{cm3}}^{\text{bladder}}, 74 - D_{2\text{cm3}}^{\text{bladder}}, 110 - D_{0.1\text{cm3}}^{\text{urethra}}, \\ 78 - D_{1\text{cm3}}^{\text{rectum}}, 74 - D_{2\text{cm3}}^{\text{rectum}}, 150 - D_{50\%}^{\text{prostate}}, \\ 200 - D_{20\%}^{\text{prostate}} \end{array}\right\}. \end{aligned} \tag{3.1}$$

An important property of the bi-objective planning model is that when both LCI $> 0\%$ and LSI > 0 Gy, the plan is satisfactory. Optimization of the bi-objective planning model thus aims directly for satisfactory plans. This is a property that does not hold for the weighted-sum approaches of IPSA and HIPO, where unfulfilled criteria can be compensated by fulfilled criteria.

Adding the sparing criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ directly to the LSI would lead to a comparison of planning criteria with different quantities (volume and dose), which could give unexpected results as these are not directly comparable. In Chapter 2, these two criteria were not included in the LSI, but added as hard constraints. It was found that these hard constraints did not influence the results, as these were never found to be violated in the considered

patient set, which is a subset of the patients considered in this chapter. However, it cannot be guaranteed that these hard constraints are not violated on additional patients and with different planning strategies considered in this chapter. For completeness, the planning criteria are therefore not added as hard constraints but rewritten in terms of dose-indices so that they can be directly added to the LSI. The sparing criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ are therefore replaced by respectively $D_{50\%}^{\text{prostate}} < 150\%$ and $D_{20\%}^{\text{prostate}} < 200\%$.

3.2.5 Trade-off set generation with IPSA and HIPO

Using the setup of the dose-penalty model in Table 3.1, four parameters need to be tuned, i.e., L_b, L_r, L_u, and W_{OARs} for either IPSA or HIPO, distinguished by a superscript. These four parameters were automatically tuned by optimizing the bi-objective planning model, i.e., by finding a set of plans with different high-quality trade-offs in the LCI and LSI. For this, the evolutionary algorithm MAMaLGaM [181] was used (full-covariance version). By automatic parameter tuning that entails running IPSA or HIPO, the dose distribution can no longer be quickly updated, and therefore, MAMaLGaM, a similar type of algorithm as MO-GOMEA, has better performance [131]. MAMaLGaM was used to perform the automatic bi-objective parameter tuning throughout this chapter, while MO-GOMEA was used only for direct optimization of the bi-objective planning model.

Per patient, MAMaLGaM was run for 25 iterations with a standard population size of 410 solutions, resulting in about 10^4 runs of each of the inverse planning method. This number of iterations was determined based on preliminary testing on a limited set of patients, and obtained trade-off curves were found to be non-improving for the last five iterations in all runs. The typical values for W_{OARs} that were found by preliminary testing were between 0 and 2. Therefore, for the start of the automatic parameter tuning, W_{OARs} was initialized randomly in the interval $[0, 2]$. An upper weight of 2 implies that dose penalties for the OARs were initially considered to be at most twice as important as the target coverage penalties. To make sure that this range is not too restrictive, we set an upper bound of 10.

The upper dose limits for the OARs were initialized around the target value stated in our clinical protocol. That is, L_b and L_r are initialized randomly in the range 60%–100%, and L_u in the range 100%–110%. Dose limits were bounded between 0% and 400%. The initialization ranges for the bladder and rectum are relatively large because both these volumes have two planning criteria specified in the clinical protocol. Additionally, these planning criteria are specified for a larger volume compared to that for the urethra, i.e., 1 cm^3 and 2 cm^3 compared to 0.1 cm^3, which implies that a higher dose than the planning criterion is allowed,

as long as the volume receiving that higher dose is not too large. To incorporate this in the dose-penalty model, we used a larger initialization range.

MAMaLGaM is a stochastic algorithm, that, depending on the initialization of the random number generator, finds a different trade-off set every run, although previous work shows the variation in outcome to be small [131]. We further reduced the effect of this randomness by running MAMaLGaM 10 times per patient for each inverse planning method, and since we aim to compare to the best possible trade-off set, the resulting 10 trade-off sets were combined into a single trade-off set by maintaining only the best plans. We refer to the resulting trade-off sets obtained as *IPSA-tuned* and *HIPO-tuned*.

As alternative to trade-off set generation via automatic parameter tuning, we generated 1000 treatment plans by randomly sampling values for the parameters of the dose-penalty model (i.e., L_b, L_r, L_u, and W_{OARs}) for both IPSA and HIPO, similar as in [17]. Parameter values were sampled uniformly random in the same ranges as those that were used for the initialization of automatic bi-objective parameter tuning. The resulting trade-off set is then the subset of plans with only the best trade-offs. As no further optimization is performed in this approach, we did not repeat it multiple times like we did for MAMaLGaM. We refer to the trade-off sets generated via this approach as *IPSA-random* and *HIPO-random*.

3.2.6 Trade-off set generation with MO-GOMEA

The bi-objective planning model is non-convex, non-linear, and non-smooth, and to optimize plans according to it, the state-of-the-art multi-objective real-valued MO-GOMEA, was used [36, 131]. MO-GOMEA exploits that the dose distribution can be quickly updated when only few dwell times change. MO-GOMEA does not require any patient-specific parameters to be set, and a single run results in a trade-off set of hundreds of plans. For the experiments in this chapter, MO-GOMEA was run for 10 minutes on a Graphics Processing Unit (GPU) (NVIDIA Titan X). In clinical practice, it was shown that a runtime of 3 minutes was sufficient [36]. We use a longer runtime however to show maximally achievable plan quality, and because we increased the number of dose-calculation points, see Section 3.2.2.

MO-GOMEA is a stochastic algorithm, like MAMaLGaM. Parameter tuning was therefore repeated 10 times per patient. The resulting trade-off sets were combined into a single trade-off set, that we refer to as *GOMEA-direct*.

3.2.7 Method comparisons

For each of the 39 patients, trade-off sets were automatically constructed using the three approaches discussed above, which results in five trade-off sets per patient

(IPSA-tuned, HIPO-tuned, IPSA-random, HIPO-random, and GOMEA-direct). These trade-off sets were visually compared as trade-off curves, together with the clinically-used plan, in the same figure. We compared the best obtained results (IPSA-tuned and HIPO-tuned) with GOMEA-direct. Moreover, we investigated the added value of performing parameter tuning by comparing IPSA-tuned with IPSA-random, and similarly HIPO-tuned with HIPO-random.

When a trade-off set is visualized as a trade-off curve, the area under this curve can be used as a performance indicator, which is known as the *hypervolume* [7, 220]. A larger hypervolume implies that the planning method achieved better trade-offs. The hypervolume is a performance measure that in itself may be hard to interpret. Its value depends on a reference point, i.e., the lower cut-off values of the trade-off curve, and it has no interpretable unit (in our case, dose times volume). It does however allow for a straightforward comparison of multiple trade-off curves, where a larger hypervolume is better. Differences in hypervolume were tested for statistical significance by a Wilcoxon signed-rank test with $\alpha = 0.05$.

Random parameter sampling versus automatic parameter tuning

All plans corresponding to the randomly sampled parameters were visualized as a scatterplot in the same figure (and not just the plans in the trade-off set, that have the best trade-offs). This gives insight in how many of the resulting plans are *satisfactory*, i.e., that satisfy all planning criteria. Additionally, it shows the distribution of obtained trade-offs. Further, the percentage of satisfactory plans was computed, which is an estimate of the probability that a randomly sampled parameter set results in a satisfactory plan, similar as in [17].

IPSA/HIPO-tuned versus GOMEA-direct

We aimed to investigate maximum performance of each inverse planning method. IPSA-tuned, HIPO-tuned, and GOMEA-direct are therefore compared in more detail. To assess quality of individual plans, three reference plans were selected from each trade-off set. For this, the following selection rules were used: a left (L) plan with an LCI-value of approximately -2%, a middle (M) plan with an LCI-value of approximately 1%, and a right (R) plan with an LCI-value of approximately 4%, corresponding to plans that have a target coverage of at least 93%, 96%, and 99% respectively. Of these selected plans, DVIs of the planning criteria on which the LCI and LSI are based (Section 3.2.4) were compared. Differences were tested for statistical significance by a Wilcoxon signed-rank test with $\alpha = 0.05$, and Bonferroni correction was applied for the 60 pairwise tests performed (IPSA/HIPO versus MO-GOMEA for plans L/M/R for each of 10 DVIs), resulting in a corrected $\alpha = 0.0008$.

3.2.8 Constructing class solutions

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

We used plans obtained via automatic parameter tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned) as reference plans to construct multiple class solutions for IPSA and HIPO. By averaging the four parameters (L_b, L_r, L_u, and W_{OARs}) of the reference plans L, M, and R over all patients, we constructed three class solutions CS-L, CS-M, and CS-R, representing different trade-offs. Plan quality of these class solutions was compared to the automatic per-patient tuned results (IPSA-tuned or HIPO-tuned).

We recorded how many of the resulting plans were satisfactory, and analyzed the obtained DVI values. To prevent overfitting to the current patient set, leave-one-out cross validation was applied. One patient was removed from the set of 39 patients, and a class solutions was constructed from the remaining 38 patients. This class solution was then validated on the patient that was left out. We repeated this procedure by leaving each patient out once. Only results on the validation sets are shown.

The data set is rather small with only 39 patients. Therefore, leave-one-out cross validation is a good choice, as the training set is kept as large as possible (38 patients). Leave-one-out cross validation is often time consuming as it needs to be repeated for each data point (patient), thus 39 times. Since constructing class solutions by taking means is extremely fast, as there is no need to re-do the tuning of IPSA or HIPO, this was not an issue here.

3.3 Results

3.3.1 Trade-off curves

Trade-off curves based on the trade-off sets generated by automatic parameter tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned), and by MO-GOMEA (GOMEA-direct) are shown in Figure 3.1. Parameter tuning of IPSA had a median run time of 3 hours per patient (range: 1.5–5h) and of HIPO 23 hours (range: 11.5–42h), using the clinical software implementation of both methods.

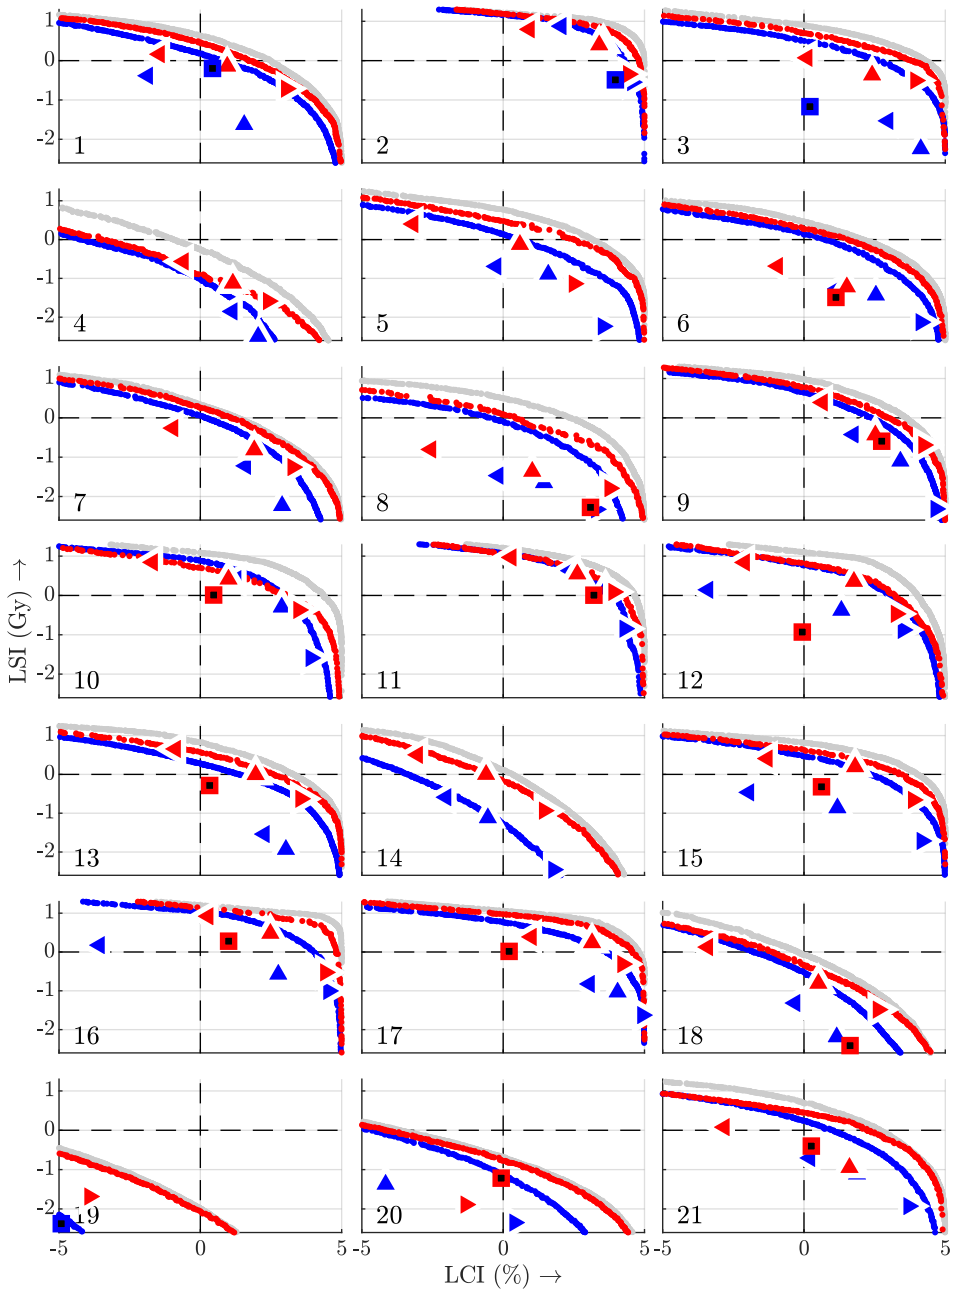

Figure 3.1: (Part 1): Trade-off curves per patient (patient number in lower left corner) based on the trade-off sets obtained with IPSA-tuned, HIPO-tuned, and GOMEA-direct. Each obtained plan is visualized as a single dot. Triangles are plans obtained by the class solutions (CS) L (<), M (∧) and R (>), IPSA in blue, HIPO in red. (...)

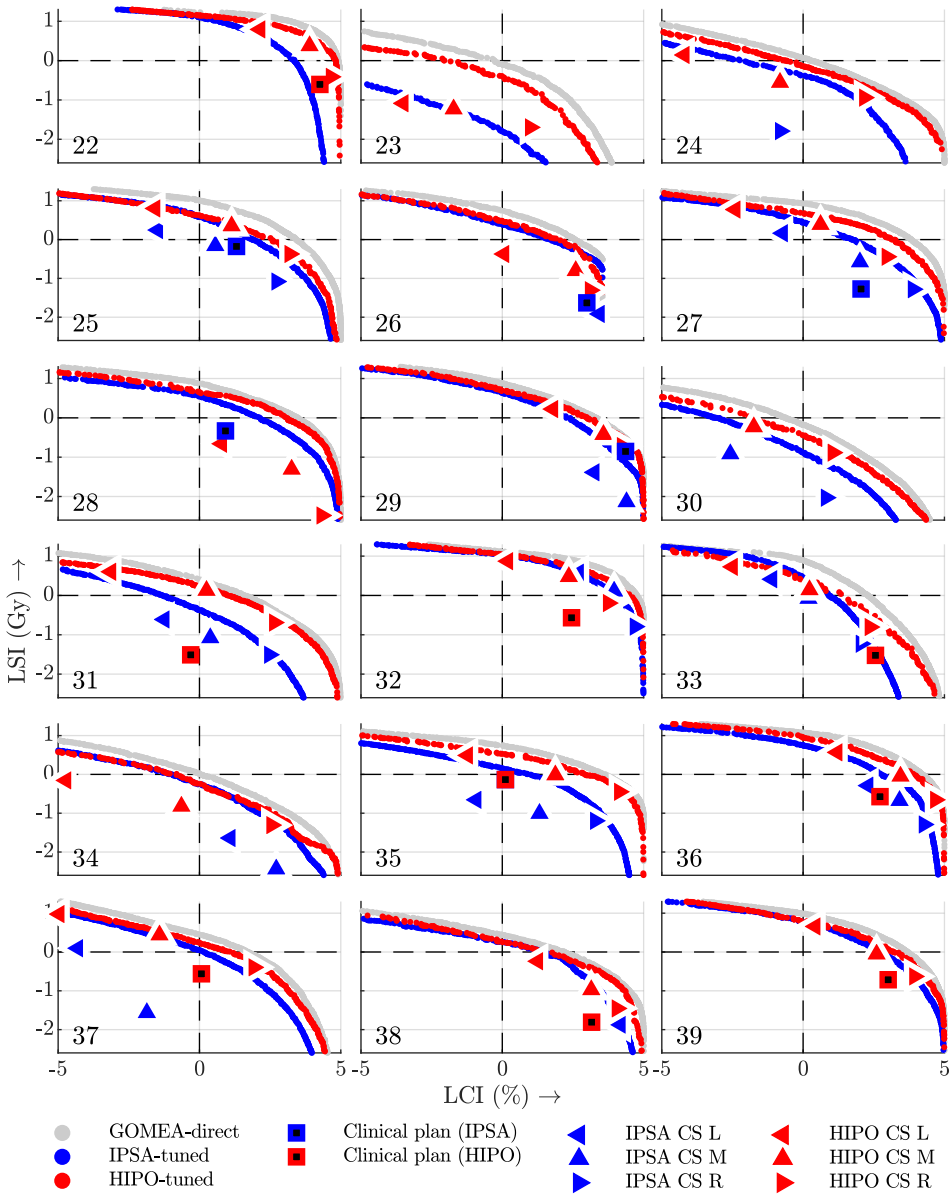

Figure 3.1: (Part 2): (...) Black squares represent the trade-off of the clinical plan, outlined in the color of the inverse planning method from which treatment planning was initialized. Some clinical plans and class solution plans fell outside the figure range because they violated one or multiple planning criteria too much.

Computations were performed on different machines, which affected run time, but computation time is patient dependent and relates mainly to the number of active dwell positions, i.e., the number of variables that IPSA or HIPO need to optimize.

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

Only for patients 10 and 33, automatic parameter tuning of IPSA achieved somewhat better trade-offs than that of HIPO for lower LCI values, i.e., on the left side of the trade-off curve. For all other patients, trade-offs obtained by automatic parameter tuning of IPSA were inferior to trade-offs obtained by both automatic parameter tuning of HIPO and by MO-GOMEA. The hypervolume of GOMEA-direct was, on average over all patients, 29.7 (range: 7.7–37.9), while the hypervolume of HIPO-tuned and IPSA-tuned were significantly lower ($p < 0.01$) with average values of respectively 27.6 (7.0–37.0) and 24.6 (0.2–35.8).

Obtained trade-off curves clearly differ for different patients in shape and maximally achievable plan quality (see Figure 3.1). For 28/39 patients (72%), all three methods obtained satisfactory plans. Plans obtained by automatic tuning of HIPO were satisfactory for two additional patients (8 and 31), resulting in 30/39 (77%) patients with satisfactory plans. MO-GOMEA obtained satisfactory plans for four additional patients (8, 14, 24, and 31), resulting in satisfactory plans for 32/39 (82%) patients. Finally, for seven (18%) patients (4, 18, 19, 20, 23, 30, and 34), none of the methods obtained satisfactory plans.

Worst plans were observed for patient 19. For this patient, the catheters were not placed deep enough in one region of the prostate. To still be able to fully cover the prostate, the planning criteria on the bladder or urethra were substantially violated. There was no structural difference observed in prostate volume between patients without satisfactory plans (median 29.3 cm^3) and the rest (median 31.5 cm^3).

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

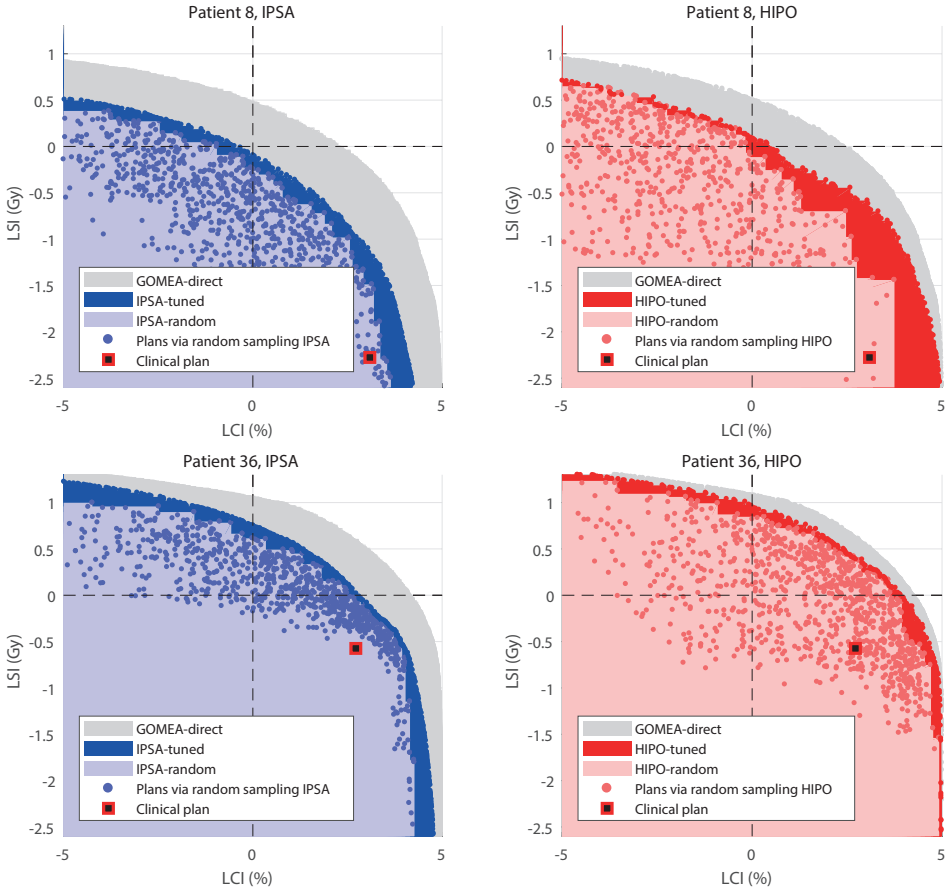

Figure 3.2: Plans obtained via random parameter sampling are illustrated by scattered dots (IPSA/HIPO-random), compared to the maximally achievable trade-offs obtained by automatically tuning IPSA/HIPO (IPSA/HIPO-tuned), and the maximally achievable trade-offs for that patient, obtained by MO-GOMEA (GOMEA-direct). Illustrated for patients 8 (top) and 36 (bottom). Squares indicate the clinical plan. To highlight the differences between these approaches, the areas under the curves are shaded. Note that the areas overlap, e.g., the area under the curve of GOMEA-direct also includes the area shaded in blue/red.

3.3.2 Method Comparisons

Random parameter sampling versus automatic parameter tuning

For two example patients (8 and 36) that are representative for the patient set, plans obtained by randomly sampling parameters are shown in Figure 3.2, together with the corresponding trade-off sets IPSA-random and HIPO-random, in which only the plans with the best trade-offs were maintained.

A gap can be observed between the IPSA-tuned and IPSA-random trade-off curves, and similar for HIPO-tuned and HIPO-random. The mean hypervolume for IPSA-random was 22.2 (range: 0.0–33.7), which is lower than the hypervolume of IPSA-tuned ($p < 0.01$). Similar, the hypervolume for HIPO-random was 25.4 (range: 5.4–36.0), which is lower than the hypervolume of HIPO-tuned ($p < 0.01$). Comparing the 28 patients for which both with automatic parameter tuning of IPSA and HIPO satisfactory plans were obtained, a median of 7.5% (range: 0–63%) of the IPSA-random plans were satisfactory. For HIPO-random, this was 14.5% (range: 0–78.5%).

IPSA/HIPO tuned versus GOMEA-direct

Table 3.2 shows the average DVI values over all patients obtained by directly optimizing the bi-objective planning model by automatic tuning of IPSA and HIPO (IPSA-tuned and HIPO-tuned), and MO-GOMEA (GOMEA-direct) for the three reference plans L, M, and R. In all cases, plans obtained by MO-GOMEA spared the urethra and bladder better while similar prostate coverage $V_{100\%}^{\text{prostate}}$ was achieved, compared to IPSA-tuned and HIPO-tuned plans.

Plans resulting from the automatic tuning of HIPO have lower high-dose regions within the prostate compared to plans resulting from direct optimization by MO-GOMEA or by automatic tuning of IPSA. IPSA-tuned and HIPO-tuned plans exhibit significantly better vesicle coverage $V_{80\%}^{\text{vesicles}}$, albeit all three methods achieved values well above the planning criterion of 95%. The values obtained for the planning criteria on the rectum were well below their aim for all patients by all three methods.

3.3.3 Class solutions

The three class solutions obtained for both IPSA and HIPO, based on the three selected plans from the trade-off sets IPSA-tuned and HIPO-tuned are shown in Table 3.3. The obtained DVI values, as a result of leave-one-out cross-validation, are shown in Table 3.4. Further, the corresponding trade-offs are visualized in Figure 3.1. Combined, the three IPSA class solutions resulted in satisfactory plans for 4 patients. HIPO class solutions resulted in satisfactory plans for 18 patients

Table 3.2: Dose-volume indices of selected plans from IPSA-tuned and HIPO-tuned. Mean (and standard deviation) shown over all 39 patients. (+) indicates a statistically significant better value compared to MO-GOMEA plans, (-) indicates a significant worse value. Volume-indices (V) in volume percentage, dose-indices (D) in Gy.

Plan	Method	$V_{100\%}^{\text{prostate}}$	$V_{80\%}^{\text{vesicles}}$	$D_{50\%}^{\text{prostate}}$	$D_{20\%}^{\text{prostate}}$	$D_{90\%}^{\text{prostate}}$
	criteria	$> 95\%$	$> 95\%$	<19.5 Gy	<26 Gy	>13 Gy
L	MO-GOMEA	93.0 (0.1)	97.0 (2.5)	15.8 (0.6)	19.9 (1.4)	13.3 (0.1)
	IPSA	93.0 (0.4)	98.6 (2.2) +	16.2 (0.9) -	20.8 (2.3) -	13.3 (0.2) +
	HIPO	93.0 (0.2)	98.1 (2.2) +	15.4 (0.5) +	19.4 (1.6) +	13.2 (0.1) -
M	MO-GOMEA	96.0 (0.1)	98.2 (1.5)	16.1 (0.7)	20.3 (1.6)	13.6 (0.3)
	IPSA	95.9 (0.8)	99.2 (1.3) +	16.8 (1.0) -	21.7 (2.4) -	13.8 (0.3) +
	HIPO	96.0 (0.0)	98.8 (1.3) +	15.7 (0.6) +	19.9 (1.8) +	13.5 (0.2) -
R	MO-GOMEA	98.9 (0.4)	99.7 (0.4)	16.8 (0.9)	21.4 (2.0)	14.2 (0.5)
	IPSA	98.6 (1.4) -	99.7 (0.7)	18.1 (1.0) -	23.4 (2.4) -	14.8 (0.5) +
	HIPO	98.9 (0.5)	99.7 (0.4)	16.4 (0.8) +	20.9 (2.1) +	13.9 (0.4) -

Plan	Method	$D_{1\text{cm}^3}^{\text{bladder}}$	$D_{2\text{cm}^3}^{\text{bladder}}$	$D_{1\text{cm}^3}^{\text{rectum}}$	$D_{2\text{cm}^3}^{\text{rectum}}$	$D_{0.1\text{cm}^3}^{\text{urethra}}$
	criteria	<11.2 Gy	<9.6 Gy	<10.1 Gy	<9.6 Gy	<14.3 Gy
L	MO-GOMEA	9.6 (0.8)	8.6 (0.7)	8.7 (0.9)	7.7 (0.9)	13.5 (0.5)
	IPSA	9.9 (0.6) -	9.0 (0.5) -	9.0 (1.1) -	8.0 (1.0) -	13.8 (0.7) -
	HIPO	9.8 (0.6) -	8.9 (0.6) -	9.0 (0.9) -	8.0 (0.9) -	13.6 (0.5) -
M	MO-GOMEA	10.2 (1.0)	9.1 (0.9)	9.1 (1.1)	8.0 (1.1)	14.0 (0.7)
	IPSA	10.5 (0.7) -	9.5 (0.7) -	9.3 (1.2) -	8.3 (1.2) -	14.4 (0.9) -
	HIPO	10.4 (0.8) -	9.4 (0.7) -	9.2 (1.1)	8.2 (1.1)	14.1 (0.7) -
R	MO-GOMEA	11.4 (1.1)	10.3 (1.0)	9.6 (1.3)	8.6 (1.3)	15.0 (1.0)
	IPSA	11.7 (0.9) -	10.5 (0.8)	9.9 (1.4)	8.8 (1.3)	15.6 (0.9) -
	HIPO	11.5 (1.0)	10.4 (0.9) -	9.9 (1.2)	8.9 (1.2) -	15.1 (0.9) -

in total, with class solutions based on plan M as the single best class solution with satisfactory plans for 13 patients. The class solutions based on plan R, aimed to result in plans with 99% coverage, violated the sparing criteria in all but one case.

For both IPSA and HIPO, standard deviations between obtained class solution parameters were larger for bladder and rectum limits L_b and L_r, and notably for the IPSA bladder limit, while the standard deviation for the urethra limit L_u was small. Compared to IPSA, HIPO class solutions had lower limits L_b and L_r for all plans, but somewhat higher L_u. In Figure 3.1, it can be seen that the ordering of

Table 3.3: Class solutions (CS) for IPSA and HIPO, obtained by averaging the parameters of selected plans from the trade-off sets IPSA-tuned and HIPO-tuned. Dose limits L_b, L_r, and L_u corresponding to the bladder, rectum and urethra respectively, are expressed in percentage of the planning-aim dose of 13 Gy. W_{OARs} is expressed as the relative importance of the organs at risk (OARs) to the target volumes, see Table 3.1. Standard deviations of the corresponding parameters over all patients are provided between brackets. 'Satisfactory' indicates for how many of the 39 patients the resulting plan satisfied all planning criteria.

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

the class solutions such that LCI(plan L) $\leq$ LCI(plan M) $\leq$ LCI(plan R) was largely maintained. The class solutions were aimed to result in plans with a coverage of 93%, 96%, and 99% respectively, which was achieved on average, but the standard deviation of $V_{100\%}^{\text{prostate}}$ and $V_{80\%}^{\text{vesicles}}$ was found to be large, compared to the standard deviations in Table 3.2.

3.4 Discussion

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

Plans obtained by automatic tuning of HIPO and by direct optimization of the bi-objective planning model with MO-GOMEA outperformed the trade-off of the clinical plan for all patients, while plans obtained by automatic tuning of IPSA outperformed the clinical plan for all but two patients (22 and 29). We would like to note that according to our clinical practice, planners often choose for a clinical plan with a higher coverage than strictly required to satisfy the coverage

Table 3.4: Dose-volume index values of class solutions (CS) for IPSA and HIPO. Results were obtained by leave-one-out cross validation. Average and standard deviation (between brackets) computed over all 39 patients, volume-indices (V) in volume percentage, and dose-indices (D) in Gy.

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

criteria, at the cost of (slightly) violating one or more sparing criterion, even though satisfactory plans might have been achieved for those patients by manual parameter tuning. For seven patients, it was not possible to obtain satisfactory plans with any of the approaches, which was caused by unfavorable implant geometry. There was no relation found between prostate volume and satisfactory plans.

The algorithms that are used for optimization in this chapter, MO-GOMEA, MAMaLGaM, and also simulated annealing within IPSA and L-BFGS within HIPO, are heuristic search algorithms that have no guarantee that maximum plan quality is reached. These methods aim to obtain a high-quality solution within a reasonable amount of time and effort. Previous work indicates that the obtained results of MO-GOMEA and MAMaLGaM, on these and similar problems, are near to maximally achievable quality [36, 38, 131, 181]. Throughout this chapter we referred for simplicity, to this as maximally achievable plan quality.

The random sampling approach we used was similar to the approach used in [17], albeit on a different protocol and patient set. In that work, no notion of distance to maximally attainable values was given. Here, we showed that plan quality can be further improved by performing automatic parameter tuning.

The ultimate goal of treatment planning is to obtain a treatment plan that represents the best trade-off between (local) tumor control and normal tissue

complications, and ideally, that is what is directly optimized for. The bi-objective parameter tuning approach presented in this chapter is a black-box method that does not rely on the particular implementation of the objectives. Instead of using the LCI and LSI as objectives, the tumor control probability (TCP) and normal tissue complication probability (NTCP) could also be used. To compute the TCP and NTCP, biological models have been developed [16]. These models however depend on a number of radio-biological parameters which are unknown or uncertain in practice. Although these models give a correct description of the main characteristics of the radiation response, caution has to be taken if these models are to be applied to patients [40, 59, 105]. When the clinical planning criteria are formulated in terms of TCP/NTCP, this can be rather straightforwardly incorporated into our automatic parameter tuning. However, to obtain the best treatment plans in terms of the current clinical protocol, which is formulated in terms of DVIs, it is logical to directly optimize for that, which is achieved by the bi-objective problem formulation.

Other multi-objective approaches exist [48, 49, 118, 155], but the bi-objective planning model specifically allows for a direct optimization of the DVI-based planning criteria without having too many objectives. Optimizing all planning criteria as separate objectives results in a many-objective optimization problem, and solving these problems by presenting a representative trade-off set is difficult and time consuming, even with state-of-the-art algorithms [132]. Additionally, visualization of the trade-off set when using four or more objectives is no longer straightforward.

In practice, besides the planning criteria, planners also look at the 3D dose distribution to assess plan quality. It was shown in the previous chapter that optimizing the bi-objective planning model with MO-GOMEA resulted in plans that were considered clinically acceptable. However, because of fundamental differences in how plans are generated by IPSA and HIPO, structural differences in the dose distribution might have occurred that are harder to quantify, besides the observed differences in the obtained values for the planning criteria.

Well-spread high-quality trade-off sets were obtained when applying automatic bi-objective parameter tuning to both IPSA and HIPO. This suggests that the chosen parameter setup, with only four parameters of the dose-penalty model that were automatically tuned, was sufficiently flexible.

We chose in this chapter not to include upper dose limits for the prostate in IPSA and HIPO, mainly based on prior observations with using MO-GOMEA to solve the bi-objective problem formulation that the planning criteria $V_{150\%}^{\text{prostate}} < 50\%$ and $V_{200\%}^{\text{prostate}} < 20\%$ are rather easy to satisfy, as shown in Chapter 2. The results in

Table 3.2 show that, even without upper dose limits on the prostate, the planning criteria are well below their aim. Hence, these planning criteria do not play a role in the automatic parameter tuning process of IPSA and HIPO, nor in the optimization process of MO-GOMEA. For HIPO, the obtained $V_{150\%}^{\text{prostate}}$ and $V_{200\%}^{\text{prostate}}$ values are even better than with the other methods. This shows that it is indeed not required for IPSA and HIPO to furthermore add upper dose limits for the prostate, and that the other upper dose limits reduce high-dose regions in the prostate below their aim. Despite this observation, including more parameters of the dose-penalty model in the automatic tuning process could potentially improve resulting plan quality. However, this increases computation time of the automatic tuning process, as more parameters need to be tuned. It would furthermore complicate the tuning process, as multiple parameters have the same effect on the obtained plan, of which we gave an example in Section 3.2.3. Increasing planning complexity increases the risk of ultimately achieving inferior plans, both for manual and automatic parameter tuning. When applying automatic bi-objective parameter tuning to a different clinical protocol, for a fair comparison between IPSA, HIPO, and MO-GOMEA, the parameter setup needs to be reconsidered. The reasoning and results in this chapter could be used as guidelines to do so.

Automatic bi-objective parameter tuning resulted for IPSA in inferior trade-offs compared to HIPO for many patients, although a similar setup of the dose-penalty model was used. For this, we used the clinical software implementation of IPSA and HIPO, and there are multiple fundamental differences between those two methods that could be the cause of these differences. In HIPO, a quadratic dose-penalty model is used, while in IPSA a linear dose-penalty model is used. Because of these model differences, different solvers were used to optimize them. In HIPO, L-BFGS is used, while simulated annealing is used in IPSA to optimize the dose-penalty model. Later improvements on IPSA showed that linear programming is capable of solving this dose-penalty model with higher accuracy, with up to about one percent improvement in DVIs of the prostate and urethra for some patients, but this version is not clinically available [4].

Other differences between the two methods are that for IPSA, dose-calculation points were only sampled on the surface of OARs, as only surface dose limits were specified, whereas for HIPO dose-calculation points were also sampled within the OAR volumes. In IPSA, the delineated contours specified by the planner were directly used to sample dose-calculations points, while in HIPO, the contours were converted in a three-dimensional mesh of the organ surface. Even when there are many delineated contours, the latter approach is closer to how DVIs are evaluated within Oncentra Brachy, and in our in-house-developed dose engine. Different

dose-point sampling and volume reconstruction methods have been shown to affect computed plan quality [202].

When looking into the DVI values of the optimized plans, it was noted that direct optimization of the bi-objective planning model with MO-GOMEA resulted in a lower vesicle dose than automatic tuning of IPSA and HIPO, albeit all obtained values were well above their aim. When MO-GOMEA optimizes the bi-objective planning model, it applies an internal weighting mechanism of the planning criteria aimed to prevent that the second-worst objective is not further optimized when the worst objective cannot be improved [36]. This mechanism was not applied in the automatic parameter tuning of IPSA and HIPO. However, it seems not to be required, as the obtained vesicle dose was found to be higher than obtained with MO-GOMEA. This might be a result of how the dose-penalty model was set up in the automatic parameter tuning approach, by using a single weight for both target volumes.

Automatic patient-specific parameter tuning, i.e., its application to daily clinical cases, as performed in this chapter, is not possible in current clinical practice due to the required computational effort. For that, GPU-based versions are required, which are starting to become available [17, 36]. A recent study showed that a GPU-based version of IPSA is able to compute 1000 plans within 10 seconds [17], which suggests that our automatic parameter-tuning approach can be run within approximately 4 minutes, and most likely even faster. MO-GOMEA was shown to be able to optimize the bi-objective planning model in 30 seconds on a GPU [36], or within 5 minutes on a Central Processing Unit (CPU) [130]. This is all well within the time limits of the clinical practice, especially since manual graphical optimization often takes 30 minutes or more, as was described in Chapter 2.

For clinical practice today, class solutions could be used for IPSA or HIPO. With the construction of class solutions based on the obtained HIPO-tuned (or IPSA-tuned) plans, we showed that multi-objective, i.e., trade-off representing, class solutions can effectively be obtained in a principled way. However, subsequent (automatic) patient-specific parameter tuning is still clearly required for IPSA and HIPO to obtain the highest-possible quality plans. Construction of class solutions can be further extended by splitting the patient set based on anatomical features, e.g., prostate size, and constructing a class solution for each subset might increase plan quality. Alternatively, patient features could be correlated to the method's parameters, as was done for IPSA in recently published works [48, 49]. Also then, high-quality reference plans are required for all patients that are generated by the inverse planning method of interest, as the corresponding parameters must be known. As discussed in Chapter 2, treatment plans may have been adapted

manually in practice, so that the corresponding parameter values are not known, or the plans may be sub-optimal because of the manual treatment planning process that has to be performed in limited time. Our bi-objective parameter-tuning approach is then still an essential tool that can generate high-quality reference plans automatically.

The three obtained class solutions for HIPO (Table 3.3) are currently being introduced in our clinical practice for which we will follow the following strategy: First, class solution M is applied. If the coverage of the resulting plan is too low, class solution R is applied, else, if the resulting coverage is too high, class solution L is applied. Then, manual parameter tuning or graphical optimization is performed to further improve plan quality. If manual parameter tuning is performed by a planner, it is recommended to follow the same strategy of adapting only the four parameters in Table 3.1. However, the observed difference in performance between the results from automatic parameter tuning and random parameter sampling do indicate that this is likely still not trivial to do.

3.5 Conclusion

We successfully developed an automatic bi-objective parameter tuning approach. When applied to IPSA and HIPO, it results in a set of plans with a wide range of trade-offs. By tuning HIPO, better trade-offs were obtained than by tuning IPSA. For most patients, automatic tuning of HIPO resulted in plans close to the maximally achievable plan quality obtained by optimizing the bi-objective planning model directly with MO-GOMEA. Automatic parameter tuning was furthermore shown to improve plan quality significantly compared to random parameter sampling. Finally, from the automatically tuned plans, three class solutions were successfully constructed representing different trade-offs that are being introduced in our clinical practice.

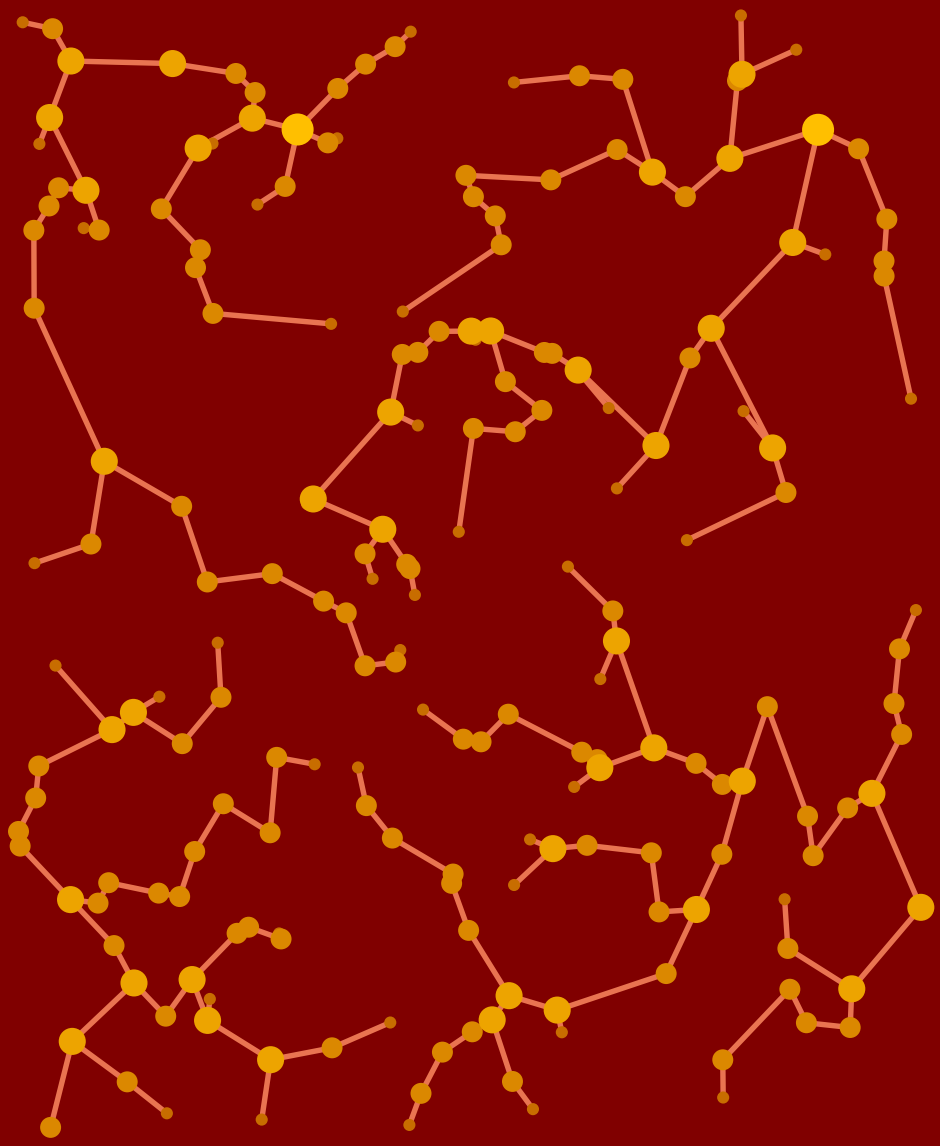

4

Multimodal optimization with the hill-valley evolutionary algorithm

Research question 3:
How can model-based evolutionary algorithms be adapted to perform efficient and effective (single-objective) multimodal optimization, without making assumptions on the number of modes, or on their shape or size?

This chapter is an adaptation of S.C. Maree, D. Thierens, T. Alderliesten, P.A.N. Bosman. (2020) *Two-phase Real-valued Multimodal Optimization with the Hill-Valley Evolutionary Algorithm.* In M. Preuss, M.G. Epitropakis, J.E. Fieldsend and X. Li (Eds.), *Metaheuristics for Finding Multiple Solutions*, Springer (in press).

Abstract

The aim of multimodal optimization is to obtain all global optima of an optimization problem. In this chapter, we introduce a general framework for two-phase algorithms for multimodal optimization, in which different high-fitness regions (niches) are located in the first phase via clustering, and each of the located niches is separately optimized with a core search algorithm in the second phase. One such two-phase algorithm is the Hill-Valley Evolutionary Algorithm (HillVallEA). In HillVallEA, the remarkably simple hill-valley clustering method is used. The idea behind hill-valley clustering is that two solutions belong to the same niche (valley) when there is no hill in between them, which can be easily tested by performing additional function evaluations. We compare hill-valley clustering to two other recently-introduced fitness-informed clustering methods: Nearest-Better Clustering and Hierarchical Gaussian Mixture Learning. We show how these clustering methods, as well as different core search algorithms, influence the resulting optimization performance of the two-phase framework on a commonly used benchmark suite. Our results show that HillVallEA, equipped with the core search algorithm Adapted Maximum-Likelihood Gaussian Model Univariate, outperforms all other algorithms for multimodal optimization, both within the limited benchmark budget, and in the long run.

4.1 Introduction

Model-based evolutionary algorithms adapt an underlying search model based on certain features of the fitness landscape [167]. Classically, these features are related to linkage, or dependence, of problem variables. The performance of many evolutionary algorithms deteriorates when the fitness landscape is multimodal, as high-quality solutions can be found in different parts of the search space, which may prevent narrowing down the search to a specific region. Being able to explicitly deal with multimodality may therefore be beneficial to many evolutionary algorithms. In addition, exploring multiple niches can provide additional insight into the structure of the problem at hand. Real-world problems are often not unimodal, and by providing the decision maker with a set of diverse high-quality solutions, a desirable solution can be chosen based on external factors that are not easily incorporated in the optimization problem and are best considered once a set of interesting alternatives is known [176].

In the field of optimization, *niching* refers to a method of obtaining and maintaining solutions in multiple niches, i.e., high-fitness parts of the search space. Niching methods originated as a tool for preserving population diversity in evolutionary algorithms, but are now generally designed for multimodal optimization [126]. Multimodal optimization is applicable to optimization problems with all types of problem variables, but is mainly applied to problems with real-valued variables, as we do in this chapter, because of a natural notion of distance and locality in real-valued fitness landscapes.

One of the difficulties of niching methods is that they often introduce additional parameters such as a minimal niche size or the (expected) number of niches [126]. How to set these parameters is however generally unknown beforehand, especially in a black-box setting. Niching methods applied in such more general optimization settings should make as few assumptions as possible on the size, shape, or number of niches.

A naive approach to multimodal optimization is restarting a *core search algorithm* randomly in different parts of the search space. In case of real-valued optimization, commonly used core search algorithms are the Covariance Matrix Self-Adaptation evolution strategy (CMSA) [22] and variants of the Adapted Maximum-likelihood Gaussian Model Iterated Density-estimation Evolutionary Algorithm (AMaLGaM) [31]. A downside of this naive serial approach to multimodal optimization is that optima within larger niches are likely to be obtained many times before a core search algorithm optimizes a smaller niche. This makes the naive approach often expensive, and it is difficult to determine when to stop restarting. A parallel—diversity-preserving—search of the search space based on crowding [50] or fitness-sharing [75] can overcome this, however, these methods do not explicitly separate solutions over multiple niches.

A recent evolutionary algorithm for multimodal optimization is the Repelling-Subpopulations CMSA (RS-CMSA) [3], in which multiple instances of the core search algorithm are randomly initialized in the search space. To prevent multiple instances from converging to the same optimum, rejection sampling is used to push them away from each other, and also from previously obtained optima.

As opposed to random initialization of core search algorithms within the search space, two-phase algorithms for multimodal optimization aim in the first phase at locating different niches in the search space. In the second phase of these algorithms, core search algorithms are initialized in the located niches. The Nearest-better Evolutionary Algorithm (NEA2+) [174, 175] is one such two-phase algorithm, which uses Nearest-Better Clustering (NBC) in the first phase to cluster an initial set of solutions. NBC is a fitness-informed clustering method that uses

a distance measure based on the concept of a nearest-better solution. The idea behind NBC is that local optima can be detected by the observation that there are no *nearby* solutions with a better fitness. However, the distinction between *nearby* and *far away* is rather difficult to make as it depends on both the number of solutions that is clustered, and the problem dimensionality [175].

In this chapter, we introduce a general framework for two-phase algorithms for multimodal optimization that can be equipped with different clustering methods and core search algorithms. Besides NBC, we consider Hierarchical Gaussian Mixture Learning (HGML) [143], which is based on the same nearest-better concept as in NBC, but the correlation between the search space and fitness values is exploited to determine a cluster set that best captures the structure of the fitness landscape. The third fitness-informed clustering method that we consider is Hill-Valley Clustering (HVC). HVC is based on the hill-valley test, which states that when there is a hill between two solutions, they belong to a different niche (valley). To test this in practice, intermediate test solutions are sampled and evaluated. The hill-valley test was first used in the multi-national evolutionary algorithm [197], where it was used to test every newly sampled solution against a population of solutions. The same test was also used by RS-CMSA, but only in a post-processing step, to determine whether two presumed optima are distinct. As core search algorithms, we consider CMSA and different variants of AMaLGaM.

The remainder of this chapter is organized as follows. In Section 4.2, we introduce a framework for two-phase algorithms for multimodal optimization. We discuss the three fitness-informed clustering-based niching methods NBC, HGML, and HVC in detail in Section 4.3. In Section 4.4, we describe the different core search algorithms. In Section 4.5, we experimentally compare the ability of the three clustering methods to locate multiple niches. We furthermore evaluate the performance of the different clustering methods and core search algorithms within the framework on the test problems of the CEC 2013 niching benchmark suite [125], and compare the best performing two-phase algorithms of the framework to other algorithms for multimodal optimization, both when allowing a smaller and a larger computational budget. We conclude in Section 4.6.

4.2 Framework for two-phase algorithms

In this chapter, we consider a general framework for two-phase algorithms for multimodal optimization (see Algorithm 4.1), that is given by,

$$\mathcal{E} = \text{TwoPhaseAlgorithm}(f, \text{ClusterMethod}, \text{CoreSearchAlgorithm}, \varepsilon), \tag{4.1}$$

where $f : X \to \mathbb{R}$ is the to-be-minimized fitness function, with $X \subseteq \mathbb{R}^d$ the d-dimensional search space. The aim of an algorithm for multimodal optization is to obtain all global optima of f, which are collected in the elitist archive $\mathcal{E} \subset X$. The user specifies the used *clustering method* and *core search algorithm*, which are discussed in respectively Section 4.3 and Section 4.4. A fundamental difficulty of real-valued multi-modal optimization is to distinguish between (high-quality) local optima and global optima. Therefore, the user specifies an acceptable fitness tolerance $\varepsilon \in \mathbb{R}_{\geq 0}$. Every solution with a fitness value less than ε worse than the best obtained fitness value is considered to be a global optimum. This tolerance can be set *large* when the aim is to also locate (high-quality) local optima.

In the first phase, an initial population $\mathcal{P} \subset X$ of $|\mathcal{P}| = N$ solutions is sampled. Truncation selection is then performed by selecting the $\lfloor \tau N \rfloor$ fittest solutions, and the set of selected solutions $\mathcal{S} \subseteq \mathcal{P}$ is clustered, i.e., $\mathcal{K} = \text{ClusterMethod}(\mathcal{S})$, with $\mathcal{K} = \{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_{K-1}\}$ being the set of K clusters $\mathcal{C}_i \subset X$. We will discuss three fitness-informed clustering methods in Section 4.3. By clustering only the selected solutions, less effort is spent on low-fitness regions in the search space. On the other hand, small niches could be accidentally discarded when the selection pressure is too high.

In the second phase, from each cluster $\mathcal{C} \in \mathcal{K}$, a core search algorithm is initialized and run sequentially until convergence. Core search algorithms return a presumed optimal solution, i.e., $\mathbf{x} = \text{CoreSearchAlgorithm}(\mathcal{C}; N_c, \rho, \mathcal{E})$. Three additional parameters need to be specified to initialize a core search algorithm. N_c is the population size used by the core search algorithm internally in the subsequent optimization of that niche. The initial search distribution multiplier $\rho \in \mathbb{R}_{>0}$ is used to initialize a cluster when it contains only one solution. In that case, the covariance matrix used by the core search algorithm is initialized by the identity matrix, multiplied by ρ. The aim is to set ρ small enough so that clusters are initialized within a single niche. In case of a box-constrained search space $X \subset \mathbb{R}^d$, when scattering N solutions equidistantly, the expected distance between two solutions is $\sqrt[d]{|X|/N}$, where $|X|$ is the volume of X. We therefore set $\rho = 0.01 \cdot \sqrt[d]{|X|/N}$, so that newly sampled solutions are expected to be closer-by than the previous set of selected solutions. Alternatively, when it is not possible to (easily) determine $|X|$, the average distance between solutions can instead be used to set ρ, as we will see in Section 4.3.3. Finally, the elitist archive $\mathcal{E}$ is provided to the core search algorithms which is used only to determine whether to terminate, as discussed in detail later in Section 4.4.1. Core search algorithms are run sequentially in the order of the fitness value of the best solution of the cluster from which they are initialized. This order prioritizes clusters that are more likely to result in a global optimum, which

Algorithm 4.1: Framework for two-phase algorithms

function: $[\mathcal{E}]$ = TwoPhaseAlgorithm(f, ClusterMethod, CoreSearchAlgorithm, ε)
input : To-be-minimized fitness function $f : X \to \mathbb{R}$, with $X \subseteq \mathbb{R}^d$
ClusterMethod; `// NBC, HGML, or HVC`
CoreSearchAlgorithm; `// See Table 4.1`
$\varepsilon \geq 0$; `// Threshold for global optima`
output : Set of presumed global optima $\mathcal{E} \subset X$;

$N = 64$; `// Initial population size`
$N_c = 0.8 \cdot N_c^{\text{rec}}$(CoreSearchAlgorithm); `// Cluster size (Table 4.1)`
$\tau = \tau^{\text{rec}}$(CoreSearchAlgorithm); `// Selection pressure (Table 4.1)`
$\mathcal{E} = \{\}, \mathcal{P} = \{\}, \mathcal{K} = \{\}$; `// Elitist archive, Population, Cluster set`

while *budget remaining* **do**

```
// First phase - locating niches
P = sample(N, P, K, d + 1, 2.0) S = truncation_selection(P, τ);
// For initializing clusters with |C| = 1
ρ = 0.01 · ∛(|X|/|S|)

// Add elites to selection
if ClusterMethod = HVC then
    S = S ∪ E
// Cluster the selection
K = ClusterMethod(S)

// Second phase - niche optimization
foreach C ∈ K do

    // Skip already-optimized niches (HVC only)
    if ClusterMethod = HVC and x_(0) ∈ C is an elite e ∈ E then
        continue

    // Run till convergence
    x = CoreSearchAlgorithm(C; N_c, ρ, E)

    // Update elitist archive
    if x is a presumed distinct global optimum then
        E = E ∪ {x}
        foreach e ∈ E that satisfies f(x) + ε < f(e) do
            E = E\{e}

// Increase initial population & cluster size
if no new solution was added to E then
    N ← 2.0 · N  N_c ← 1.1 · N_c
```

Algorithm 4.2: Hill-Valley Test [144, 197]

function: $[B]$ =Hill-Valley$(\mathbf{x}, \mathbf{y}, N_t, f)$
input : Solutions $\mathbf{x}, \mathbf{y} \in X$;
Number of test solutions $N_t \geq 0$;
To-be-minimized fitness function $f : X \to \mathbb{R}$
output : $\mathbf{x}$ and $\mathbf{y}$ belong to the same niche?

for $k = 0, \ldots, N_t - 1$ **do**
  $\mathbf{x}_{\text{test}} = \mathbf{y} + \frac{k+1}{N_t+1}(\mathbf{x} - \mathbf{y})$;
  if $\max(f(\mathbf{x}), f(\mathbf{y})) < f(\mathbf{x}_{\text{test}})$ **then**
    return false;
return true;

is important when optimization is performed with a limited budget, as is the case in this chapter.

Each run of a core search algorithm results in a presumed (locally) optimal solution $\mathbf{x}$. The user-defined fitness tolerance $\varepsilon \geq 0$ specifies whether a solution is added to the elitist archive, or whether existing solutions need to be removed from the elitist archive $\mathcal{E}$, similar to the post-processing step in [3]. Additionally, even though clusters are aimed to be initialized in different niches, it could be that specific global optima are obtained multiple times. To prevent this, without having to specify a niche-radius that determines whether two solutions are distinct, we use the hill-valley test (see Algorithm 4.2) to determine whether two presumed optima belong to different niches. In the hill-valley test, a straight line is drawn between two solutions in the search space, and the fitness is evaluated on $N_t = 5$ equidistantly located test solutions along this line. If the fitness of any of the test solutions is worse than the fitness of both solutions, there is a *hill* between the two solutions, and the two solutions presumably belong to different *valleys* (niches). As we will see in Section 4.3.3, the hill-valley test is also an essential part of HVC.

If there is budget remaining (in terms of function evaluations or computation time) after all core search algorithms are terminated, a restart is performed with a larger population size. Schemes like these that increase the population size of an evolutionary algorithm over time, i.e., with restarts, are common to overcome setting these parameters [9, 37, 87].

Here, the initial population size is set to $N = 64$ (independent of the problem dimensionality d). The initial cluster size is set to $N_c = 0.8 \cdot N_c^{\text{rec}}$, where N_c^{rec} is the recommended population size of the corresponding core search algorithm, which we will discuss in Section 4.4. Core search algorithms are initialized within a niche, which aims to make the problem landscape locally unimodal, and therefore,

a cluster size is used that is smaller than recommended. At the end of a run, only if there were no new elites obtained in that run, we increase N by a factor of 2.0 and N_c by a factor of 1.1. This is to aid with locating smaller niches and optimizing more difficult niches over time. These values were based on previously-obtained empirical results [141, 144].

4.2.1 Initial population sampling

It has been shown that better spreading the initially sampled population, compared to uniform random sampling, can improve the performance of evolutionary algorithms for multimodal optimization [205]. Here, we use a combination of greedy scattered subset selection and rejection sampling (see Algorithm 4.3). Greedy scattered subset selection aims to select a diverse set of solutions from a population [181]. It starts with the solution that has the largest parameter value in a randomly chosen dimension. Then, iteratively, the solution that is furthest away from all previously selected solutions is added, until the desired number of solutions in the subset is reached. When the two-phase algorithm is initialized, $\psi \cdot N$ solutions are sampled, which is reduced to N solutions using greedy scattered subset selection, where we use $\psi = 2$, empirically determined. After all core search algorithms have been terminated, and there is budget remaining, a new initial population is sampled. By rejection sampling, we aim to reduce the number of solutions in niches that are already explored. For this, let us denote the i-th nearest neighbor (nn_i) of a solution $\mathbf{x} \in X$ within a set of solutions $\mathcal{P} \subset X$ by,

$$\text{nn}_i(\mathbf{x}; \mathcal{P}) = \underset{\mathbf{y} \in \mathcal{P} \setminus \{\mathbf{x}\}}{\arg\min} \left(\|\mathbf{x} - \mathbf{y}\| : \mathbf{y} \notin \{\text{nn}_j(\mathbf{x}; \mathcal{P})\}_{j=0}^{i-1} \right), \tag{4.2}$$

where $\|\cdot\|$ is the Euclidean distance, and we simply write $\text{nn}(\mathbf{x}; \mathcal{P}) = \text{nn}_0(\mathbf{x}; \mathcal{P}) = \arg\min_{\mathbf{y} \in \mathcal{P} \setminus \{\mathbf{x}\}} (\|\mathbf{x} - \mathbf{y}\|)$ for the nearest neighbor of $\mathbf{x}$ within $\mathcal{P}$. We store for each solution $\mathbf{x}$ of the previous initial population $\mathcal{P}^{\text{prev}}$ to which cluster it belonged. For this, we introduce the notation $\mathcal{C}_{\langle \mathbf{x} \rangle}$ to the denote the cluster that contains $\mathbf{x}$. Then, again, $\psi \cdot N$ solutions are sampled, but now based on rejection sampling, where a sample is rejected with probability $p = 0.9$ if its nearest $N_{\min} = d + 1$ solutions of the previous initial population belonged to the same cluster, i.e., when,

$$\mathcal{C}_{\langle \text{nn}(\mathbf{x}; \mathcal{P}^{\text{prev}}) \rangle} = \mathcal{C}_{\langle \text{nn}_1(\mathbf{x}; \mathcal{P}^{\text{prev}}) \rangle} = \cdots = \mathcal{C}_{\langle \text{nn}_{N_{\min}-1}(\mathbf{x}; \mathcal{P}^{\text{prev}}) \rangle}. \tag{4.3}$$

Note that only the selection of fittest solutions are clustered, and not the entire population $\mathcal{P}^{\text{prev}}$. In case the nearest solution from the previous generation does not belong to any cluster, the newly sampled solution is always accepted, as this

Algorithm 4.3: Initial Population Sampling [141]

function: $[\mathcal{P}]$ = sample(N, $\mathcal{P}^{\text{prev}}$,$\mathcal{K}^{\text{prev}}$, $N_{\min}$, ψ)
input : Desired number of solutions $N \geq 1$;
Previous population $\mathcal{P}^{\text{prev}} \subset X \subseteq \mathbb{R}^d$;
Previous cluster set $\mathcal{K}^{\text{prev}} = \{\mathcal{C}_0^{\text{prev}}, \mathcal{C}_1^{\text{prev}}, \ldots\}$, with $\mathcal{C}_i^{\text{prev}} \subset X$;
Number of nearest neighbors $N_{\min} \geq 1$;
Sample ratio $\psi \geq 1.0$;
output : Population $\mathcal{P} \subset X$ of size $|\mathcal{P}| = N$;

```
// Sample ψ · N solutions with rejection sampling
```
for $i = 0, \ldots, \psi \cdot N - 1$ **do**
  $\mathbf{x}_i$ = sample_solution_uniformly();
  if $|\mathcal{P}^{\text{prev}}| > 0$ **then**
    if $\text{nn}(\mathbf{x}_i; \mathcal{P}^{\text{prev}}) \notin \mathcal{C}^{\text{prev}}$ **for all** $\mathcal{C}^{\text{prev}} \in \mathcal{K}^{\text{prev}}$ **then**
      continue; `// accept`
    if $\mathcal{C}^{\text{prev}}_{\langle \text{nn}(\mathbf{x}_i; \mathcal{P}^{\text{prev}}) \rangle} \neq \mathcal{C}^{\text{prev}}_{\langle \text{nn}_j(\mathbf{x}_i; \mathcal{P}^{\text{prev}}) \rangle}$ **for any** $j = 1, \ldots, N_{\min} - 1$ **then**
      if $\text{UniformRandom}_{[0,1]} < 0.9$ **then** set $i = i - 1$; `// reject`

```
// Select final N solutions with greedy scattered subset selection
```
$\mathcal{P} = \{\mathbf{x}$ with maximum parameter value in random dimension $j = 1, \ldots, d\}$;
for $i = 1, \ldots, N - 1$ **do**
  $\mathcal{P} = \mathcal{P} \cup \{\arg\max_{\mathbf{x}}(\|\mathbf{y} - \mathbf{x}\| : \mathbf{x} \notin \mathcal{P}, \mathbf{y} \in \mathcal{P})\}$; `// add furthest solution`
return $\mathcal{P} = \{\mathbf{x}_i, \ldots, \mathbf{x}_N\}$;

means that a high-fitness solution was obtained in a region that was considered of low fitness in the previous generation, which was therefore not explored.

For only the resulting N solutions, the fitness value is computed. Note that this is however a relatively expensive approach with a computational complexity of $\mathcal{O}(N^2 d)$, and when computation time is limited, standard uniform sampling may be preferred at the cost of a slight performance reduction.

4.3 Fitness-informed clustering

We now discuss three fitness-informed clustering methods that aim to cluster a set of solutions, such that each cluster resides in a single niche. These clustering methods are fitness-informed in the sense that they are not only based on distances in the search space, but also incorporate corresponding fitness values.

4.3.1 Nearest-better clustering

In Nearest-Better Clustering (NBC) [174–176] (see Algorithm 4.4), a spanning tree of solutions called the *nearest-better tree* is constructed. In this three, each solution

is connected to the nearest solution in the search space that has better fitness, i.e., to its *nearest-better solution*. We can formally define the nearest better $\text{nb}(\mathbf{x})$ solution of a solution $\mathbf{x} \in X$ within a set of solutions $\mathcal{S} \subset X$ by first defining,

$$\mathcal{S}^+(\mathbf{x}) = \{\mathbf{y} \in \mathcal{S} \backslash \{\mathbf{x}\} : f(\mathbf{y}) \leq f(\mathbf{x})\}, \tag{4.4}$$

as the subset of $\mathcal{S}$ of all solutions with equal or better fitness than $\mathbf{x}$. Then, let,

$$\begin{aligned} \text{nb}(\mathbf{x}) &= \text{nn}(\mathbf{x}; \mathcal{S}^+(\mathbf{x})) = \underset{\mathbf{y} \in \mathcal{S}^+(\mathbf{x})}{\arg\min}(\|\mathbf{x} - \mathbf{y}\|), \\ \delta(\mathbf{x}) &= \|\mathbf{x} - \text{nb}(\mathbf{x})\|, \end{aligned} \tag{4.5}$$

where we refer to $\delta(\mathbf{x})$ as the nearest-better distance of solution $\mathbf{x}$ (within $\mathcal{S}$). Ranking the solutions based on fitness allows for an efficient construction of the nearest-better tree. Denote the solution with rank i by $\mathbf{x}_{(i)}$, with $\mathbf{x}_{(0)}$ being the best solution, then $\mathcal{S}^+(\mathbf{x}_{(i)}) = \{\mathbf{x}_{(0)}, \ldots, \mathbf{x}_{(i-1)}\}$. Note that the best solution in the population, $\mathbf{x}_{(0)}$, has no nearest-better solution.

The rationale behind the nearest-better tree is that the best solutions found so far in the niche of a local optimum have no nearby solutions with better fitness, and therefore have a relatively long outgoing edge, i.e., a large nearest-better distance. By removing these edges, a number of disconnected sub-trees remains, and each sub-tree forms a cluster. To detect long edges, we compute the mean edge length,

$$\mu_\delta = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|-1} \delta(\mathbf{x}_{(i)}), \tag{4.6}$$

and remove edges with length $\delta(\mathbf{x}) \geq \phi \cdot \mu_\delta$. It is not straightforward to calibrate these cutting rules, as distances between solutions depend both on the problem dimensionality and the number of solutions [176]. We set $\phi = 2.0$ in this work, adhering to literature on NBC [174]. In terms of computational complexity, generating the NBC can be performed in $\mathcal{O}(|\mathcal{S}|^2 d)$.

4.3.2 Hierarchical Gaussian mixture learning

In Hierarchical Gaussian Mixture Learning (HGML) [143] (see Algorithm 4.5), solutions are clustered such that, when a Gaussian Mixture Model (GMM) is fitted to the clusters, it correlates well with the fitness landscape. In that way, when sampling from the GMM, high-fitness solutions are expected. A GMM was used because in the second phase, niches will be explored with core search algorithms that are also Gaussian-based.

Algorithm 4.4: Nearest-Better Clustering (NBC) [174]

function: $[\mathcal{K}] = \text{NBC}(\mathcal{S})$
input : Set of solutions $\mathcal{S} \subset X$;
output : Set of clusters $\mathcal{K} = \{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_{K-1}\}$, with $\mathcal{C}_i \subset X$;

Rank solutions $\mathbf{x} \in \mathcal{S}$ by fitness value, such that $\mathbf{x}_{(0)}$ is the best solution;
$\mathcal{C}_0 := \{\mathbf{x}_{(0)}\}$; $K = 1$;
$\phi = 2.0$;
$\mu_\delta = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|-1} \delta(\mathbf{x}_{(i)})$; `// Note that` $\delta(\mathbf{x}_{(0)})$ `does not exist`
for $\mathbf{x} \in \mathcal{S}, \mathbf{x} \neq \mathbf{x}_{(0)}$ **do**
  if $\delta(\mathbf{x}) < \phi \cdot \mu_\delta$ **then**
    Add $\mathbf{x}$ to cluster $\mathcal{C}_{\langle \text{nb}(\mathbf{x}) \rangle}$ of $\text{nb}(\mathbf{x})$;
  else
    New cluster $\mathcal{C}_K := \{\mathbf{x}\}$; $K = K + 1$;
return $\mathcal{K} = \{\mathcal{C}_0, \ldots, \mathcal{C}_{K-1}\}$;

Gaussian mixture model

Let $\mathcal{N}_\theta$ be a d-dimensional Gaussian distribution, parameterized by $\theta = (\boldsymbol{\mu}, \Sigma)$, where $\boldsymbol{\mu} \in \mathbb{R}^d$ is the distribution mean and $\Sigma \in \mathbb{R}^{d \times d}$ the covariance matrix. We denote its probability distribution function for $\mathbf{x} \in \mathbb{R}^d$ by $p(\mathbf{x}; \theta)$. To fit a Gaussian distribution to a set of data points, one could use the maximum likelihood estimator, which yields a closed form solution [81, 89].

A natural extension of the Gaussian distribution is the GMM, given by,

$$p_K(\mathbf{x}; \Theta) = \sum_{k=0}^{K-1} w_k p(\mathbf{x}; \theta_k), \tag{4.7}$$

with K mixture components, where $\Theta = \{(w_k, \theta_k)\}_{k=0,\ldots,K-1}$ is the set of distribution parameters and w_k are the positive mixing weights, summing up to one. For a GMM, there is no closed form maximum likelihood estimator, but estimates can for instance be found via expectation-maximization [81].

Expectation-maximization is however a computationally expensive algorithm, even for a fixed number of components K. It is computationally cheaper to first cluster the data into K clusters. As each solution then belongs to only one cluster, we can subsequently use the maximum likelihood estimator to estimate a mixture component on each of the clusters, and set the weights inversely proportional to the number of solutions in each cluster. By using this approach, a set of clusters can be directly and uniquely associated with a GMM. Due to this simplification, overlapping mixture components are no longer possible, which is desirable for niching approaches, as niches are also non-overlapping.

Algorithm 4.5: Hierarchical Gaussian Mixture Learning (HGML)

function: $[\mathcal{K}] = \text{HGML}(\mathcal{S})$
input : Set of solutions $\mathcal{S} \subset X$;
output : Set of clusters $\mathcal{K} = \{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_{K-1}\}$, with $\mathcal{C}_i \subset X$;

Rank solutions $\mathbf{x} \in \mathcal{S}$ by fitness value, such that $\mathbf{x}_{(0)}$ is the best solution;
$\mathcal{C}_0 = \{\mathbf{x}_{(0)}\}$;
$\mathbf{E} = \{\}$; `// List of edges of the nearest-better tree`

for $i = 1, \ldots, |\mathcal{S}| - 1$ **do**
  $\mathcal{C}_i = \{\mathbf{x}_{(i)}\}$;
  $\mathbf{E} = \mathbf{E} \cup \{(\mathcal{C}_i, \mathcal{C}_{\langle \text{nb}(\mathbf{x}_{(i)}) \rangle})\}$; `// Edge from` $\mathcal{C}_i$ `to cluster of nearest-better`

$\mathcal{K}_0 := \{\mathcal{C}_0, \ldots, \mathcal{C}_{|\mathcal{S}|-1}\}$;

for $n = 1, \ldots, N - 1$ **do**
  Find the shortest edge $(\mathcal{C}_f, \mathcal{C}_t) \in \mathbf{E}$;
  Create new cluster by merging $\mathcal{C} = \mathcal{C}_f \cup \mathcal{C}_t$;
  Initialize new cluster set $\mathcal{K}_n := \mathcal{K}_{n-1} \backslash \{\mathcal{C}_f, \mathcal{C}_t\} \cup \mathcal{C}$;
  Delete edge $(\mathcal{C}_f, \mathcal{C}_t)$ from $\mathbf{E}$;
  Replace all occurrences of $\mathcal{C}_f$ and $\mathcal{C}_t$ in $\mathbf{E}$ by the new $\mathcal{C}$;

$n^+ = \arg\max_n (\text{Round}(\text{DFC}(\mathcal{K}_n), 0.05))$;
if $\text{DFC}(\mathcal{K}_{n^+}) < 0$ **then**
  $n^+ = N - 1$;

return $\mathcal{K} = \mathcal{K}_{n^+}$;

A fundamental difficulty of mixture models is how to determine the number of mixture components K, as more mixture components will inherently result in a better fit of the data [89]. In our case, rather than fitting the data points, the primary target is to fit a GMM that maximizes the probability of sampling high-fitness solutions. We will use this idea to determine a suitable number of mixture components.

Density-fitness correlation

The Density-Fitness Correlation (DFC) was introduced to determine whether the probability distribution of an estimation-of-distribution algorithm needed to be adapted for a better match with the fitness landscape [30]. Here, we use the DFC to test, in a greedy way, which GMM is the best match with the structure of the problem at the current state of the search.

Given a GMM, parameterized by Θ, as in Equation (4.7), and a set of solutions $\mathcal{S} \subset X$, we compute the DFC as follows. For each solution $\mathbf{x} \in \mathcal{S}$, let $D(\mathbf{x}; \Theta, \mathcal{S}) = \text{Rank}(p_K(\mathbf{x}; \Theta) | \mathcal{S})$ be the density rank of $\mathbf{x}$ under the GMM parameterized by Θ, such that the solution with the highest probability density has rank 0. Furthermore, let $F(\mathbf{x}; \mathcal{S}) = \text{Rank}(f(\mathbf{x}) | \mathcal{S})$ be the fitness rank of $\mathbf{x}$, such that the best solution

has rank 0. Then, the DFC is given by the Spearman rank correlation between the density and fitness ranks,

$$\mathrm{DFC}(\Theta, \mathcal{S}) = 1 - \frac{6 \sum_{i=0}^{|\mathcal{S}|-1} \left[F(\mathbf{x}_i; \mathcal{S}) - D(\mathbf{x}_i; \Theta, \mathcal{S})\right]^2}{|\mathcal{S}|(|\mathcal{S}|^2 - 1)}, \tag{4.8}$$

and takes values in $[-1, 1]$. The larger the DFC, the higher the probability that high-fitness solutions are sampled. In our case, as the solutions are clustered first, the cost of computing the DFC can be reduced by computing the DFC per cluster, i.e., per Gaussian $\mathcal{N}_{\theta_k}$. The DFC of the GMM is then the average of the clusters DFC values, weighted by the number of solutions in each cluster, i.e., for a cluster set $\mathcal{K} = \{\mathcal{C}_k\}_{k=0}^{K-1}$, with $\theta_k = (\boldsymbol{\mu}_k, \Sigma_k)$ corresponding to cluster $\mathcal{C}_k$, we obtain,

$$\mathrm{DFC}(\mathcal{K}) = \sum_{k=0}^{K-1} \frac{1}{|\mathcal{C}_k|} \mathrm{DFC}(\theta_k, \mathcal{C}_k). \tag{4.9}$$

Using the DFC, we can select the best GMM out of a set of candidate GMMs.

Hierarchical clustering

To generate a set of candidate GMMs, a bottom-up hierarchical clustering approach is used where, initially, each solution is considered to be a separate cluster. Then, iteratively, clusters are merged until one large cluster remains that contains all solutions. After each merge, a GMM is fitted to the current cluster set, and the corresponding DFC is computed using Equation (4.9). Clustering $|\mathcal{S}|$ solutions results in $|\mathcal{S}|$ cluster sets, each corresponding to a GMM. The cluster set with the best DFC is selected as final cluster set.

The merge order is however important. Solely merging based on distance in the search space might result in solutions of two different niches being merged into the same cluster only because they are close in the search space, while they may in fact belong to different niches. To reduce this effect, we make use of the nearest-better tree that we discussed in Section 4.3.1. Edge lengths in the nearest-better tree were only defined as the (Euclidean) distance between two solutions. We naturally extend this definition by defining the distance between two clusters as the Euclidean distance between the cluster means in the search space. We then iteratively merge the two clusters corresponding to the shortest edge in the nearest-better tree. We recall from Section 4.3.1 solutions with long outgoing edges are expected to be (local) optima, and by iteratively merging the shortest edge in the tree, these solutions get merged latest.

Note that at least $N_{\min} = d + 1$ solutions are required for each cluster in the cluster set to estimate a stable $d \times d$ covariance matrix. We therefore ignore all cluster sets that contain a cluster with less than $N_{\min}$ solutions, by setting the corresponding DFC to -1, i.e., the worst attainable value.

GMM estimates are prone to statistical noise, especially in smaller clusters, thus if two clusters have a similar DFC, we prefer the larger cluster. To realize this, the DFCs are rounded with an accuracy of 0.05, which was empirically determined [143]. If two rounded DFCs have the same value, the larger cluster is chosen. If all the DFCs are negative, which implies that all of the estimated GMMs correlate poorly with the problem structure, we fall back to using a single cluster that contains all solutions. When using a lower bound of $N_{\min} = d + 1$ solutions per cluster, HGML has a computational complexity of $\mathcal{O}(|\mathcal{S}|d^2 + |\mathcal{S}|^2 d)$ [143].

4.3.3 Hill-Valley clustering

The final clustering method that we discuss in this chapter is Hill-Valley Clustering (HVC) [141, 144] (see Algorithm 4.6). As discussed previously, in NBC, solutions were clustered together with their nearest-better solution, as long as these are *nearby*. It might however be that even nearby solutions belong to a different niche. In HVC, we explicitly check for this using the hill-valley test (see Algorithm 4.2), which is also used to determine whether two obtained elites are distinct global optima. In HVC, if a solution does not belong to its nearest-better solution, it is checked whether it belongs to its next-nearest better. For this, we define the i-th nearest-better solution of a solution $\mathbf{x} \in X$ within a set of solutions $\mathcal{S} \subset X$ as $\text{nb}_i(\mathbf{x}) = \text{nn}_i(\mathbf{x}; \mathcal{S}^+(\mathbf{x}))$, with $\mathcal{S}^+(\mathbf{x})$ defined as in Equation (4.4).

To cluster a set of solutions $\mathcal{S} \subset X$, HVC is initialized by forming the first cluster $\mathcal{C}_0 = \{\mathbf{x}_{(0)}\}$ from the best solution $\mathbf{x}_{(0)} \in \mathcal{S}$. Then, we consider the second-best solution $\mathbf{x}_{(1)} \in \mathcal{S}$, and test whether it belongs to the same niche as its nearest-better solution, $\text{nb}(\mathbf{x}_{(1)}) = \mathbf{x}_{(0)}$, using the hill-valley test. When it does, $\mathbf{x}_{(1)}$ is added to the cluster of $\mathbf{x}_{(0)}$. Otherwise, as there are no other better solutions, a new cluster is formed, i.e., $\mathcal{C}_1 = \{\mathbf{x}_{(1)}\}$. Next, the third-best solution $\mathbf{x}_{(2)} \in \mathcal{S}$ is tested against the nearest solution that has better fitness, which can either be $\mathbf{x}_{(0)}$ or $\mathbf{x}_{(1)}$, depending on which one is nearer. If $\mathbf{x}_{(2)}$ does not belong to the same niche as its nearest-better solution $\text{nb}(\mathbf{x}_{(2)})$, we test it against $\text{nb}_1(\mathbf{x}_{(2)})$. If it also does not belong to that niche, we create a new cluster from $\mathbf{x}_{(2)}$. For each solution, we check at most its $N_{\min} = d + 1$ nearest-better solutions. In this way, even for the one-dimensional problems, at least two neighboring solutions are considered. This procedure is then repeated for all solutions.

Algorithm 4.6: Hill-Valley Clustering (HVC) [144]

function: $[\mathcal{K}] = \text{HVC}(\mathcal{S}, f)$
input : Set of solutions $\mathcal{S} \subset X$;
To-be-minimized fitness function $f : X \to \mathbb{R}$, with $X \subseteq \mathbb{R}^d$
output : Set of clusters $\mathcal{K} = \{\mathcal{C}_0, \mathcal{C}_1, \ldots, \mathcal{C}_{K-1}\}$, with $\mathcal{C}_i \subset X$;

Rank solutions $\mathbf{x} \in \mathcal{S}$ by fitness value, such that $\mathbf{x}_{(0)}$ is the best solution;
$\mathcal{C}_0 := \{\mathbf{x}_{(0)}\}$; $K = 1$;
for $\mathbf{x} \in \mathcal{S}$ **do**
  for $j = 0, \ldots, d-1$ **do**
    $\mathbf{y} = \text{nb}_j(\mathbf{x})$; `// Break if there are no more better solutions`
    if *already checked cluster of* $\mathbf{y}$ **then**
      continue;
    $N_t = 1 + \left\lfloor \|\mathbf{x} - \mathbf{y}\| / \sqrt[d]{|X|/|\mathcal{S}|} \right\rfloor$; `// Or` $N_t = 1 + \lfloor \|\mathbf{x} - \mathbf{y}\| / \mu_\delta \rfloor$
    if Hill-Valley$(\mathbf{x}, \mathbf{y}, N_t, f)$ **or** $(N_t = 1$ **and** Rank$(f(\mathbf{x})|\mathcal{S}) \geq 0.5 \cdot |\mathcal{S}|)$ **then**
      Add $\mathbf{x}$ to cluster of $\mathbf{y}$;
      break;
  if $\mathbf{x}$ *was not added to any cluster* **then**
    $\mathcal{C}_K := \{\mathbf{x}\}$; $K = K + 1$;
return $\mathcal{K} = \{\mathcal{C}_0, \ldots, \mathcal{C}_{K-1}\}$

Number of test solutions N_t

In the post-processing step to filter out global optima that are obtained multiple times, the hill-valley test was used with $N_t = 5$ test solutions. We would like to use fewer test solutions during the clustering process however, but to have more test solutions when two solutions are further apart. We therefore increase the number of test solutions with the distance between the two solutions, scaled by expected distance between two solutions in the selection, similar to how ρ was set in Section 4.2, i.e.,

$$N_t = 1 + \left\lfloor \frac{\|\mathbf{x} - \mathbf{y}\|}{\sqrt[d]{|X|/|\mathcal{S}|}} \right\rfloor . \tag{4.10}$$

When the search space is unbounded, or its volume $|X|$ cannot be determined easily, we can replace the expected distance $\sqrt[d]{|X|/|\mathcal{S}|}$ by the mean distance between nearest-better solutions μ_δ as given in Equation (4.6), i.e., by setting $N_t = 1 + \lfloor \|\mathbf{x} - \mathbf{y}\| / \mu_\delta \rfloor$.

Reducing function evaluations within HVC

Two extensions are added to HVC to reduce the number of function evaluations. First, if the nearest-better solution and a next-nearest-better solution belong to the

same cluster, we only perform the hill-valley test on the nearest-better solution, and if it is rejected, we reject the next-nearest-better solution without performing the hill-valley test.

Second, it is generally not required to accurately model the low-fitness search space. We therefore automatically cluster two solutions $\mathbf{x}, \mathbf{y} \in \mathbb{R}$, with $f(\mathbf{y}) \leq f(\mathbf{x})$ together when they are close to each other and both belong to the worst 50% of the set $\mathcal{S} \subset X$ that is being clustered. This is the case when $N_t = 1$ and $\text{Rank}(f(\mathbf{x})|\mathcal{S}) \geq 0.5 \cdot |\mathcal{S}|$, without performing the hill-valley test.

Additionally, when the hill-valley test has been performed, and when two solutions are clustered together, all test solutions generated by the hill-valley test are added to the cluster as well. When the hill-valley test rejected a solution pair, the test solutions that have been evaluated so far are added to the cluster of the worst, which is the endpoint of the line from which sampling starts. However, the solution that violated the hill-valley test, and thus has worse fitness than the two solutions, is discarded.

Preventing re-optimization of niches

To prevent niches from being re-explored after a restart of the framework, all elites are added back to the population before clustering. After clustering, when the best solution in a cluster is an elite, no core search algorithm is initialized from that cluster. This approach can however only be applied when HVC is used as clustering method, as HVC tests whether solutions actually belong to the same niche as an elite, while the other clustering methods simply cluster them to any nearby solution.

4.4 Core search algorithms

We consider the core search algorithms listed in Table 4.1, which are CMSA [22] and different versions of AMaLGaM [31]. AMaLGaM is an estimation of distribution algorithm, where a Gaussian distribution is fitted to selected solutions with maximum likelihood and adapted based on whether improvements are found near or far from the mean. In AMaLGaM, a full covariance matrix is estimated, while a univariate covariance matrix is estimated in AMaLGaM-univariate. The latter is therefore cheaper in terms of memory and computational effort, as only the diagonal of the covariance matrix needs to be computed and stored. This is especially relevant for problems with a larger problem dimensionality d. More relevant in our case is however that estimating an accurate full-rank covariance matrix requires a larger population size, which implies that fewer generations of AMaLGaM can

Table 4.1: Different core search algorithms with corresponding recommended settings as taken from literature: selection pressure $\tau \in (0, 1]$, population size N_c^{rec} for problem dimensionality d.

Core search algorithm	Abbreviation	τ^{rec}	N_c^{rec}
CMSA [22]	CMSA	0.50	$3 \log d$
AMaLGaM [31]	AM	0.35	$17 + 3 \cdot d\sqrt{d}$
AMaLGaM-univariate [31]	AMu	0.35	$10\sqrt{d}$
iAMaLGaM [31]	iAM	0.35	$10\sqrt{d}$
iAMaLGaM-Univariate[31]	iAMu	0.35	$4\sqrt{d}$

be run compared to AMaLGaM-univariate for the same computational budget. In iAMaLGaM(-univariate), the covariance matrix is estimated incrementally over multiple generations, which allows for a smaller population size.

CMSA was used instead of the more common CMA-ES [86], as it was suggested to perform better when adding elitism [3]. CMSA fits a Gaussian using the population mean of the previous generation, essentially modeling the direction of improvements.

4.4.1 Termination criteria for core search algorithms

As core search algorithms are run in serial, additional termination criteria are required besides the budget defined in terms of function evaluations. AMaLGaM-variant core search algorithms are terminated if the maximum standard deviation of solutions in the search space is too low (10^{-12}), or if the standard deviation of fitness values is too low (10^{-12}) [31]. CMSA is terminated using the recommended criteria [86], and parameters set as in RS-CMSA [3], that is, if the improvement in fitness value over the last $10 + \lfloor 30d/N_c \rfloor$ generations is less than 10^{-5}. Furthermore, fail-safe termination criteria are added to terminate CMSA if the standard deviation of solutions in the search space reaches machine accuracy (10^{-15}) and if the condition number of the covariance matrix is larger than 10^{14}.

Termination when converging to elite

In an attempt to prevent core search algorithms from re-exploring niches, in every fifth generation of a core search algorithm the best obtained solution is compared to the nearest-better solution in the elitist archive using the hill-valley test (with $N_t = 5$ test solutions). When it is found that these two solutions belong to the same niche, the core search algorithm is terminated.

Termination when converging to local optimum

The final stopping criterion that we introduce is used to detect whether a core search algorithm is converging to a local minimum. As a proxy for the global minimum, let b be the fitness value of the best solution in the elitist archive. Let a_g be the average fitness of the selected solutions in generation g and $\Delta_g := a_g - (b+\varepsilon)$ the fitness difference with the best elite, incorporating the user-defined tolerance ε of whether a solution can be considered a global optimum. If $\Delta_g \leq 0$, the current selection is better than the so-far obtained elites, and we do not terminate.

To predict the value of Δ_{g+1}, we use the fact that AMaLGaM was shown to have exponential convergence on smooth unimodal functions such as the sphere function [31]. Under this assumption, we can write $\Delta_{g+1} = \Delta_g(1-r)$, where r is the rate of convergence. We approximate r by r_n over the previous n generations by,

$$r \approx r_n = 1 - \left(1 - \frac{\Delta_{g-n} - \Delta_g}{\Delta_{g-n}}\right)^{1/n}, \tag{4.11}$$

with $n = 5$ to reduce statistical noise. To prevent premature termination, this criterion is only applied when Δ_g decreased in the most recent $n = 5$ generations. Finally, we estimate the *time to optimum* (t_o) in order to achieve an accuracy of $\Delta_{g+t_o} = \Delta_g(1-r)^{t_o} = 10^{-12}$, again under the assumption of exponential convergence. Rewriting this in terms of t_o gives,

$$t_o = \frac{\log\left(10^{-12}/\Delta_g\right)}{\log\left(1-r\right)} \approx \frac{\log\left(10^{-12}/\Delta_g\right)}{\log\left(1-r_n\right)} = \frac{\log\left(10^{-12}/\Delta_g\right)}{\frac{1}{n}\log\left(\Delta_g/\Delta_{g-n}\right)}. \tag{4.12}$$

The idea is that t_o increases rapidly when converging to a local optimum. Therefore, a core search algorithm is terminated if $g + t_o$ exceeds 50 times the maximum number of generations it took to find any elite in the elitist archive.

4.5 Experiments

We compare the discussed methods on the CEC 2013 niching benchmark suite [125]. This benchmark consists of 20 problems, to be solved within a specified budget in terms of function evaluations (see Table 4.2). All benchmark problems are defined on a box-constrained domain. For each of the benchmark problems, the location of the optima and the corresponding fitness values are known. These are only used to measure performance, and not during optimization. We use the benchmark guidelines to determine whether a global optimum is attained [125], at an accuracy level of $\varepsilon = 10^{-5}$. All experiments in this chapter are run with this tolerance level.

Table 4.2: CEC 2013 niching benchmark suite [125]. For each problem, the function name, problem dimensionality d, number of global optima $\#gopt$, local optima $\#lopt$, and budget in terms of function evaluations are given.

P	Function name	d	#gopt	#lopt	Budget
1	Five-Uneven-Peak Trap	1	2	3	50K
2	Equal Maxima	1	5	0	50K
3	Uneven Decreasing Maxima	1	1	4	50K
4	Himmelblau	2	4	0	50K
5	Six-Hump Camel Back	2	2	5	50K
6	Shubert	2	18	many	200K
7	Vincent	2	36	0	200K
8	Shubert	3	81	many	400K
9	Vincent	3	216	0	400K
10	Modified Rastrigin	2	12	0	200K
11	Composition Function 1	2	6	many	200K
12	Composition Function 2	2	8	many	200K
13	Composition Function 3	2	6	many	200K
14	Composition Function 3	3	6	many	400K
15	Composition Function 4	3	8	many	400K
16	Composition Function 3	5	6	many	400K
17	Composition Function 4	5	8	many	400K
18	Composition Function 3	10	6	many	400K
19	Composition Function 4	10	8	many	400K
20	Composition Function 4	20	8	many	400K

We can then make use of a key performance measure in multimodal optimization, the peak ratio (PR), which is defined as the fraction of obtained distinct global optima over the total number of global optima for the problem at hand. Besides the peak ratio, a secondary performance measure that is commonly used is defined as the fraction of obtained *distinct* global optima within the provided solution set. This measure aims to overcome a scenario in which an algorithm simply returns all obtained solutions, without giving insight whether solutions actually belong to different modes. In our two-phase framework, the elitist archive is post-processed to remove local optima and duplicate global optima, in a similar fashion as in RS-CMSA. Because of the post-processing step, the perfect score for this measure is always achieved, and it is therefore left out of this chapter.

Unless mentioned otherwise, all experiments are repeated 50 times, and resulting performance measures are averaged over all repetitions. Results are tested for statistical significance using the one-sided unpaired Wilcoxon rank-sum test at $\alpha = 0.05$. Additionally, the Bonferoni correction is applied by dividing α by the number of tests performed in each experiment. Source code of HillVallEA and the two-phased framework is available online [138].

4.5.1 Clustering comparison

We investigate how many unique niches (corresponding to global optima) can be located when clustering a population of a given population size with the clustering methods in Section 4.3, without performing subsequent optimization of the located niches. We define a niche to be located when the best solution of a cluster is in that niche. For smooth functions (i.e., benchmark problems P1–P10), this is easily tested for with the hill-valley test as the locations of their optima are known. The niche ratio (NR) is then given by the number of located niches as a fraction of the total number of niches of the problem. Furthermore, we define the success ratio (SR) as the number of different niches located as a fraction of the number of clusters. Both NR and SR should be maximized and when $NR = SR = 1$, all niches are located exactly once.

We compare the clustering methods on three problems with $d = 2$ variables that exhibit different landscape features. The Modified Rastrigin function (P10) has no local optima and all niches are the same shape and size. The Vincent function (P7) also has only global optima, but of different shape and size, of which a few very small. Finally, the Shubert function (P6) has many local optima and global optima with small niches.

All methods cluster the same set of initial solutions, and when the population size is increased, more solutions are added instead of re-sampling the entire population, to reduce randomness. The extra function evaluations used by HVC are included in the results. Additionally, we show the improvement in performance for HVC when using a better spread of the sampled population as described in Section 4.2.1.

Results

Results are shown in Figure 4.1. Overall, the Modified Rastrigin function is a rather easy problem, as it can be clustered with only few solutions, and both HGML and HVC obtain the perfect score, i.e., $NR = SR = 1$. NBC also obtains a perfect score for a population size of around 10^3 solutions, but as the population size increases, it tends to overestimate the number of clusters. This same behaviour is observed for all three test problems. This is because of the cutting rule, as described in Section 4.3.1, that has a fixed threshold. The results clearly indicate that it should be adapted to the population size, as this influences the distribution of distances between neighboring solutions.

HVC achieves the best NR for the three test problems. The additional function evaluations are included. This shows that it is worth spending additional evaluations during clustering in order to better locate niches versus simply

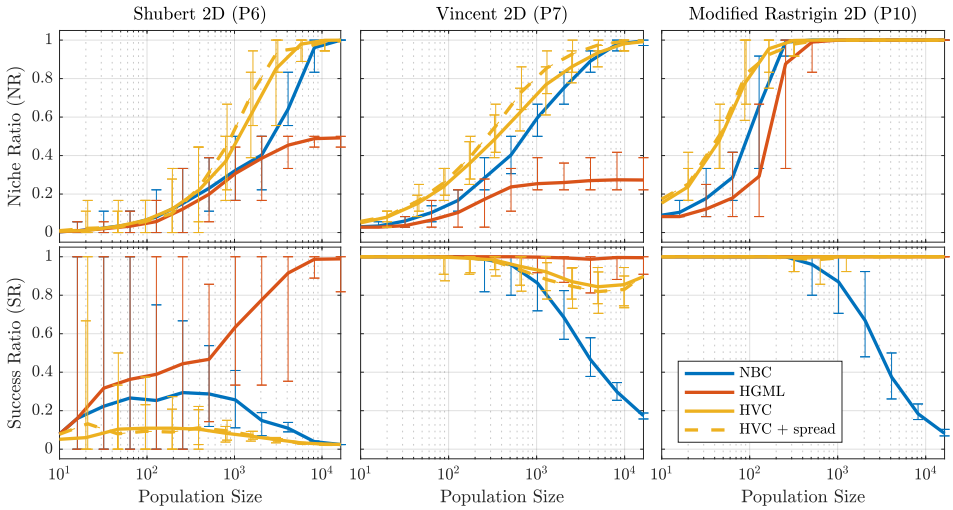

Figure 4.1: Niche Ratio and Success Ratio for the three clustering methods discussed in Section 4.3, for different population size. Error bars represent min/max scores. *HVC + spread* is HVC applied on a better spread of the initial solutions in the population as described in Section 4.2.1.

increasing the population size. Additionally, better spreading the initial sampled population shows to further improve HVC.

A drawback is that HVC aims to locate all niches, including the ones corresponding to local optima. This can be observed for the Shubert function, for which HVC obtains a large number of clusters. The Vincent and Modified Rastrigin problem have no local optima, therefore, the SR of HVC is closer to 1.0 for these problems. For the Vincent function, the SR deteriorates slightly when the population size increases. This happens because HVC only checks a limited number of nearby solutions ($N_{\min} = d + 1 = 3$ solutions). For the very stretched niches of the Vincent function, and when the population size is not too large, it might be that all those solutions belong to a different niche. Therefore, multiple clusters are (incorrectly) initialized in these stretched niches. Checking more nearby solutions by increasing $N_{\min}$ reduces this effect, at the cost of additional function evaluations.

HGML is not able to distinguish peaks that are close to each other for both the Shubert and Vincent problem. This is a result of the hierarchical clustering process that is based on the nearest-better distances. For the Shubert function, peaks are divided over nine groups of two optima. HGML obtains all 9 groups correctly, but, as the two peaks in each group are very close to each other, a single Gaussian fits rather almost as good to these two peaks as to two separate Gaussians would. Additionally, the hierarchical clustering order is based on nearest-better distances.

Table 4.3: Peak ratios obtained by equipping the framework for two-phase algorithms with different core search algorithms and clustering methods, averaged over all 20 benchmark problems.

Clustering	AM	AMu	CMSA	iAM	iAMu
HVC	0.869	**0.892**	0.855	0.887	0.748
NBC	0.857	0.873	0.868	0.873	0.798
HGML	0.702	0.715	0.727	0.720	0.707

Since these two peaks are very close to each other, these solutions get often wrongly merged into the same cluster early in the merging process. This is also the reason why NBC performs worse on the Shubert function than on the Vincent function compared to HVC. On the other hand, HGML performs best in terms of the SR, as it obtains $SR = 1$ when the population size is large enough, as each Gaussian it fits accurately captures one (or multiple) modes. To improve the niche ratio of HGML, clustering can be applied every generation, as in [143], and not in a two-phase manner ad done in this work.

4.5.2 Core search algorithms and clustering methods

We equip the two-phase framework with each of the five core search algorithms in Table 4.1 and each of the three clustering methods of Section 4.3. The resulting 15 different two-phase algorithms are run on the 20 benchmark problems.

Results

Table 4.3 shows the obtained average peak ratios for the 15 different two-phase algorithms. We refer to the algorithms equipped with HVC as the Hill-Valley Evolutionary Algorithm (HillVallEA). The best average peak ratio (0.89) was obtained by HillVallEA with AMu as core search algorithm (HillVallEA-AMu).

HVC performed generally best, closely followed by NBC, and with quite an improvement over HGML. This is in line with the results obtained in the previous experiment. In terms of core search algorithms, AMu and iAM have the same recommended population size, which allows for the same number of restarts, and the resulting performance is very similar in all cases. A reason for the bad performance of iAMu could be that the used cluster size, $N_c = 0.8 \cdot N_C^{\text{rec}}$, is too small, and therefore deteriorates performance. CMSA sometimes *jumps* out of its own niche shortly after initialization. Because of that, even when small niches are located in the clustering process, they might get lost during optimization, resulting in a lower peak ratio when using HVC or NBC. HGML however finds fewer clusters

Table 4.4: Peak ratios obtained by equipping the framework for two-phase algorithms with AMu and the three clustering methods, per problem, and the average (avg) over all 20 problems. All scores are averaged over 50 runs. Bold scores are the best obtained per problem, and those not significantly different from it. For problems 1–6 and 10, all instances achieves a perfect score of 1, and are therefore not shown.

Problem	AMu+HVC	AMu+NBC	AMu+HGML
7	**1.00**	**1.00**	0.67
8	0.97	**0.98**	0.92
9	**0.97**	0.71	0.37
11	**1.00**	**1.00**	0.95
12	**1.00**	**1.00**	0.90
13	**1.00**	**1.00**	0.64
14	0.92	**0.98**	0.78
15	**0.75**	**0.75**	**0.73**
16	**0.72**	**0.77**	**0.53**
17	**0.75**	**0.75**	**0.42**
18	**0.67**	**0.67**	0.19
19	**0.59**	**0.61**	**0.48**
20	**0.48**	0.24	0.12
avg.	**0.892**	0.873	0.715

upon initialization, and jumping out of the located niche seems therefore beneficial, resulting in the best performance of HGML when combining it with CMSA.

In Table 4.4, the peak ratio per problem obtained with two-phase algorithms equipped with AMu is shown. For problems 8 and 14, NBC performs slightly better than HVC. However, for problems 9 and 20, HVC greatly outperforms the other two clustering methods. The restart scheme largely overcame that HGML was unable to locate all peaks for problems 6 and 7, but this approach still resulted in a lower peak ratio for almost all problems.

4.5.3 Algorithm comparison

We compare HillVallEA-AMu, i.e., the two-phase algorithm equipped with HVC and AMu, to other algorithms that performed well in the niching competitions held at the GECCO and CEC conferences. The included algorithms are, NEA2+ [175], RLSIS [205], and RS-CMSA [2].

Results

Table 4.5 shows the peak ratio per problem. HillVallEA-AMu obtains the best peak ratios for all problems, followed by RS-CMSA, NEA2+, and finally RLSIS. Problems 1–5 and 10 are fully solved by all methods (except for two runs of

Table 4.5: Peak ratios per problem (P) for different algorithms. Higher scores are better and 1 is the maximum achievable score. Mean peak ratios over 50 runs are shown, together with the standard deviation. Average (avg.) computed over all 20 problems. Bold scores are the best obtained per problem, and those not significantly different from it.

P	HillVallEA-AMu	RS-CMSA	NEA2+	RLSIS
1	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000
2	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000
3	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000
4	**1.000** ±0.000	**1.000** ±0.000	**0.990** ±0.049	**1.000** ±0.000
5	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000	**1.000** ±0.000
6	**1.000** ±0.000	**0.999** ±0.008	**0.991** ±0.023	0.872 ±0.078
7	**1.000** ±0.000	**0.997** ±0.008	0.810 ±0.046	0.920 ±0.025
8	**0.975** ±0.016	0.871 ±0.032	0.567 ±0.045	0.189 ±0.039
9	**0.972** ±0.011	0.730 ±0.018	0.539 ±0.014	0.584 ±0.016
10	**1.000** ±0.000	**1.000** ±0.000	**0.988** ±0.033	**1.000** ±0.000
11	**1.000** ±0.000	**0.997** ±0.023	0.943 ±0.086	**1.000** ±0.000
12	**1.000** ±0.000	0.948 ±0.062	0.785 ±0.094	0.950 ±0.066
13	**1.000** ±0.000	**0.997** ±0.023	0.937 ±0.094	0.793 ±0.103
14	**0.923** ±0.090	0.810 ±0.082	0.810 ±0.075	0.703 ±0.069
15	**0.750** ±0.000	**0.748** ±0.017	0.718 ±0.060	0.720 ±0.059
16	**0.723** ±0.079	0.667 ±0.000	**0.683** ±0.050	0.670 ±0.023
17	**0.750** ±0.000	0.703 ±0.066	0.723 ±0.057	**0.738** ±0.037
18	**0.667** ±0.000	**0.667** ±0.000	**0.650** ±0.050	**0.667** ±0.000
19	**0.593** ±0.078	0.502 ±0.017	0.505 ±0.075	0.515 ±0.041
20	**0.480** ±0.046	**0.482** ±0.043	0.380 ±0.093	0.422 ±0.075
avg.	**0.892** ±0.030	0.856 ±0.024	0.801 ±0.032	0.787 ±0.032

NEA2+). These problems are too simple to give insight in the performance of different algorithms. None of the algorithms can obtain the final two Weierstrass peaks of Composition Function 4 for any dimensionality (problems 15, 17, 19, and 20) within the given computational budget. Specifically, for problems 15, 17, and 18, the standard deviation of the obtained peak ratio for HillVallEA-AMu is zero. This suggests that the final peaks cannot be obtained even when run with a larger budget, which is something we investigate in the next experiment. Again, note that, due to the post-processing step in HillVallEA-AMu and RS-CMSA, both algorithms only obtain distinct global optima, and the number of obtained solutions is therefore equal to the number of located optima.

4.5.4 Larger budget

To show and study maximum performance, we compare HillVallEA-AMu with the best competitor, RS-CMSA, on a computational budget 100 times larger than the original budget. Due to computational limits, this experiment was repeated only

20 times. We consider the problems on which HillVallEA-AMu was not able to obtain a peak ratio of 1 in the previous experiment, i.e., the Shubert and Vincent problems in 3D, and CF3 and CF4 with different problem dimensionalities.

Results

The result for the larger budget is shown in Figure 4.2. HillVallEA-AMu outperforms RS-CMSA in all cases by achieving a higher peak ratio for the same number of evaluations. For problem 8 and 9, both methods obtain all peaks when the computational budget is large enough, which suggests that both algorithms are fundamentally sound. The difference between HillVallEA-AMu and RS-CMSA is most clear for the Vincent function (P9), which has a number of very small niches that are difficult to locate, and lacks local optima. Due to its design, HillVallEA-AMu performs especially well in these two scenarios.

For CF3, the peak ratio curve shows that even for $d = 2$, two of the peaks (corresponding to the Weierstrass function) are difficult to obtain. This effect is enhanced to the extent that for problem 18, with $d = 10$, these two peaks are never obtained, even for the case of the extended computational budget. For CF4, the two peaks of the Weierstrass functions are never obtained, not even for $d = 2$. This suggests that these optima might be a needle in a haystack, and therefore fundamentally virtually unobtainable. For increasing problem dimensionality, the two Griewank optima furthermore also become harder to obtain.

4.6 Conclusion

We introduced a framework for two-phased algorithms for multimodal optimization, that equips a clustering method to locate niches in the first phase, and a core search algorithm in the second phase to optimize each niche locally. We showed that hill-valley clustering outperforms other fitness-informed clustering methods by locating more niches when allowing the same number of function evaluations. The best performing instance of the framework for two-phase algorithms is obtained by equipping the framework with hill-valley clustering, which we refer to as the Hill-Valley Evolutionary Algorithm (HillVallEA). HillVallEA in particular outperforms all other algorithms on a commonly used niching benchmark suite when combined with the core search algorithm AMaLGaM-Univariate (AMu), both within the limited benchmark budget and when allowing many more function evaluations.

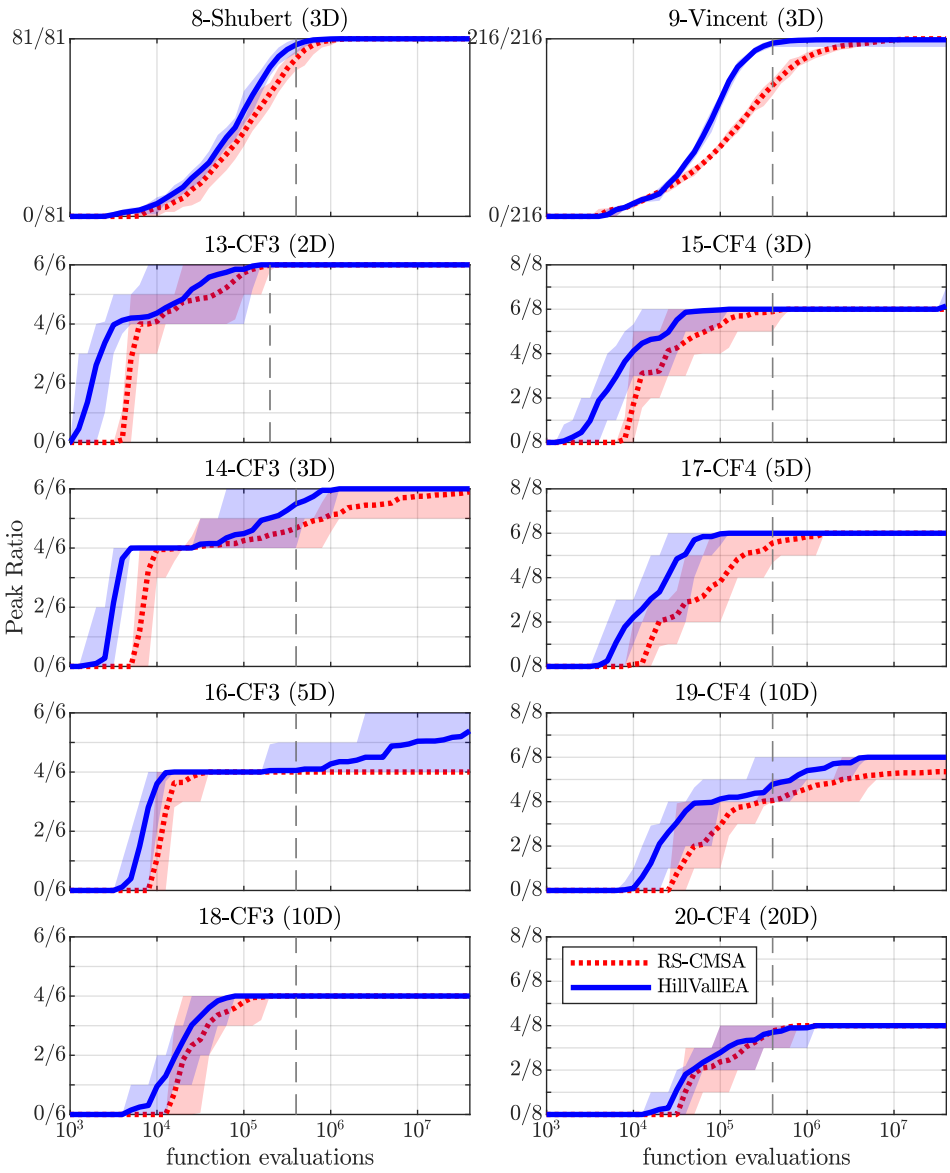

Figure 4.2: Peak ratio shown as a function of the number of function evaluations, for RS-CMSA and HillVallEA-AMu. Shaded areas represent min-max range over multiple repetitions of the experiment. Problems for which HillVallEA did not obtain a peak ratio of 1.0 within the original budget (dashed vertical line) are shown. The allowed budget in terms of function evaluations was 100 times the original budget. Note that the problems are sorted so that all variants of CF3 are in the left column and all variants of CF4 in the right column.

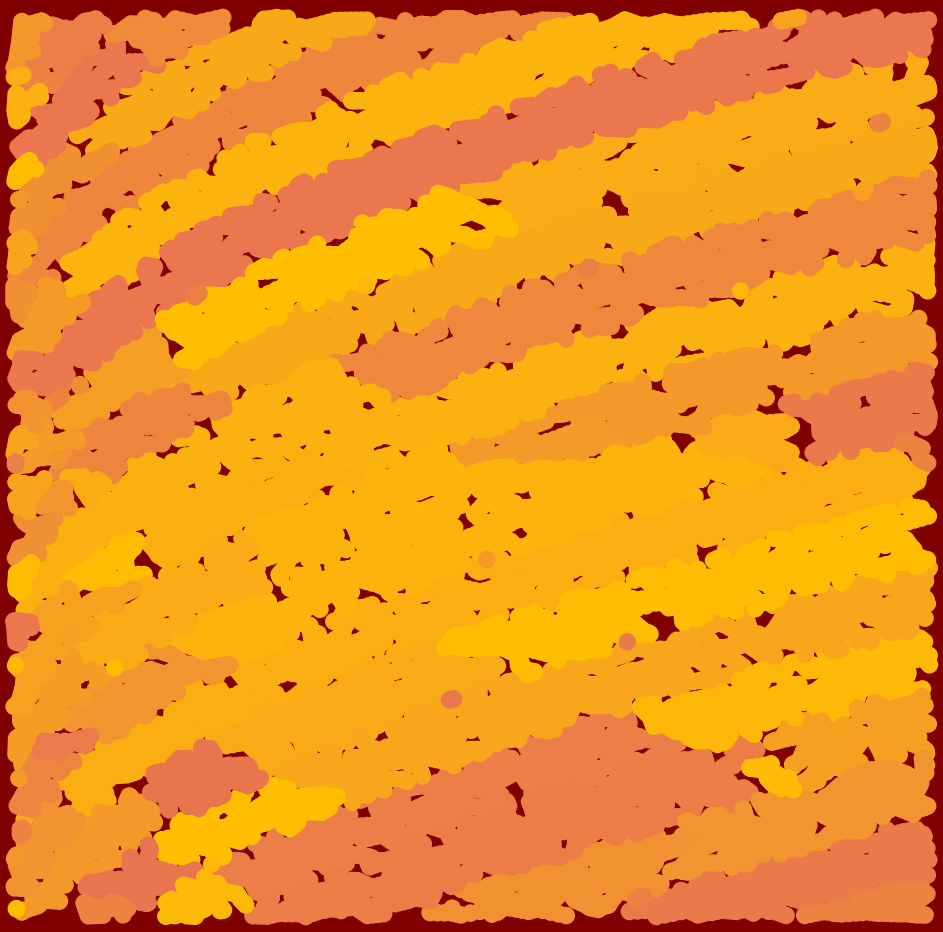

5

Multimodal multi-objective optimization by hill-valley clustering

Research question 4:

How can niching via hill-valley clustering be applied to multi-objective optimization problems to allow for efficient and effective multi-objective multimodal optimization, without making assumptions on the number of modes, or on their shape or size?

This chapter is an adaptation of S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2019) *Real-valued evolutionary multi-modal multi-objective optimization by hill-valley clustering.* In Proceedings of the Genetic and Evolutionary Computation Conference, 568–576, ACM Press, New York, New York.

Abstract

In model-based evolutionary algorithms, the underlying search distribution is adapted to the problem at hand, for example based on dependencies between decision variables. Hill-valley clustering is an adaptive niching method in which a set of solutions is clustered such that each cluster corresponds to a single mode in the fitness landscape. This can be used to adapt the search distribution of an evolutionary algorithm to the number of modes, exploring each mode separately. Especially in a black-box setting, where the number of modes is a priori unknown, an adaptive approach is essential for good performance. In this work, we introduce multi-objective hill-valley clustering and combine it with MAMaLGaM, a multi-objective evolutionary algorithm, into the Multi-Objective Hill-Valley Evolutionary Algorithm (MO-HillVallEA). We empirically show that MO-HillVallEA outperforms MAMaLGaM and other well-known multi-objective optimization algorithms on a set of benchmark functions. Perhaps most important, we show that MO-HillVallEA is capable of obtaining and maintaining multiple approximation sets simultaneously over time.

5.1 Introduction

A multi-objective optimization problem comprises two or more objective functions that need to be optimized simultaneously. When these objectives are conflicting, instead of a single optimal solution, multiple *Pareto-optimal* solutions exist. Without further information, none of these solutions is better than any other. The aim of multi-objective optimization is to obtain as many *diverse* Pareto-optimal solutions as possible, to be presented to the user for decision making. Multi-objective evolutionary algorithms (MOEAs) are aimed to find a set of Pareto-optimal solutions with different trade-offs in the objectives [27, 53, 212]. Distinct solutions that have similar objective values often get lost during optimization, as there is no mechanism to maintain all of these [27, 53]. Maintaining these could however provide useful insight for decision making and improve performance.

The aim of multimodal multi-objective optimization is to obtain a good approximation of the set of *all* Pareto-optimal solutions [126, 192, 211]. In other words, while multi-objective optimization is aimed to approximate the *Pareto front* in objective space, multimodal multi-objective optimization is aimed to approximate the *Pareto set* in decision space.

Recently, a number of MOEAs for multimodal optimization (MMOEAs) have been introduced, by applying niching techniques to existing MOEAs. DN-NSGA-II is a niched NSGA-II using crowding, which shows better decision-space diversity compared to standard NSGA-II [127]. MO_Ring_PSO_SCD is a niched particle swarm optimizer using an index-based ring topology [211]. MOEA/D-AD is a niched MOEA/D, which was shown to outperform MO_Ring_PSO_SCD in terms of decision-space diversity [192].

The contribution of this chapter is twofold. First, we introduce a novel niching technique, hill-valley clustering, for multi-objective optimization in Section 5.2. Hill-valley clustering, previously introduced for single-objective optimization in Chapter 4, is an adaptive clustering method in which a set of solutions is clustered into niches by spending additional function evaluations to determine whether two solutions belong to the same niche. Second, in Section 5.3, we combine hill-valley clustering with the MOEA MAMaLGaM [181] into an MMOEA, which we refer to as MO-HillVallEA. In Section 5.4, MO-HillVallEA is benchmarked against MAMaLGaM and other (M)MOEAs on different (multimodal) multi-objective optimization problems. We discuss the results of this chapter in Section 5.5, and conclude in Section 5.6.

5.2 Multi-objective niching

We define multimodal multi-objective optimization by a to-be-minimized function $\mathbf{f} : \mathcal{X} \to \mathbb{R}^m$, where $\mathbf{f} = [f_0, \ldots, f_{m-1}]$ is an m-dimensional objective function and where $\mathcal{X}$ is the n-dimensional decision space $\mathcal{X} \subseteq \mathbb{R}^n$. A solution $\mathbf{x} \in \mathcal{X}$ is said to *dominate* another solution $\mathbf{y} \in \mathcal{X}$, when $\mathbf{x}$ is better than $\mathbf{y}$ in at least one objective, and not worse in the others. The set of all non-dominated solutions is called the *Pareto set*. The image of the Pareto set under $\mathbf{f}$ is called the *Pareto front*. The aim of multi-objective optimization is to find an approximation set $\mathcal{A}$ of solutions that approximates the Pareto front, while the aim of multimodal multi-objective optimization is to approximate the Pareto set.

In optimization, a *niche* is a subset of the decision space where only one mode resides. In this chapter, we consider *multi-objective niching* as the partitioning of the decision space into the minimum number of niches required so that all objectives are uni-modal within a niche. Each niche then contains one *local Pareto set*. The Pareto set is a subset of the union of all local Pareto sets. When a local Pareto set maps to (part of) the global Pareto front, we refer to it as a *global Pareto set*. A formal definition of local Pareto sets, or locally efficient sets, is given in [107].

Let us demonstrate the concept of multi-objective niching using the minimum distance (MinDist) problem [27, 103], which is based on multiple objective functions of the form $f_i(\mathbf{x}|\mathbf{c}_0, \mathbf{c}_1) = \min\{\|\mathbf{x} - \mathbf{c}_0\|, \|\mathbf{x} - \mathbf{c}_1\|\}$, for center points $\mathbf{c}_0, \mathbf{c}_1 \in \mathcal{X}$. We configure the MinDist problem with two objectives as follows,

$$\begin{aligned} \mathbf{f} &= [f_0(\mathbf{c}_0, \mathbf{c}_1), f_1(\mathbf{c}_2, \mathbf{c}_3)], \\ \mathbf{c}_0 &= [-2, -1, 0, \ldots, 0], \quad \mathbf{c}_1 = [2, 1, 0, \ldots, 0], \\ \mathbf{c}_2 &= [-2, 1, 0, \ldots, 0], \quad \mathbf{c}_3 = [2, -1, 0, \ldots, 0]. \end{aligned} \tag{5.1}$$

This instance of the MinDist problem is visualized in Figure 5.1. The overlapping niches in the two objectives result in four multi-objective niches. Two local Pareto sets map to the global Pareto front, while the other two map to local Pareto fronts. A local Pareto set can be understood as a curve in the decision space, connecting optima of different objectives. However, parts of these the curves resulting in the local Pareto sets are dominated by the global Pareto sets. Note that, even though both objective functions have no local optima, the multi-objective problem has both local and global Pareto sets.

5.2.1 Single-objective hill-valley clustering

Hill-valley clustering is a niching approach that was introduced for single-objective multimodal optimization in Chapter 4. It can be used to cluster a set of solutions such that each cluster resides in a single niche. To determine whether two solutions belong to the same niche, the *hill-valley test* is used [197], of which pseudo code is given in Algorithm 5.1. The idea behind this test is that, when there is a *hill* between two solutions, they belong to different *valleys* (niches). For this, a number of solutions are evaluated along the line segment connecting the two solutions in decision space. If any of these test solutions has worse fitness than the two solutions, the two solutions belong to different niches.

To cluster a population of solutions, an approach inspired by nearest-better clustering [174] is used. First, the population is sorted on fitness, starting with the best solution, which forms the first cluster. Then, iteratively, each next-best solution is tested to determine whether it belongs to the niche of the nearest solution that has better fitness. When a solution-pair is rejected, the next-nearest solution with better fitness is considered. It could be that the nearest solution belongs to a different niche, but a next-nearest solution that is located slightly further, but in a different direction, belongs to the same niche. If a solution does not belong to the same niche as its $N_n = n + 1$ nearest-better neighbors, a new cluster is formed from that solution. This is repeated until the entire population is clustered.

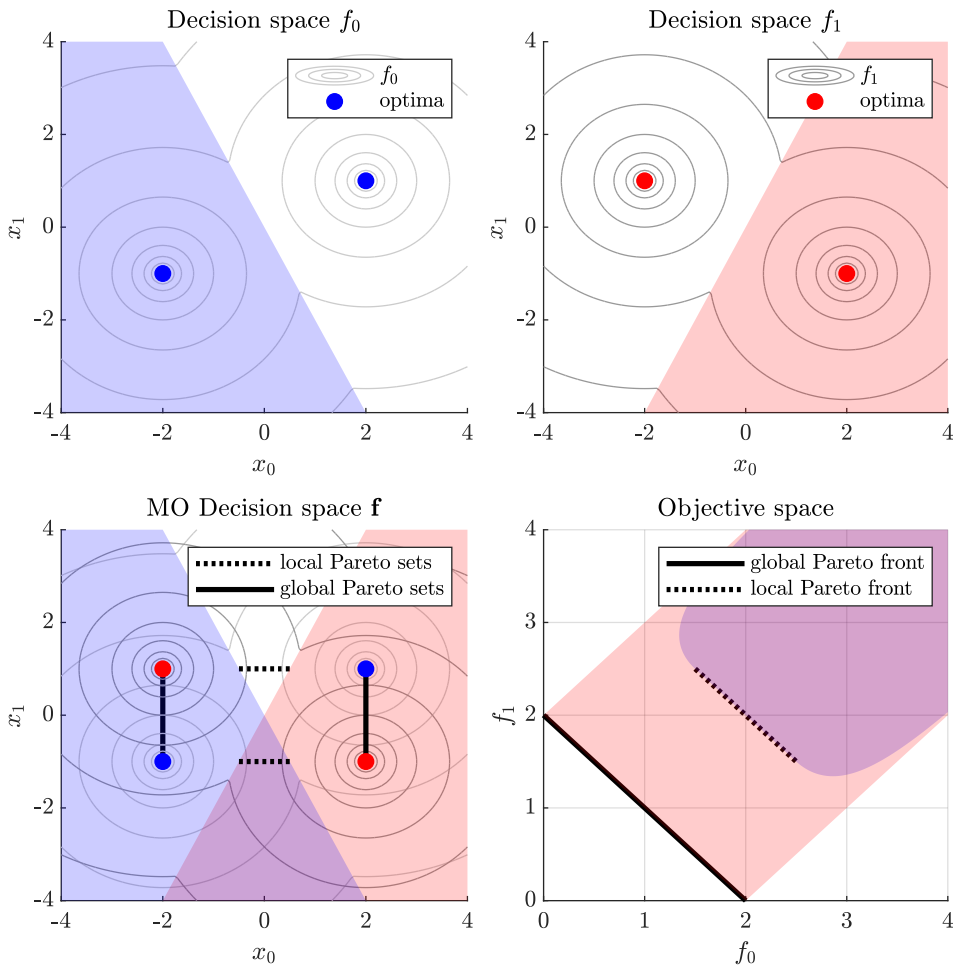

Figure 5.1: Niches for the bi-objective MinDist problem. The top subfigures show the single-objective niches for both objectives in different colors. The lower left subfigure show the four niches for the multi-objective problem $\mathbf{f} = [f_0, f_1]$, together with the four local Pareto sets. The lower right subfigure shows objective space (the blue niche overlaps exactly with red niche, and white with purple).

5.2.2 Multi-objective hill-valley clustering

We extend hill-valley clustering to the multi-objective case using the insight gained from Figure 5.1. Single-objective hill-valley clustering is performed once for each objective. The result is m (different) clusterings of the population. These are then reduced to a single cluster set by taking intersections, similar as the differently colored regions illustrated in Figure 5.1. That is, solutions that belong to the same niches for all objectives are grouped together. See pseudo code in Algorithm 5.2.

Algorithm 5.1: Hill-valley Test

function: [boolean] = HillValleyTest($\mathbf{x}, \mathbf{y}, N_t, f$)
input : Solutions $\mathbf{x}, \mathbf{y}$, Number of test points N_t,
Objective function f (to be minimized)
output : $\mathbf{x}, \mathbf{y}$ belong to the same niche? (boolean)

for $k = 1, \ldots, N_t$ **do**
  $\mathbf{x}_k = \mathbf{x} + \frac{k}{N_t+1}(\mathbf{y} - \mathbf{x})$;
  if $\max\{f(\mathbf{x}), f(\mathbf{y})\} < f(\mathbf{x}_k)$ **then**
    return false;
return true;

To improve efficiency, test solutions evaluated by the hill-valley test are stored to prevent evaluating them twice for the clustering of different objectives. Furthermore, test solutions for solution-pairs that ended up in the same cluster are added to that cluster as well. Test solutions for solution-pairs that ended up in different clusters are discarded, as it is not clear to which cluster they should be added without further testing.

Hill-valley clustering uses at least N function evaluations to cluster a population of size N, as at least one neighbor needs to be checked for each solution. An upper bound is more difficult to formulate because the number solutions used in the hill-valley test depends on the distance between nearest better solutions, which is problem- and sample-dependent. As at most $n + 1$ neighbors are considered per solution, the number of function evaluations is roughly $\mathcal{O}((n + 1)N)$ in practice.

5.3 MO-HillVallEA

We use hill-valley clustering to construct a multimodal multi-objective evolutionary algorithm that we refer to as the Multi-Objective Hill-Valley Evolutionary Algorithm (MO-HillVallEA). The idea behind the algorithm is that every generation, the population is clustered using hill-valley clustering. Then, each cluster is explored with a generation of a core optimization algorithm. For this we use different versions of MAMaLGaM. MAMaLGaM is an estimation-of-distribution algorithm, i.e., a type of model-based evolutionary algorithm that learns a probability distribution to subsequently sample new solutions from [181].

The population $\mathcal{P}$ is initialized by uniform sampling, after which hill-valley clustering is used to obtain clusters $\mathbb{C} = \{\mathcal{C}_0, \mathcal{C}_1, \ldots\}$. For each cluster $\mathcal{C}_i$, a generation of the core optimization algorithm is performed to generate offspring $\mathcal{O}_i$.

Algorithm 5.2: Hill-valley Clustering

function: $\mathbb{C} = \text{HillValleyClustering}(\mathcal{P}, \mathbf{f})$
input : Population of solutions $\mathcal{P}$,
Objective function $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^m$
output : Set of clusters $\mathbb{C} = \{\mathcal{C}_0, \ldots, \mathcal{C}_{s-1}\}$

$\Delta = \sqrt[n]{V_{\mathcal{P}}/|\mathcal{P}|}$; // $V_{\mathcal{P}}$ = volume of bounding box of $\mathcal{P} \subset \mathbb{R}^n$
for $l = 0, \ldots, m-1$ **do**
  // Single-objective hill-valley clustering for objective l
  Sort $\mathbf{x} \in \mathcal{P}$ on fitness value of l-th objective, fittest first;
  $\mathcal{C} = \{\mathbf{x}_0\}$; $\mathbb{K}_l = \{\mathcal{C}\}$; // Init cluster from best solution
  for $i = 1, \ldots, |\mathcal{P}| - 1$ **do**
    // Compute Euclidean distances to all better $\mathbf{x}_p$
    $\delta_p = \|\mathbf{x}_i - \mathbf{x}_p\|$, for $p = 0, \ldots, i-1$;
    for $j = 0, \ldots, \min\{i-1, n\}$ **do**
      k = index of j-th nearest better solution according to $\{\delta_p\}_{p=0}^{i-1}$;
      if *HillValleyTest*$(\mathbf{x}_k, \mathbf{x}_i, 1 + \lfloor \delta_k/\Delta \rfloor, f_l)$ **then**
        $\mathcal{C}_{(\mathbf{x}_k)} = \mathcal{C}_{(\mathbf{x}_k)} \cup \{\mathbf{x}_i\}$; // Add to cluster of $\mathbf{x}_k$
        break;
    if $\mathbf{x}_i$ *was not added to any cluster* **then**
      $\mathcal{C} := \{\mathbf{x}_i\}$; $\mathbb{K}_l = \mathbb{K}_l \cup \mathcal{C}$; // Init new cluster
$\mathbb{C} = \bigcap\{\mathbb{K}_l\}_{l=0}^{m-1}$; // all intersections of all clusters in $\mathbb{K}_l$

Core optimization algorithms of the model-building type, such as MAMaLGaM, generally estimate model parameters over the course of multiple generations. Therefore, a set of model parameters ρ_i is maintained for each cluster $\mathcal{C}_i$. To smoothly transfer model parameters over generations, clusters are linked to the nearest cluster from the previous generation. This distance is measured by the Euclidean distance between the cluster means in decision space. Note that this supports having a variable number of clusters every generation.

Besides the main population, an elitist archive $\mathbb{E}$ is maintained that contains the best solutions over time. To allow for multimodal optimization, all non-dominated solutions within a niche need to be maintained, even if they are dominated by a solution from a different niche. Therefore, a local elitist archive, or subarchive, $\mathcal{E}_i$ is generated from each cluster $\mathcal{C}_i$ by maintaining only the non-dominated solutions within that cluster. The (global) elitist archive $\mathbb{E}$ is then the set of all subarchives, i.e., $\mathbb{E} = \{\mathcal{E}_0, \mathcal{E}_1, \ldots\}$. Note that in this way, dominated solutions can end up in the elitist archive, but only if the dominating solution is from a different subarchive, and thus presumably from a different niche. In order to incorporate subarchives from

the previous generation into the clusters of the next generation, one must know to which cluster each subarchive corresponds. Therefore, all elites are added to the population before hill-valley clustering is applied. In that way, clusters contain all new offspring and all elites. After the clustering process and the construction of the elitist archive, the elites are removed from the clusters, but maintained in the elitist archives. Note that MAMaLGaM, within core_opt_generation($\mathcal{C}, \mathcal{E}, \rho$), does again add a few elites back in the clusters to improve convergence.

To reduce the number of function evaluations spent during hill-valley clustering, whenever two elites are considered in the hill-valley test that originate from the same subarchive, they are said to be part of the same niche, without further testing. However, when two elites originate from different subarchives, the hill-valley test is performed. Due to the discrete nature of the hill-valley test, solutions close to the boundaries of a niche can be clustered incorrectly, resulting in small, low-fitness clusters. By testing elites from different niches every generation, these clusters are more often merged with the correct neighboring clusters.

If, at the end of the generation, the elitist archive size exceeds a user-defined target size $N_{\mathbb{E}}$, adaptive objective-space discretization is performed, to reduce computational cost while maintaining diversity in the archive [134, 181]. All subarchives are discretized using the same grid size, which is adapted until the total archive size is less than the given target size $N_{\mathbb{E}}$. Note that this maintains diversity within subarchives, but does not focus on total diversity over all archives. The main generational loop for MO-HillVallEA is inspired by MAMaLGaM, and pseudo code is given in Algorithm 5.3.

5.3.1 MAMaLGaM

We briefly describe a generation of MAMaLGaM, and necessary adaptations required for its use as a core optimization algorithm.

MAMaLGaM is initialized with a cluster $\mathcal{C}_i \in \mathbb{C}$ as its population, from which offspring are generated. For this, $\mathcal{C}_i$ is clustered once more. To prevent confusion with hill-valley clustering, we refer to these subclusters within MAMaLGaM as *subsets*. MAMaLGaM combines two clustering approaches to obtain a predefined number of subsets k. First, for each of the m objectives, a subset is formed by the N_c best solutions in that objective. Second, the population is sorted based on domination rank, and balanced k-leader-means clustering [181] is used to cluster the τN best-ranked solutions into the remaining $k-m$ subsets, where τ is the selection fraction. This results in a set of k overlapping subsets. A Gaussian distribution is then estimated for each subset, from which new offspring are sampled. For this, mechanisms from different evolutionary algorithms have been used previously,

Algorithm 5.3: MO-HillVallEA

function: $[\mathbb{E}] = \text{MO-HillVallEA}(\mathbf{f}, N, N_{\mathbb{E}})$
input : Objective function $\mathbf{f}$, Population size N, Elitist archive target size $N_{\mathbb{E}}$
output : Elitist archive $\mathbb{E} = \{\mathcal{E}_0, \mathcal{E}_1, \ldots\}$

$\mathcal{P} = \text{UniformSampling}(N)$; `// Also evaluates solutions`
$\mathbb{C} = \text{HillValleyClustering}(\mathcal{P}, \mathbf{f})$; `// See Algorithm 5.2`
$\mathbb{E} = \text{ConstructLocalElitistArchives}(\mathbb{C}, N_{\mathbb{E}})$;
$\rho = \text{InitModelParameters}(\mathbb{C})$; `// For each` $\mathcal{C}_i$`, modelparams` ρ_i

while *budget remaining* **do**
  `// Copy all elites to the population`
  $\mathcal{P} = \bigcup_i \{\mathcal{E}_i\}$;
  `// Run a generation of each core optimizer`
  for $(\mathcal{C}_i, \mathcal{E}_i, \rho_i) \in (\mathbb{C}, \mathbb{E}, \rho)$ **do**
    $(\mathcal{O}_i, \rho_i) = \text{core\_opt\_generation}(\mathcal{C}_i, \mathcal{E}_i, \rho_i)$;
    $\mathcal{P} = \mathcal{P} \cup \mathcal{O}_i$; `// Collect offspring` $\mathcal{O}_i$
  $\mathbb{C}_{\text{prev}} = \mathbb{C}$; `// Backup old clusters`
  $\mathbb{C} = \text{HillValleyClustering}(\mathcal{P}, \mathbf{f})$; `// See Algorithm 5.2`
  $\mathbb{E} = \text{ConstructLocalElitistArchives}(\mathbb{C}, N_{\mathbb{E}})$;
  $\mathbb{C} = \text{RemoveElitesFrom}(\mathbb{C})$
  $(\mathbb{C}, \rho) = \text{LinkClusters}(\mathbb{C}, \mathbb{C}_{\text{prev}}, \rho)$;

such as, CMA-ES, iAMaLGaM, and AMaLGaM [181]. Subsets are registered to subsets from the previous generation and a one-to-one subset registration is applied. Besides sampling new solutions every generation, a maximum of τN elites is put back into the population, selected from the elitist archive based on objective-space diversity.

To initialize MAMaLGaM, the offspring size needs to be set, for which we use $N_i = \lfloor N/|\mathbb{C}| \rfloor$, i.e., we sample an equal number of offspring for each cluster, with the purpose of better maintaining smaller niches. Furthermore, we fix the total number of subsets k over all clusters within MO-HillVallEA. For an instance of MAMaLGaM initialized with $\mathcal{C}_i \in \mathbb{C}$ as population, we set the number of subsets k_i by $k_i = \lfloor k \cdot |\mathcal{C}_i| / \sum_j |\mathcal{C}_j| \rfloor$, with $k_i \geq 1$. Only when $k_i > m$, m out of k_i subsets are formed by performing single-objective selection. Because k_i can now vary per generation, the one-to-one subset registration that was previously used is no longer applicable. It is therefore replaced by simply registering each subset to the nearest subset from the linked cluster of the previous generation. As subsets are formed based on objective-space distances, subset distances are computed by the distance of the subset means in objective space as well. Additionally, single-objective subsets are linked with each other if they existed in both generations. This approach is significantly faster than the original one-to-one subset registration.

5.3.2 Post-processing the approximation set

In practice, a user can often process only a limited number of solutions during decision making, and the size of the resulting approximation set $\mathcal{A}$ that an algorithm obtains must thus be limited. Directly restricting the size of the elitist archive $\mathbb{E}$ during optimization has a risk of deteriorating performance, especially when the desired number of solutions is small. Therefore, a post-processing step is applied when the desired approximation set size $N_{\mathcal{A}}$ is smaller than the elitist archive target size $N_{\mathbb{E}}$. In that case, the approximation set is formed by combining all local elitist archives that contain at least one solution that is non-dominated within the global elitist archive. If the approximation set still exceeds $N_{\mathcal{A}}$, a greedy scattered subset selection is performed to reduce the archive size while preserving decision-space diversity as good as possible. This is the same subset selection algorithm as was used for balanced k-leader-means clustering [181]. A similar post-processing step was used for MOEA/D-AD in [192], with the difference that MO-HillVallEA can maintain dominated solutions as long as they are from different subarchives.

5.3.3 Multi-start scheme

To overcome the need for parameter tuning of the population size N, a multi-start scheme is applied, similar to [38], where multiple instances of MO-HillVallEA are run simultaneously, each with a larger population size. Recursively, after 8 generations of an instance of MO-HillVallEA with population size N, one generation of an instance with population size $2N$ is performed. The first instance of MO-HillVallEA is initialized with population size $N_{\text{base}} = 10 \cdot (1+m) \cdot (1+\log n)$ and $k_{\text{base}} = 1+m$ subsets, where m is the number of objectives and n the decision-space dimensionality. Each subsequent instance has a population size increased by a factor 2 and number of subsets increased by a factor 1.5. The number of subsets increases slower than the population size, so that the subset size increases over time. A single global elitist archive $\mathbb{E}$ is maintained for all instances. Whenever a new instance is initialized, it is counted how many solutions the smaller instances contributed to the elitist archive in their most recent generation. An instance that is responsible for less than 25% of the total contribution is terminated, except for the latest instance.

5.4 Experiments

We empirically benchmark MO-HillVallEA against comparable (M)MOEAs. For this, MAMaLGaM was implemented in C++, and MO-HillVallEA was implemented

in the same framework to reduce behavioral differences due to implementation choices. Results of other (M)MOEAs were provided by the authors of [192]. Source code of MO-HillVallEA in C++ is available online [139].

5.4.1 Test problems

We consider the MinDist problem with $m = 2$ objectives described in Section 5.2 with $n = \{10, 20\}$ decision variables, initialized in the box $[-4, 4]^n$. Additionally, we consider the MinDist problem with $m = 3$ objectives, given by,

$$\begin{aligned} \mathbf{f} &= [f_0(\mathbf{c}_0, \mathbf{c}_1), f_1(\mathbf{c}_2, \mathbf{c}_3), f_2(\mathbf{c}_4, \mathbf{c}_5)], \\ \text{with, } \mathbf{c}_0 &= [-4, -4, 0], \mathbf{c}_1 = [2, 2, 0], \mathbf{c}_2 = [-2, -4, 0], \\ \mathbf{c}_3 &= [4, 2, 0], \mathbf{c}_4 = [-3, -2, 1], \mathbf{c}_5 = [3, 4, 1]. \end{aligned} \tag{5.2}$$

This configuration results in two triangle-shaped global Pareto sets. Similarly for the $m = 2$ MinDist problem, zeros are appended to the centers if the number of decision parameters is larger.

Furthermore, we consider six bi-objective multimodal test problems with $n = 2$ decision variables that are frequently considered in multimodal literature: Two-On-One [177], SYM-PART{1,2,3} [184], and SSUF{1,3} [211]. The Two-On-One problem and the SSUF1 problem both have two symmetrical global Pareto sets that are connected in decision space and that map onto the same Pareto front in objective space. The SYM-PART problems have nine global Pareto sets, where SYM-PART2 has rotated global Pareto sets and SYM-PART3 has non-linear global Pareto sets. The SSUF3 problem has two global Pareto sets that are shifted, and many local Pareto sets. All problems have a box-constrained decision space, and boundary repair is performed when solutions violate the boundary conditions. Initial populations are sampled uniformly in the entire domain.

5.4.2 Performance measures

We use the Inverted Generational Distance (IGD) [32] to measure objective-space diversity and the IGDX [213] to measure decision-space diversity,

$$\begin{aligned} \text{IGD}_{\mathcal{A}^*}(\mathcal{A}) &= \frac{1}{|\mathcal{A}^*|} \sum_{\mathbf{y} \in \mathcal{A}^*} \min_{\mathbf{x} \in \mathcal{A}} \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})\|, \\ \text{IGDX}_{\mathcal{A}^*}(\mathcal{A}) &= \frac{1}{|\mathcal{A}^*|} \sum_{\mathbf{y} \in \mathcal{A}^*} \min_{\mathbf{x} \in \mathcal{A}} \|\mathbf{x} - \mathbf{y}\|, \end{aligned} \tag{5.3}$$

where $\|\cdot\|$ is the Euclidean norm, and $\mathcal{A}^*$ is the reference Pareto set, for which 5000 Pareto-optimal solutions are sampled using the analytical expression of the Pareto

set. For the Two-On-One problem, an approximation of the Pareto set is used [177]. Both IGD and IGDX measures should be minimized, and a perfect score of 0 is achieved when each solution in $\mathcal{A}^*$ is also in $\mathcal{A}$. Note that for multimodal fitness landscapes, a low IGD does not imply a low IGDX. For example, for the MinDist problem in Figure 5.1, IGD = 0 can be achieved for IGDX = 2, when one of the global Pareto sets is perfectly approximated, while no solution has been obtained in the other set. Best achievable scores for the IGD and IGDX depend on the maximum number of solutions $N_{\mathcal{A}}$ desired in the approximation set. To compute achievable limits for these scores, the reference Pareto set is compared to a subset of it that contains $N_{\mathcal{A}}$ solutions. This subset is generated using greedy scattered subset selection [27]. Subset selection has been performed with objective-space distances for the IGD and decision-space distances for the IGDX.

Additionally, we introduce a novel performance measure that we refer to as the *mode ratio* (MR), which is the ratio of attained modes. We say that a Pareto set, or mode, is attained if the IGDX for that mode is smaller than a predefined threshold ε, which we set to $\varepsilon = 0.05$ for problems with two objectives, and $\varepsilon = 0.1$ for problems with three objectives. To compute the MR, the reference Pareto set $\mathcal{A}^*$ is partitioned up into different modes by clustering it with hill-valley clustering. The MR should be maximized, and has value 1.0 when modes are attained.

We aim in this chapter to improve decision space diversity, measured by the IGDX and MR. To indicate how this affects objective space diversity, the IGD is also included. It should be noted, however, that it is not the primary aim of the multimodal approaches to optimize the IGD.

5.4.3 Visualization of hill-valley clustering

Hill-valley clustering is visualized for the 2D benchmark problems in Figure 5.2 and Figure 5.3 with initial population sizes $N = 250$ and $N = 10^4$. Tree-like clusters are formed because test solutions that are used by the hill-valley test are added to the clusters. A well-structured clustering can be observed for the Two-On-One, SYM-PART3, and three-objective MinDist problems. For the SYM-PART3 problem, low-fitness clusters with only a few solutions are formed on local Pareto sets that exist in between the global Pareto sets, which is caused by the limited number of test points in the hill-valley test. For SSUF1, the domain boundary results in that Pareto sets are connected only by a single point, resulting in more clusters than niches. For the SSUF3 problem, hill-valley clustering is not always able to connect the long stretched local Pareto sets. This can be reduced by increasing the number of neighbors considers in the clustering process.

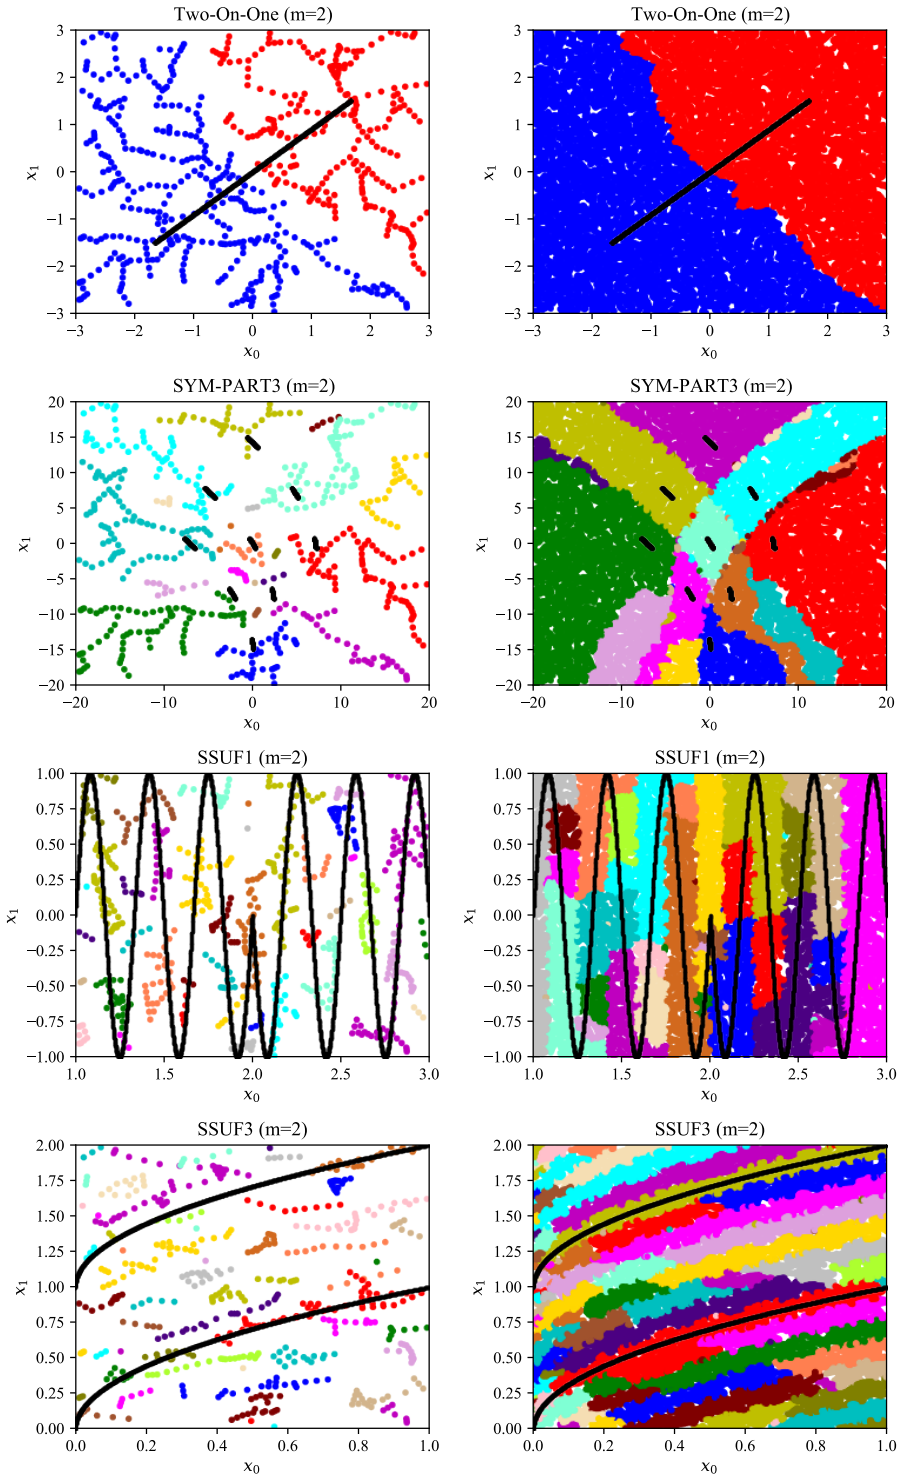

Figure 5.2: Hill-valley clustering of the decision space for different problems with two objectives, with population size of $N = 250$ (left column) and $N = 10^4$ (right column). The global Pareto sets are shown in black.

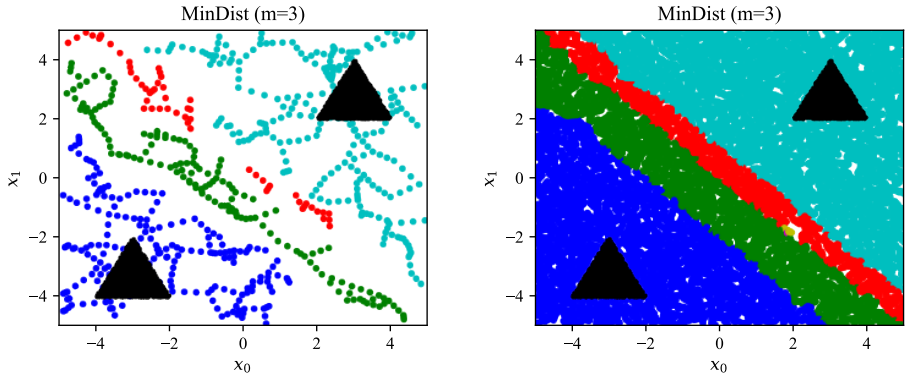

Figure 5.3: Hill-valley clustering of the decision space for the MinDist problem with three objectives, with population size of $N = 250$ (left) and $N = 10^4$ (right). The global Pareto sets are shown in black.

5.4.4 Comparing core optimization algorithms

We compare different core optimization algorithms and their performance within MO-HillVallEA in combination with the multi-start scheme. All parameters of MO-HillVallEA and MAMaLGaM are set according to literature [181]. For the bi-objective problems, the elitist archive target size is set to $N_{\mathbb{E}} = 1000$ and for the three-objective problems to $N_{\mathbb{E}} = 2500$. To show the full potential of MO-HillVallEA, post processing is disabled, by setting the approximation set equal to the union of all subarchives $\mathcal{A} = \bigcup_i \mathcal{E}_i$.

We compare four versions of MAMaLGaM as core search algorithm. MAMaLGaM (MAM) and MAMaLGaM-univariate (MAMu) estimate a Gaussian distribution with respectively a full-rank and univariate covariance matrix. Similarly for iMAMaLGaM (iMAM) and iMAMaLGaM-univariate (iMAMu), but which estimate the mean and covariance matrix incrementally over the course of multiple generations. Limiting the search distribution to a univariate covariance matrix or estimating it incrementally typically requires a smaller population size, but results in a worse fit of the fitness landscape.

Results

Results in terms of the three performance measures for MO-HillVallEA with different core optimization algorithms and MAMaLGaM are shown in Figure 5.4 and Figure 5.5. In terms of the IGD, all versions of MO-HillVallEA perform similar, and a similar rate of convergence can be observed for MO-HillVallEA and MAMaLGaM. MAMaLGaM ultimately achieves a better IGD score for all problems, which is mainly due to the limited elitist archive size. When the archive is full, discretization of the archive takes place. MO-HillVallEA aims to

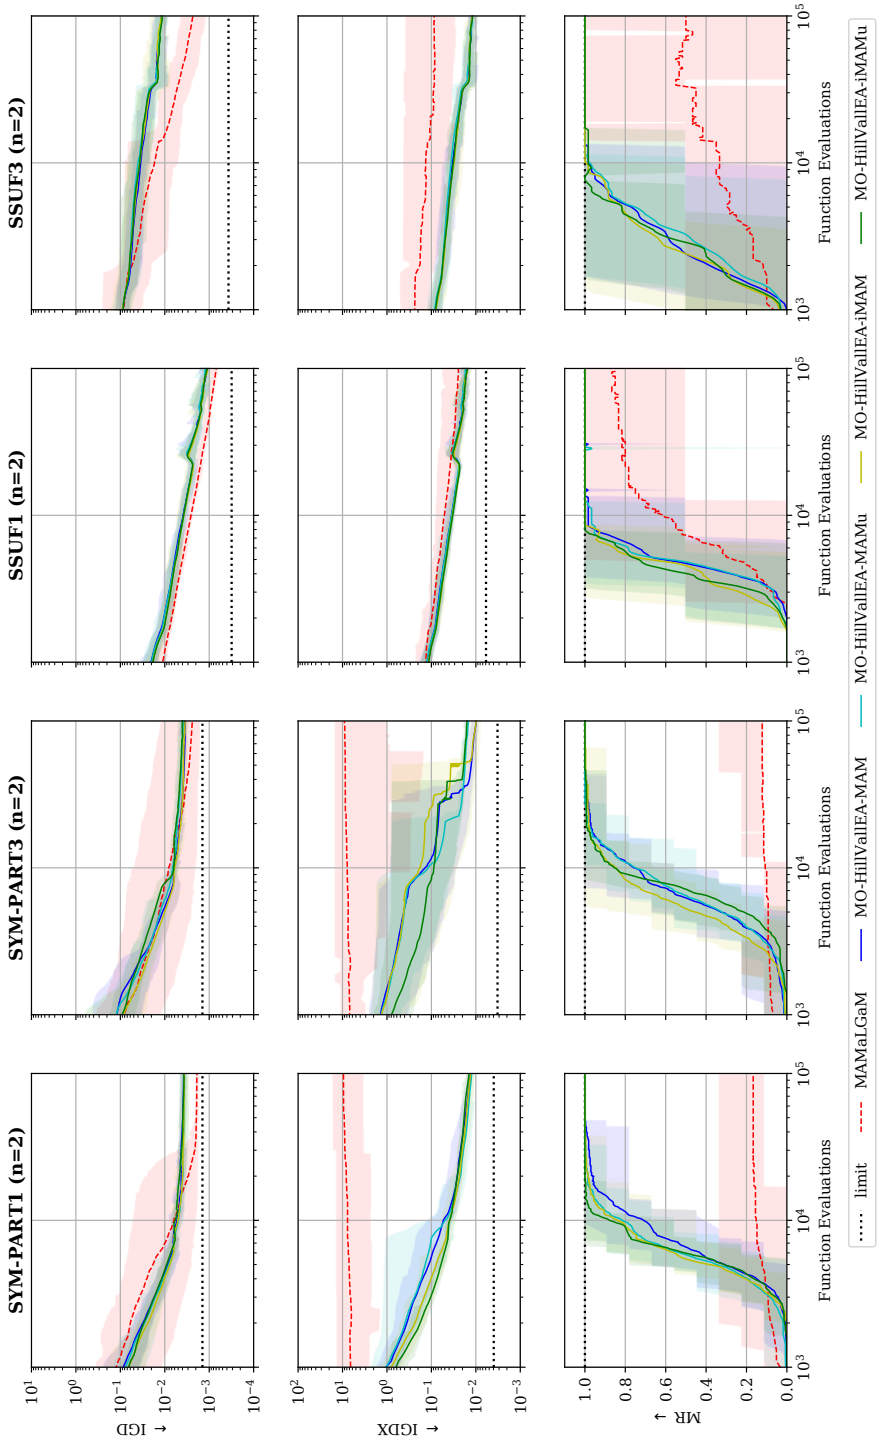

Figure 5.4: MO-HillVallEA with different core optimization algorithms compared to MAMaLGaM. Results are averaged over 30 runs, shaded areas are min/max scores.

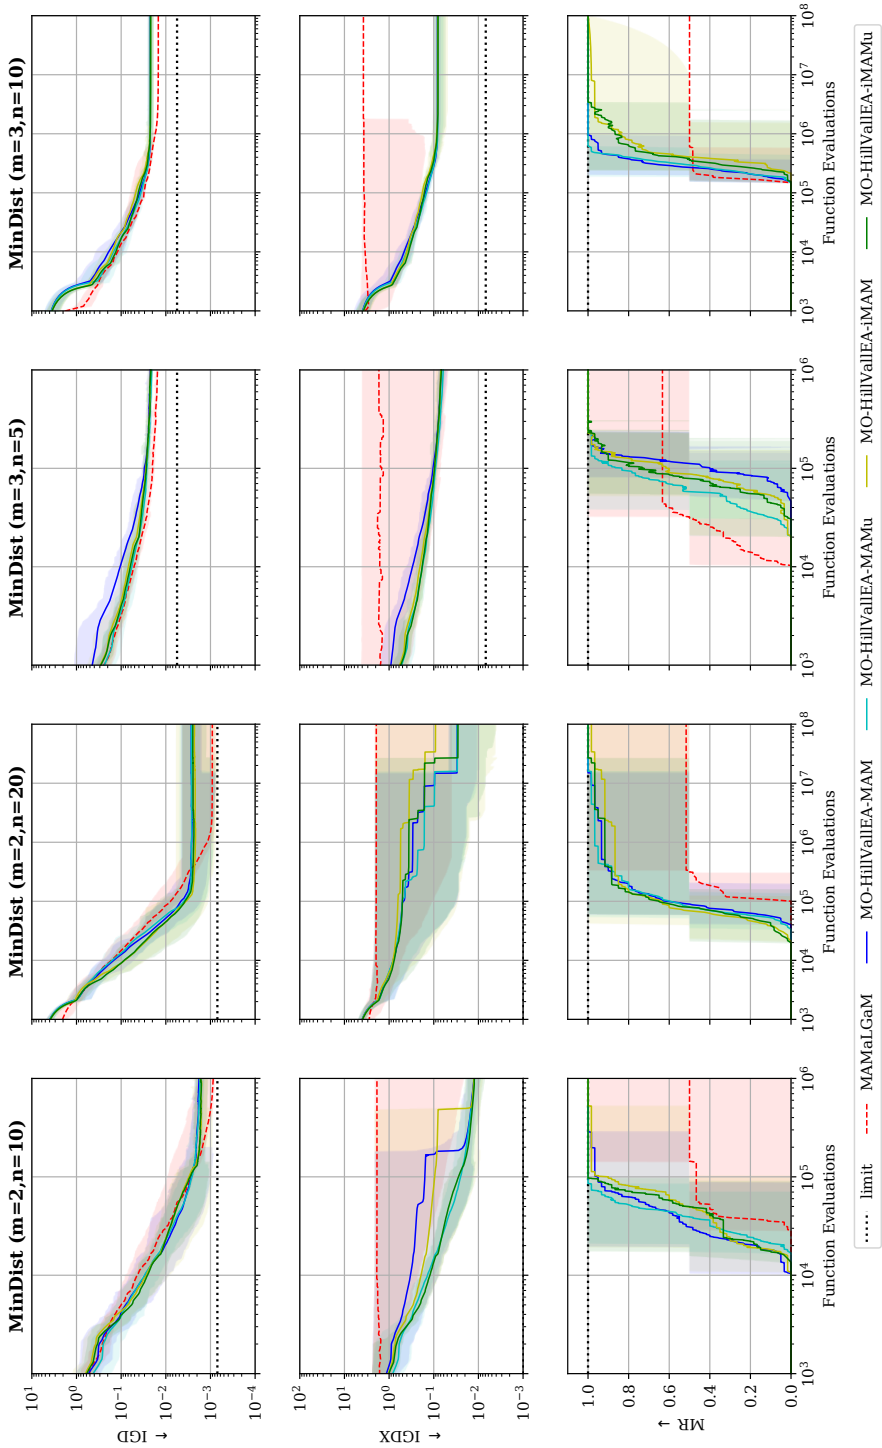

Figure 5.5: IMO-HillVallEA with different core optimization algorithms compared to MAMaLGaM. Results are averaged over 30 runs, shaded areas are min/max scores.

maintain parameter space diversity within the archive while reducing its size, while MAMaLGaM aims to maintain objective-space diversity. This results in stagnating IGD convergence for MO-HillVallEA. Initially, performance of MO-HillVallEA on the MinDist problem ($m = 2, n = 20$) is better in terms of the IGD, which can be attributed to the observation that MO-HillVallEA splits the search distribution into two, each focusing on converging locally, while MAMaLGaM tries to approach both modes with a single search distribution, resulting in a slower convergence. As soon as one of the modes is discarded due to generational drift, MAMaLGaM achieves a higher accuracy. In terms of decision-space performance, measured by the IGDX and MR, MO-HillVallEA outperforms MAMaLGaM for all problems. This holds especially for SYM-PART and MinDist problems, where MAMaLGaM ultimately obtains only a single mode, while MO-HillVallEA obtains all modes. The two modes of the MinDist problem ($m = 2, n = 20$) appear to be close to each other, due to the decision space dimensionality. Distinguishing them is hard with a small population size. Therefore, restarts are required to obtain both modes, resulting in the step-like IGDX curves.

Differences between core optimization algorithms within MO-HillVallEA can be observed for the SYM-PART1 and the bi-objective MinDist problems in IGDX and MR. The Pareto sets for these problems are linear and non-rotated, so the univariate MAMu and iMAMu outperform the full-covariance MAM and iMAM, as can be seen in terms of MR performance and a lower IGDX. For SYM-PART3, with rotated Pareto sets, the full-covariance MAM and iMAM outperform the univariate core optimization algorithms by achieving a higher MR.

5.4.5 Benchmark comparison with budget

Next, we benchmark MO-HillVallEA in the same experiment setup and budget as in [192], which allows 30 000 function evaluations and approximation set size of at most $N_{\mathcal{A}} = 100$ solutions. From the proposed problems, the Omni-Test [55] was omitted, as it has $3^5 = 243$ global Pareto sets, which is more than the number of solutions allowed in the current experiment setup. Therefore, higher IGDX scores would be achieved when each solution is a non-Pareto optimal solution that is in between Pareto sets. This is not in line with the purpose of multimodal optimization. The IGD and IGDX are used as performance measures. The MR is omitted because the limited approximation set size $N_{\mathcal{A}}$ does not allow all modes to be obtained simultaneously with a sufficient number of solutions to achieve the desired accuracy of $\varepsilon = 0.05$. Therefore, higher MR scores would be obtained by algorithms that obtain only a subset of the modes, which is again not the purpose of this experiment.

MO-HillVallEA is equipped with MAM, based on its better IGDX performance on the MinDist problems shown in Figure 5.5. The MMOEAs MO_Ring_PSO_SCD [211] and MOEA/D-AD [192], and the MOEAs MAMaLGaM, NSGA-II [53], and MOEA/D [192] are used for comparison. We set the elitist archive size to $N_{\mathbb{E}} = 1000$. The approximation set is post-processed using greedy scattered subset selection in *decision* space. The multi-start scheme was not used for MO-HillVallEA in this experiment as it is used to prevent the need to set the population size, but generally lowers performance in a limited budget setting as considered. Similar to the compared algorithms, MO-HillVallEA was ran with the standard population size recommended in literature. The population size N and number of clusters is fixed according to the recommendations in [181] to $k = 20$ clusters and population size $N = \frac{1}{2}k\lfloor 17 + 3n^{\frac{3}{2}} \rfloor = 250$, where n is the number of decision variables. For MAMaLGaM, the same setup as MO-HillVallEA is used, but the final elitist archive is post-processed with greedy scattered subset selection in *objective* space. Obtained scores are tested for significance with the Wilcoxon rank-sum test with Bonferroni correction for the 60 tests performed, resulting in a significance level of $\alpha = \frac{0.01}{60} = 0.00017$.

Results

Table 5.1 shows that MAMaLGaM comes very close to the best achievable IGD for all problems, outperforming all other algorithms. This suggests that MAMaLGaM is a sensible choice to use as a starting point for MO-HillVallEA. MO-HillVallEA outperforms all algorithms in terms of the IGDX on all problems but SYM-PART1. In all cases, MO-HillVallEA obtained all nine Pareto sets of the SYM-PART problems, which becomes increasingly harder for respectively SYM-PART1–3. The reference Pareto set of the Two-On-One problem is an approximation that is accurate up to 0.0045 [177], and multiple algorithms have IGD and IGDX scores close to this accuracy, which may obfuscate true performance.

5.5 Discussion

We showed that hill-valley clustering can be used to successfully distinguish multi-objective niches. These niches are then explored separately in MO-HillVallEA. By considering Pareto domination per niche, local Pareto sets can be maintained. However, for problems with a large number of niches, (e.g. Omni-Test [55], SSUF1, SSUF3), this approach results in a large number of clusters, and a very large population is required to explore all niches simultaneously. In that case, controlling

Table 5.1: Mean IGD and IGDX for MMOEAs and MOEAs over 31 runs and standard deviation in brackets. Bold scores are best per problem or not significantly worse than any other. In the *limit* column, best scores achievable with $N_{\mathcal{A}} = 100$ are shown.

Problem	Limit	MO-HillVallEA	MOEA/D-AD	MO_Ring_PSO_SCD
		IGD (objective space)		
Two-On-One	0.004	0.006 *(0.000)*	0.064 *(0.008)*	0.061 *(0.006)*
SYM-PART1	0.018	0.047 *(0.009)*	0.030 *(0.002)*	0.028 *(0.001)*
SYM-PART2	0.018	0.044 *(0.006)*	0.031 *(0.002)*	0.031 *(0.002)*
SYM-PART3	0.018	0.043 *(0.005)*	0.031 *(0.002)*	0.032 *(0.003)*
SSUF1	0.004	0.008 *(0.001)*	0.007 *(0.001)*	0.006 *(0.001)*
SSUF3	0.004	0.012 *(0.001)*	0.019 *(0.006)*	0.011 *(0.002)*
		IGDX (parameter space)		
Two-On-One	0.013	**0.026** *(0.001)*	0.035 *(0.003)*	0.037 *(0.002)*
SYM-PART1	0.051	**0.073** *(0.007)*	**0.069** *(0.003)*	0.148 *(0.024)*
SYM-PART2	0.052	**0.070** *(0.006)*	0.078 *(0.003)*	0.161 *(0.026)*
SYM-PART3	0.042	**0.053** *(0.004)*	0.148 *(0.209)*	0.491 *(0.369)*
SSUF1	0.055	**0.057** *(0.001)*	0.076 *(0.008)*	0.086 *(0.006)*
SSUF3	0.008	**0.016** *(0.003)*	0.030 *(0.009)*	0.020 *(0.006)*

Problem	Limit	MAMaLGaM	NSGA-II	MOEA/D
		IGD (objective space)		
Two-On-One	0.004	**0.004** *(0.000)*	0.049 *(0.002)*	0.045 (0.001)
SYM-PART1	0.018	**0.018** *(0.000)*	0.021 *(0.001)*	0.047 *(0.001)*
SYM-PART2	0.018	**0.018** *(0.000)*	0.023 *(0.001)*	0.047 *(0.008)*
SYM-PART3	0.018	**0.019** *(0.001)*	0.023 *(0.001)*	0.045 *(0.005)*
SSUF1	0.004	**0.004** *(0.000)*	0.006 *(0.000)*	0.006 *(0.003)*
SSUF3	0.004	**0.005** *(0.001)*	0.007 *(0.002)*	0.063 *(0.052)*
		IGDX (parameter space)		
Two-On-One	0.013	0.043 *(0.012)*	0.148 *(0.118)*	0.281 *(0.164)*
SYM-PART1	0.051	9.427 *(1.520)*	7.929 *(2.343)*	9.155 *(2.748)*
SYM-PART2	0.052	9.410 *(1.082)*	5.371 *(1.964)*	9.483 *(2.191)*
SYM-PART3	0.042	8.335 *(3.050)*	5.841 *(1.892)*	7.397 *(1.965)*
SSUF1	0.055	0.142 *(0.040)*	0.132 *(0.022)*	0.244 *(0.065)*
SSUF3	0.008	0.162 *(0.077)*	0.071 *(0.048)*	0.308 *(0.109)*

the maximum number of clusters, or a serial search, as employed for single objective optimization in Chapter 4, might be preferred.

The visualization of hill-valley clustering on problems with two decision variables provided useful insight in its behavior. The performance of MO-HillVallEA on the MinDist problem with up to 20 variables shows that this technique extends to higher-dimensional problems. However, its performance and scalability has yet to be shown for more complex problems and even higher-dimensional decision spaces.

The current problem formulation of multimodal optimization, where one is interested in only locating global Pareto sets, poses a fundamental difficulty. In practice, especially when a problem contains noise, the global Pareto front is never exactly attained, which makes it difficult or impossible to distinguish between global and high-quality local Pareto sets.

It would furthermore be interesting to see how hill-valley clustering can be applied to improve performance when the aim is objective-space diversity in a multimodal fitness landscape, rather than decision space diversity. By aiming MO-HillVallEA for decision-space diversity, objective-space diversity deteriorated for most problems, compared to MAMaLGaM. This deterioration was found to relate mainly to the size limit of the elitist archive. If a balance between diversity in objective and decision space is desired, different size-control mechanisms could be applied to the elitist archive.

MOEA/D-AD [192] is equipped with a population size growing scheme by which the population size is adapted to the problem at hand. A population-sizing scheme was used, but adapting this scheme or the population size itself on the number of detected niches could further improve performance. Both MOEA/D-AD and MO-HillVallEA use a post-processing step to construct a limited-size approximation set while maintaining a larger elitist archive during optimization, which seems to be beneficial.

5.6 Conclusion

We introduced hill-valley clustering for multi-objective niching by clustering the search space. We combined hill-valley clustering with MAMaLGaM into the multimodal multi-objective evolutionary algorithm MO-HillVallEA. We empirically show that MO-HillVallEA outperforms MAMaLGaM and other multi-objective optimization algorithms in multimodal optimization on a set of multimodal benchmark functions. Furthermore, and perhaps most importantly, we show that MO-HillVallEA is capable of obtaining, maintaining, and exploiting multiple local Pareto sets simultaneously over time.

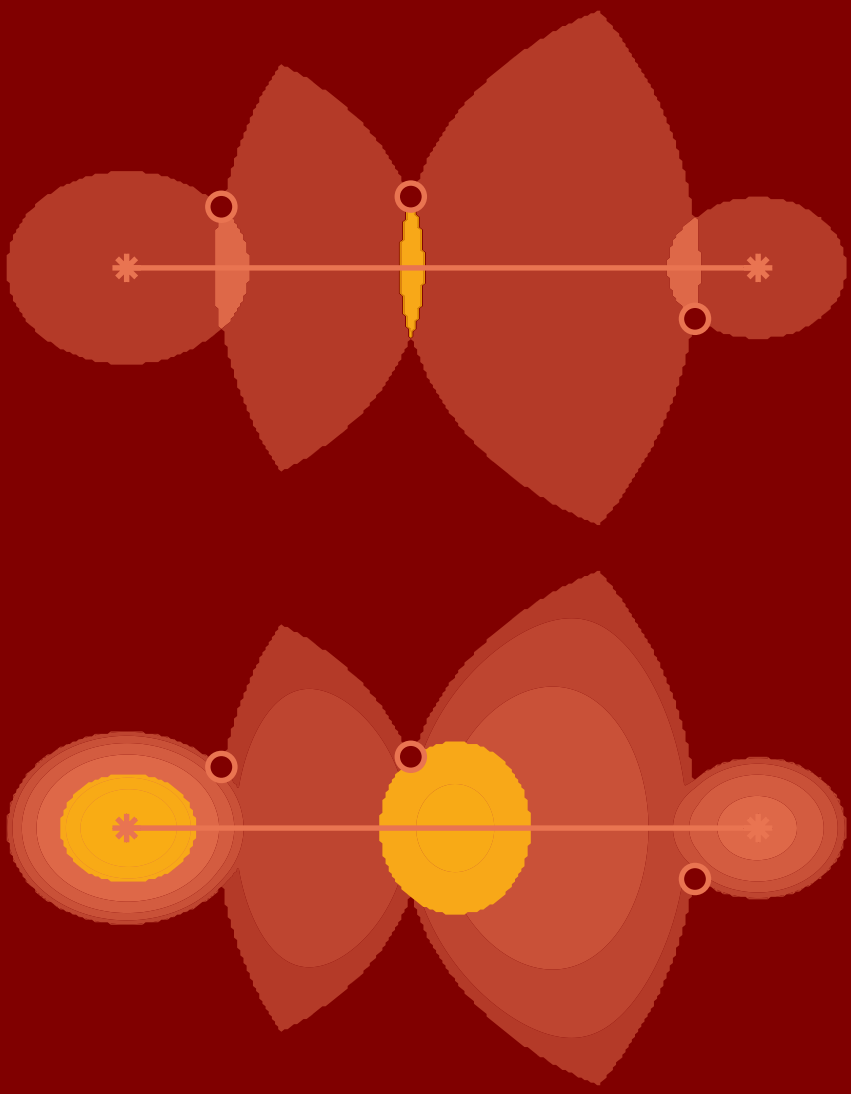

6

Uncrowded hypervolume-based multi-objective optimization

Research question 5:
How can evolutionary multi-objective optimization approaches be designed where the individuals in the population represent entire approximation sets and fitness is directly based on the hypervolume indicator or extensions thereof so as to ensure convergence to a subset of the Pareto set?

This chapter is an adaptation of S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2020) *Uncrowded Hypervolume-based Multi-objective Optimization with Gene-pool Optimal Mixing.* (submitted for journal publication, preprint arXiv:2004.05068).

Abstract

Domination-based Multi-Objective Evolutionary Algorithms (MOEAs) are today arguably the most frequently used type of MOEA. These methods however stagnate when the majority of the population becomes non-dominated, preventing convergence to the Pareto set. Hypervolume-based multi-objective optimization has shown promising results to overcome this. Direct use of the hypervolume however results in no selection pressure for dominated solutions. The recently introduced Sofomore framework overcomes this by solving multiple interleaved single-objective dynamic problems that iteratively improve a single approximation set, based on the Uncrowded Hypervolume Improvement (UHVI). It thereby however loses many advantages of population-based multi-objective optimization, such as handling multimodality. In this chapter, we reformulate the UHVI as a quality measure for approximation sets, called the Uncrowded Hypervolume (UHV), which can be used to directly solve multi-objective optimization problems with a single-objective optimizer. We use the state-of-the-art Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) that is capable of efficiently exploiting the intrinsically available grey-box properties of this problem. The resulting algorithm, UHV-GOMEA, is compared to Sofomore equipped with GOMEA, and the domination-based MO-GOMEA. In doing so, we investigate in which scenarios either domination-based or hypervolume-based methods are preferred. Finally, we construct a simple hybrid approach that combines MO-GOMEA with UHV-GOMEA and outperforms both.

6.1 Introduction

A multi-objective optimization problem is given by a to-be-minimized objective function $\mathbf{f} : \mathcal{X} \to \mathbb{R}^m$, with $\mathbf{f} = [f_1, \ldots, f_m]$ being an m-dimensional vector function and $\mathcal{X} \subseteq \mathbb{R}^n$ the n-dimensional decision space. Multi-objective problems do not naturally imply a complete ordering of solutions $\mathbf{x} \in \mathcal{X}$ based on their objective values. When the objectives contradict each other, no single solution minimizes all objectives simultaneously. The optimum of a multi-objective problem can then be defined in terms of *Pareto optimality* [113]. A solution is said to *dominate* another solution when it is strictly better in one or more objectives, and is not worse in any of the other objectives. A solution is *Pareto optimal* when there exist

no solutions that dominates it. The *Pareto set* is the set of all Pareto optimal solutions, and its image under $\mathbf{f}$ is known as the *Pareto front*. In practice, the aim of multi-objective optimization is to provide a decision maker with a set of non-dominated solutions, known as an *approximation set* $\mathcal{A} \subset \mathcal{X}$, whose image under $\mathbf{f}$, the *approximation front*, approximates the Pareto front. The decision maker then selects a preferred solution from this set. Real-valued multi-objective problems generally have an infinite number of Pareto-optimal solutions. As it is impossible to obtain all of these, multi-objective optimization algorithms generally attempt to find an approximation set containing a manageable number of solutions that forms a good representation of the entire front. This results in an inherent trade-off in the two-sided optimization goal of multi-objective optimization, since it is desirable to obtain a diverse approximation set as well as an approximation set that is close to the Pareto set [33].

Multi-Objective Evolutionary Algorithms (MOEAs) have shown to be very successful for (black-box) multi-objective optimization in practice [52]. These algorithms maintain a population of solutions, and generally equip a domination-based fitness selection scheme [38, 53, 218], such as in the Non-dominated Sorting Genetic Algorithm (NSGA-II) [53], arguably the best-known MOEA. However, if the population size is (much) smaller than the number of Pareto-optimal solutions, at some point in the optimization process the majority of solutions in the population will be non-dominated. MOEAs then typically aim to improve diversity in the population, for example using the crowding distance in NSGA-II. This could lead to cyclic behavior, with the MOEA improving diversity but worsening proximity. Consequently, even though the approximation sets obtained with domination-based MOEAs are often sufficient for practical use, these algorithms do not result in approximation sets converging to the Pareto set [19].

Indicator-based multi-objective optimization, especially based on the hypervolume measure [219], has shown promising results to overcome this limitation [21, 102]. Practically used indicators in indicator-based optimization are the R2 indicator [83], the epsilon indicator [217], and the hypervolume measure [220]. The hypervolume measure is particularly interesting as it is the only known measure that is strictly monotonic with respect to Pareto-dominance [68, 112]. This means that the approximation set with optimal hypervolume is a subset of the Pareto set. A limitation of the hypervolume measure is that dominated solutions have no contribution, in the sense that there is no selection pressure for dominated solutions towards a non-dominated region in the decision space [63]. Therefore, in the first usages of the hypervolume measure in the optimization process, it was only partially used to guide the search [62, 102, 161, 164, 216]. Specifically, in the

$\mathcal{S}$-Metric Selection Evolutionary Multi-objective Optimization Algorithm (SMS-EMOA) [21], indicators such as the hypervolume measure, are therefore used as a secondary fitness after non-dominated sorting. To overcome this limitation of the hypervolume, without relying on domination-based properties, a fitness value can be assigned to dominated solutions based on their distance to the boundary of the dominated area, i.e., the boundary surface of the hypervolume [63]. The resulting Uncrowded Hypervolume Improvement (UHVI) is not a set-based measure, and cannot be used directly to optimize approximation sets. The Sofomore framework [195] was therefore formulated, in which multiple optimizers are interleaved that each solve a dynamic single-objective problem. In this framework, a single approximation set is iteratively optimized, but by doing so, it loses some of the advantages of population-based algorithms, such as the ability to escape local optima.

In this chapter, we show how hypervolume-based multi-objective optimization can be used to replace or supplement domination-based MOEAs. To this end, we formulate the Uncrowded Hypervolume (UHV) measure, which can be directly used to achieve population-based hypervolume-driven multi-objective optimization using a single-objective problem formulation.

We use the Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) [37, 38] as baseline algorithm to compare domination-based and hypervolume-based multi-objective optimization. This algorithm has both a single-objective [37] and multi-objective [38] version, with recently published excellent results. In both versions, essentially, the same variation operators are used, which provides a more fair and modern comparison than when e.g., NSGA-II and the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) [86] were to be used. Furthermore, in this chapter we will show that the single-objective GOMEA can be used to efficiently solve the UHV problem formulation by exploiting grey-box properties that are intrinsically present in the UHV, while still assuming the multi-objective problem itself to be black box.

The remainder of this chapter is organized as follows. In Section 6.1.1, we discuss the hypervolume and related measures. In Section 6.2, we introduce an indicator-based multi-objective optimization problem formulation based on the UHV, and how we can efficiently solve this problem with GOMEA, resulting in UHV-GOMEA. In Section 6.3, we furthermore incorporate GOMEA in the Sofomore framework (Sofomore-GOMEA), and in Section 6.4, we discuss MO-GOMEA. In Section 6.5, we empirically compare these three algorithms and discuss which of the approaches is preferred in which scenario, and how to combine them into a simple hybrid approach, on both simple benchmark problems and the commonly used WFG test suite [101]. We discuss the results in Section 6.6 and conclude in Section 6.7.

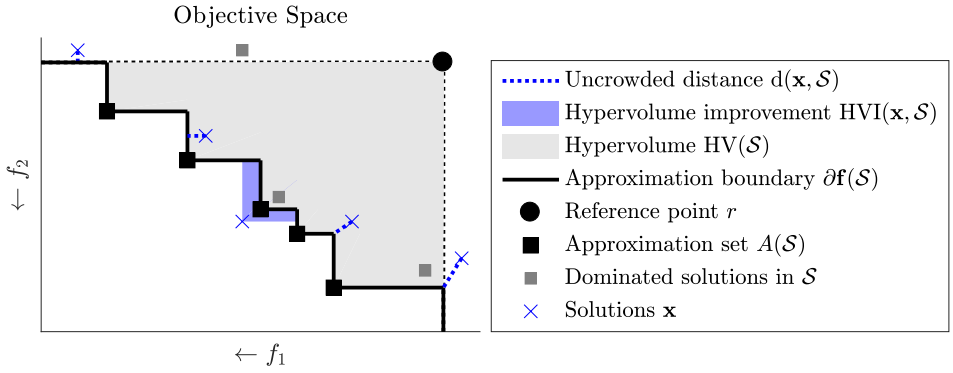

Figure 6.1: Illustration of the hypervolume measure for a bi-objective minimization problem, together with the uncrowded distance and hypervolume improvement of some example solutions $\mathbf{x}$ with respect to solution set $\mathcal{S}$.

6.1.1 Preliminaries: The hypervolume measure

A solution $\mathbf{x} \in \mathcal{X}$ is said to *weakly dominate* another solution $\mathbf{y} \in \mathcal{X}$, written as $\mathbf{x} \preceq \mathbf{y}$, if and only if $f_i(\mathbf{x}) \leq f_i(\mathbf{y})$ for all $i \in \{1, \ldots, m\}$. When the relation $f_i(\mathbf{x}) < f_i(\mathbf{y})$ is furthermore strict for at least one i, we say that $\mathbf{x}$ *dominates* $\mathbf{y}$, written as $\mathbf{x} \prec \mathbf{y}$ or, with a slight abuse of notation, as $\mathbf{f}(\mathbf{x}) < \mathbf{f}(\mathbf{y})$. A solution that is not dominated by any other solution in $\mathcal{X}$ is called *Pareto optimal*. The *Pareto set* $\mathcal{A}^\star$ can then be formulated as,

$$\mathcal{A}^\star = \{\mathbf{x} \in \mathcal{X} : \nexists \mathbf{y} \in \mathcal{X} : \mathbf{y} \prec \mathbf{x}\} \subset \mathcal{X},$$

and the corresponding *Pareto front* is given by $\{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in \mathcal{A}^\star\} \subset \mathbb{R}^m$.

Let $\wp(\mathcal{X})$ be the powerset of $\mathcal{X}$, i.e., the set of all solution sets $\mathcal{S} \subseteq \mathcal{X}$. The *hypervolume* measure $\mathrm{HV} : \wp(\mathcal{X}) \to \mathbb{R}$ [7, 219, 220] of a solution set $\mathcal{S}$ with respect to a reference point $r \in \mathbb{R}^m$ measures the volume dominated by all $\mathbf{x} \in \mathcal{S}$, and bounded by r, as is illustrated in Figure 6.1. Let $A : \wp(\mathcal{X}) \to \wp(\mathcal{X})$ be the *approximation set* of $\mathcal{S}$,

$$A(\mathcal{S}) = \{\mathbf{x} \in \mathcal{S} : \mathbf{f}(\mathbf{x}) < r, \nexists \mathbf{y} \in \mathcal{S} : \mathbf{y} \prec \mathbf{x}\}.$$

That is, $A(\mathcal{S})$ is the largest subset of $\mathcal{S}$ that contains only non-dominated solutions within the region defined by r. Additionally, let the *hypervolume improvement* $\mathrm{HVI} : \mathcal{X} \times \wp(\mathcal{X}) \to \mathbb{R}$ of a solution $\mathbf{x}$ with respect to a solution set $\mathcal{S}$ be defined as the increase in hypervolume when $\mathbf{x}$ is added to $\mathcal{S}$, which can be formulated as $\mathrm{HVI}(\mathbf{x}, \mathcal{S}) = \mathrm{HV}(\mathcal{S} \cup \{\mathbf{x}\}) - \mathrm{HV}(\mathcal{S})$, as shown in Figure 6.1.

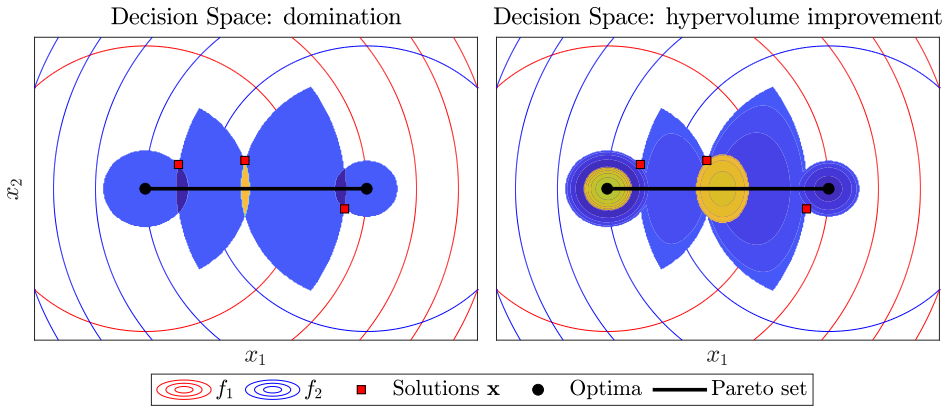

Figure 6.2: Domination region (left) compared to the hypervolume improvement region (right) for the bi-sphere problem. Solutions in the light blue region are non-dominated with respect to the current solutions. The dark blue regions dominate one of the current solutions. In the right subfigure, solutions in the blue region improve the hypervolume if added to the solution set, and darker blue represents larger improvements. Additionally, in both subfigures, the orange region is the improvement region when the middle solution is being replaced.

The difference between domination-based improvements and hypervolume-based improvements is illustrated in Figure 6.2 for the bi-sphere problem with hypervolume reference point $r = (11, 11)$. This problem is composed of two single-objective sphere problems, $f_{\text{sphere}}(\mathbf{x}) = \sum_{i=1}^{n} x_i^2$, of which one is translated, i.e., $\mathbf{f}_{\text{bi-sphere}}(\mathbf{x}) = [f_{\text{sphere}}(\mathbf{x}) \; ; \; f_{\text{sphere}}(\mathbf{x} - \mathbf{e}_1)]$, where $\mathbf{e}_i$ is the i^{th} unit vector, with all zeros except a one in the i^{th} position. Figure 6.2 shows that it holds that,

$$\text{HVI}(\mathbf{x}, \mathcal{S}) \geq 0 \Leftrightarrow \{\nexists \mathbf{y} \in \mathcal{S} : \mathbf{y} \prec \mathbf{x}\},$$

as both approaches have the same improvement region (in blue). However, when comparing the *improvement region for the middle solution* (in orange), i.e., when that solution is replaced in $\mathcal{S}$, we see that its hypervolume improvement region is larger, essentially making it easier to find improvements. Additionally, the hypervolume improvement region consists of two disconnected subsets, indicating that it takes diversity into account, in contrast to domination-based improvements.

When $\mathbf{x}$ is dominated by any solution in $\mathcal{S}$, it has zero hypervolume improvement, i.e., $\text{HVI}(\mathbf{x}, \mathcal{S}) = 0$. The Uncrowded Hypervolume Improvement (UHVI) [195] was recently introduced to overcome this. Let $\partial\mathbf{f}(\mathcal{S})$ be the *approximation boundary*, i.e., the boundary between the dominated and non-dominated region in objective space, bounded by the reference point r, as illustrated in Figure 6.1.

Let $\mathrm{ud}(\mathbf{x}, \mathcal{S})$ be the *uncrowded distance* [63], which measures the shortest Euclidean distance between $\mathbf{x}$ and $\partial\mathbf{f}(\mathcal{S}) = \partial\mathbf{f}(A(\mathcal{S}))$, when $\mathbf{x}$ is dominated by any solution in $\mathcal{S}$ or outside the region defined by r. Else, we set $\mathrm{ud}(\mathbf{x}, \mathcal{S}) = 0$. It is called the uncrowded distance as the shortest distance to $\partial\mathbf{f}(\mathcal{S})$ is obtained for a point on $\partial\mathbf{f}(\mathcal{S})$ that is not in $\mathcal{S}$ itself. The UHVI can then be defined as,

$$\mathrm{UHVI}(\mathbf{x}, \mathcal{S}) = \mathrm{HVI}(\mathbf{x}, \mathcal{S}) - \mathrm{ud}(\mathbf{x}, \mathcal{S}).$$

It is interesting to note that the uncrowded distance is similar to the distance measure d^+ that was used to construct a weak Pareto compliance version of the inverted generational distance indicator, called the IGD^+ [104]. In the IGD^+, a reference set $\mathcal{Z}$ of non-dominated solutions is used (i.e., a subset of the known Pareto set), and the distance of reference solutions $\mathbf{z} \in \mathcal{Z}$ towards the approximation boundary $\partial\mathbf{f}(\mathcal{S})$ is computed. There, the solutions $\mathbf{z}$ are non-dominated with respect to $\mathcal{S}$. Here however, we are particularly interested in the other case, where solutions are dominated with respect to $\mathcal{S}$.

6.2 The uncrowded hypervolume measure

The UHVI is a measure of solution quality with respect to a solution set $\mathcal{S}$, but not a measure of quality for $\mathcal{S}$ itself. We therefore introduce the *Uncrowded Hypervolume* (UHV) in this chapter as the hypervolume of $\mathcal{S}$ penalized by all uncrowded distances,

$$\mathrm{UHV}(\mathcal{S}) = \mathrm{HV}(\mathcal{S}) - \frac{1}{|\mathcal{S}|} \sum_{\mathbf{x} \in \mathcal{S}} \mathrm{ud}(\mathbf{x}, \mathcal{S})^m. \tag{6.1}$$

The uncrowded distances are taken to the power m such that they have the same unit as the hypervolume. As improving a non-dominated solution in $\mathcal{S}$ could increase both its hypervolume and the uncrowded distances, the factor $1/|\mathcal{S}|$ is added to guarantee that an improvement in hypervolume is not negatively influenced by the increase in uncrowded distances. The uncrowded hypervolume is a strictly monotonic indicator on the space of approximation sets (i.e., sets containing only non-dominated solutions), as it is equal to the hypervolume for those sets. However, on the entire space of solution sets, strict monotonicity does not hold.

6.2.1 UHV-maximization

In an Indicator-Based Multi-objective Optimization Problem (IBMOP), a quality indicator is used to assign a quality value to a solution set. The underlying idea is that a single-objective optimizer then can be used to explicitly search for a solution set $\mathcal{S}$ that maximizes this indicator [21, 217]. Let $I_{\mathbf{f}} : \wp(\mathcal{X}) \to \mathbb{R}$ be such an indicator, with respect to the multi-objective problem given by $\mathbf{f}$. To be able to search the space of solution sets, we parameterize solutions sets by considering sets of fixed size $|\mathcal{S}_p| = p \geq 1$. Let ϕ be a vector of concatenated decision variables of the solutions in $\mathcal{S}_p$, i.e., $\phi = [\mathbf{x}_1 \cdots \mathbf{x}_p] \in \mathbb{R}^{p \cdot n}$. Inversely, let $S(\phi) = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ be the operator that transforms ϕ into a solution set. To avoid confusion, we will refer to a solution ϕ of the IBMOP with objective function g as a g-solution, while solutions of the original multi-objective optimization problem given by $\mathbf{f}$ are called MO-solutions from now on. IBMOPs are then formulated as,

$$\begin{aligned} \text{maximize} \quad & g_{\mathbf{f},p} : (\mathcal{X})^p \to \mathbb{R}, \\ \text{with} \quad & g_{\mathbf{f},p}(\phi) = I_{\mathbf{f}}(S(\phi)), \\ & \mathbf{f} : \mathcal{X} \subseteq \mathbb{R}^n \to \mathbb{R}^m, p \geq 1 \\ & \phi = [\mathbf{x}_1 \cdots \mathbf{x}_p] \in (\mathcal{X})^p \subseteq \mathbb{R}^{p \cdot n}. \end{aligned} \tag{6.2}$$

This IBMOP formulation fully specifies the multi-objective optimization problem, as well as its optimum, i.e., the resulting distribution of solutions along the front. In case of the (uncrowded) hypervolume as quality indicator, this optimal distribution known as the optimal μ-distribution [8]. The density of solutions along the front is inverse proportional to the negative slope of the front [8]. If and only if the front is linear are the solutions in the optimal distribution equally spaced.

This clear and unique definition of the optimum of the IBMOP allows us to discuss convergence to optimality, in contrast to the general aim of multi-objective optimization that translates to a trade-off between proximity and diversity [33]. Furthermore, since the hypervolume measure is strictly monotonic with respect to Pareto dominance, optimality here implies that the obtained solutions are a subset of the Pareto set. Note that the optimal distribution of solutions along the front is determined by the choice of reference point, which we fix during the course of the optimization run. If a dynamic reference point would have been used, cyclic behavior could again occur.

6.2.2 UHV-GOMEA

We use the real-valued single-objective GOMEA to solve IBMOPs with the UHV as indicator, which we call UHV-GOMEA. We use GOMEA as published in [37],

and make only minor adaptations to better align the algorithm with IBMOPs. We discuss the outline of GOMEA here, and refer the reader to [37] for a full description.

GOMEA is a model-based evolutionary algorithm that maintains a population of N solutions. For the IBMOP these are the g-solutions $\phi^j \in \mathbb{R}^{p \cdot n}$, $j = 1, \ldots, N$. The variation operator in GOMEA is called gene-pool optimal mixing (GOM). GOM was designed specifically to perform variation by adapting only a few decision variables at a time, and thereby exploiting that objective functions can often be quickly updated if only a few variables change. This notion of *partial evaluations* typically requires that some problem knowledge is known, resulting in a *grey-box* scenario. IBMOPs are by definition grey-box, in the sense that it is known how to update the indicator value when the decision variables corresponding to only one (or a few) MO-solution change, without having to re-evaluate all other MO-solutions, or re-compute the indicator value from scratch. This means that grey-box properties of the IBMOP can be exploited while still considering $\mathbf{f}$ as a black-box.

GOMEA is equipped with a *linkage model* that specifies which subsets of decision variables must be adapted simultaneously. A linkage model L is a subset of the power set $\wp(I)$ of all decision variables, which is $I = \{1, \ldots, pn\}$ in case of an IBMOP. A linkage model is therefore also known as a *family of subsets*. In GOM, variation is performed by iteratively updating only the decision variables specified by a linkage subset $l \in L$, with $l \subseteq I$. We denote the subset of decision variables in a g-solution ϕ specified by l with $\phi_{\langle l \rangle} \in \mathbb{R}^{|l|}$. An $|l|$-dimensional Gaussian distribution is then estimated from the τN best g-solutions (with $\tau = 0.35$). From this distribution, to update each g-solution in the population, new values for the decision variables specified by l are sampled. In GOM, only improvements to g-solutions are accepted. Else, the proposed update is discarded.

Intermediate covariance updates

Originally in GOM, the sample distributions for all linkage subsets $l \in L$ were estimated at the beginning of each generation [37]. Instead, here, we estimate the sample distribution for each linkage subset right before sampling new values for the variables in that subset. By doing so, the sample model is learned based on the most recent set of g-solutions. We found this to improve the rate of convergence, at increased computationally complexity of only $\mathcal{O}((|L|-1) \cdot N \log N)$ per generation, as selection now needs to be performed for each linkage subset. A Cholesky decomposition [91, 122] of the covariance matrix is then performed which is required for the sampling. This decomposition requires a positive definite covariance matrix. Due to numerical errors or small population size, it might be

that the decomposition fails [137]. If this happens once for a linkage subset, we from then on perform a regularization using the Ledoit-Wolf shrinkage estimator (LWSE) [123]. The LWSE estimates a covariance matrix based on a convex combination of the maximum likelihood estimator and a prior matrix, for which we use the diagonal variance matrix (i.e., the diagonal of the covariance matrix). Additionally, when $|l| > \tau N - 1$, we always use a diagonal variance matrix without attempting a decomposition.

IBMOP linkage models

Any linkage model can be used in GOMEA, but we discuss three models that we employ in UHV-GOMEA in this chapter. Let $I_{(i)} = \{\phi_{(i-1)n+1}, \ldots, \phi_{in}\}$ for $i = 1, \ldots, p$ be the decision variables corresponding to the MO-solution $\mathbf{x}_i$ of g-solution $\phi = [\mathbf{x}_1 \cdots \mathbf{x}_p]$. The three linkage models we consider are given by,

$$\begin{aligned} Lm &= \left\{I_{(1)}, I_{(2)}, \ldots, I_{(p)}\right\}, \\ Lf &= \left\{I_{(1)} \cup I_{(2)} \cup \cdots \cup I_{(p)}\right\} = I, \\ Lt &= \text{UPGMA}(Lm). \end{aligned} \tag{6.3}$$

In the *marginal linkage model* Lm, only decisions variables corresponding to the same MO-solution are considered to be dependent, i.e., $\phi_{\langle l_i \rangle} = \mathbf{x}_i$ for $l_i \in Lm$. In the *full linkage model* Lf, the linkage model contains only a single subset, in which all decision variables are considered to be dependent, i.e., $\phi_{\langle l \rangle} = \phi_{\langle I \rangle} = \phi$. Finally, the *linkage tree model* Lt contains multiple levels of linkage, constructed by hierarchically clustering the subsets of the marginal model Lm using UPGMA [78], where the two *nearest* subsets are merged iteratively and added to Lt. Since each subset in $l_i \in Lm$ corresponds to the parameters of a single MO-solution $\mathbf{x}_i$, we can compute the mean objective values of each l_i,

$$\mathbf{m}_i = \frac{1}{N} \sum_{j=1}^{N} \mathbf{f}\left(\phi^j_{\langle l_i \rangle}\right) = \frac{1}{N} \sum_{j=1}^{N} \mathbf{f}\left(\mathbf{x}_i^j\right). \tag{6.4}$$

We use the distance between the means $\mathbf{m}_i$ as distance measure between linkage subsets. Since these change over time, the linkage tree is re-constructed every generation. Linkage tree construction results in $|Lt| = 2p - 1$ subsets, and Lt always contains all elements of Lm as well as Lf, but the other subsets depend on the merge order defined by this distance measure. Additionally, in GOM, all linkage subsets for which $|l| > \tau N - 1$ are skipped, as for these relatively large linkage subsets, dependencies cannot be estimated, as a full-rank covariance matrix cannot be estimated.

Permuting MO-solutions

Since permuting solutions in $\mathcal{S}$ does not influence its corresponding hypervolume, we aim to reorder the MO-solutions in $S(\phi)$, and therefore g-solution ϕ, at the beginning of a generation, such that all $\phi^j_{\langle l_i \rangle} = \mathbf{x}^j_i$ for each g-solution $j = 1, \ldots, N$ belong to a similar part of the approximation front. Therefore, we again compute the objective-space means $\mathbf{m}_i$ for $i = 1, \ldots, p$. For each g-solution ϕ^j, the MO-solutions are re-ordered in a greedy fashion by iteratively finding the (next-)nearest MO-solution-mean pair.

Elitist archive

An elitist archive $\mathcal{E} \subset \mathcal{X}$ is maintained that contains all non-dominated MO-solutions $\mathbf{x} \in \mathcal{X}$ that were evaluated during optimization. To keep the archive size tractable, adaptive objective space discretization was used as presented in [134]. If, with the archiving scheme, the archive exceeds the target size of $N_\mathcal{E} = 1000$ MO-solutions, the objective space is discretized into boxes, and only one MO-solution per box is maintained. Newly obtained non-dominated MO-solutions are then only added to the archive when they end up in an empty objective-space box, or when they dominate the MO-solution in that box. This archiving scheme is also used in MO-GOMEA [38], and the target size of the elitist archive $\mathcal{E}$ is set to $N_\mathcal{E} = 1000$ MO-solutions. The elitist archive is not an essential part of UHV-GOMEA, but merely added to allow for a comparison to archive-based MOEAs.

6.2.3 Generational computational complexity

A main limitation of the usage of the hypervolume measure in the optimization process is its computational complexity. Computation of the UHV of a solution set $\mathcal{S}_p$ containing p MO-solutions consists of three steps. First, the set of non-dominated solutions $A(\mathcal{S}_p)$ needs to be constructed, which can be performed in $\mathcal{O}(mp^2)$ time. Then, the hypervolume of $A(\mathcal{S}_p)$ needs to be computed. In the case of $m = 2$ two objectives, this can be done by simply sorting the solutions, resulting in a computational complexity of $\mathcal{O}(p \log p)$ time. Finally, using the sorted approximation set, computing the uncrowded distances can be performed in $\mathcal{O}(mp)$ time. This results in an overall computational complexity of $\mathcal{O}(mp^2)$ for computing the UHV with $m = 2$ objectives. Note however that p is typically rather small here, (e.g., $p = 9$). Note further that we did not yet explore potential speedups by considering that the hypervolume measure does not have to be computed from scratch every generation, but can be updated.

In contrast, "Fast Non-Dominated Sorting" in NSGA-II with a population of N MO-solutions can be performed in $\mathcal{O}(mN^2)$ time [53]. There, N is however

typically a magnitude larger (e.g., $N = 100$). Also, maintaining an elitist archive of size $N_{\mathcal{E}} \leq 1000$ has a computation complexity of $\mathcal{O}(mN_{\mathcal{E}})$, and is essential for the performance of MO-GOMEA [38].

To compute the hypervolume with $m \geq 3$ objectives, different algorithms have been proposed [20, 69], yielding a computational complexity of $\mathcal{O}(p^{m-2} \log p)$, which becomes the dominating term when m is large. Additionally, an efficient algorithm to compute the uncrowded distance for $m > 2$ has still to be derived.

6.3 Sofomore-GOMEA

While UHV-GOMEA optimizes the UHV by manipulating sets of MO-solutions in a population-based approach, the Sofomore framework [195] was introduced to iteratively optimize the UHVI one MO-solution at a time. Specifically, Sofomore performs a search around a single solution set $\mathcal{S}_p$ of fixed size p. For each MO-solution $\mathbf{x}_i \in \mathcal{S}_p$, a single-objective optimizer is initialized that solves the single-objective dynamic optimization problem,

$$h_i : \mathcal{X} \to \mathbb{R}, \text{ given by, } h_i(\mathbf{x}_i|\mathcal{S}_p \backslash \{\mathbf{x}_i\}) = \text{UHVI}(\mathbf{x}_i, \mathcal{S}_p \backslash \{\mathbf{x}_i\}).$$

Note that the decision space of h_i and the multi-objective problem given by $\mathbf{f}$ are the same, and an h_i-solution is thus also an MO-solution, but with different fitness value assignment. Steps of individual optimizers are then interleaved. Intuitively, each optimizer in turn aims to replace $\mathbf{x}_i$ by the solution with maximal hypervolume contribution with respect to $\mathcal{S}_p$, while keeping the other MO-solutions in $\mathcal{S}_p$ fixed.

To solve the p dynamic single-objective optimization problems given by h_i, we again use GOMEA to get a fair comparison basis for comparing different optimization approaches. Our implementation largely agrees with the combination of Sofomore and CMA-ES (called COMO-CMA-ES), as presented in [195], but a few minor changes were made that we discuss here. We refer the reader to [195] for further details on Sofomore and COMO-CMA-ES.

We assumed that the multi-objective problem is black-box, and we therefore use a full linkage model for Sofomore-GOMEA, as h_i cannot be partially evaluated. We maintain the default uniform initialization of GOMEA. In this way, Sofomore-GOMEA and HV-GOMEA are initialized from exactly the same set of MO-solutions (when the same random seed is used). Additionally, and most important, the Sofomore approach results in a set of *dynamic* optimization problems. As GOMEA accepts only improvements to solutions, it is beneficial to recompute all fitness values (i.e., h_i) at the beginning of each generation, as they might have

changed since the last generation. Note that this requires $h_i(\mathbf{x}|\mathcal{S}_p\backslash\{\mathbf{x}_i\})$ to be updated, but there is no need to re-evaluate $\mathbf{f}(\mathbf{x}_i)$ itself.

We evaluate performance based on the number of evaluations of the multi-objective problem (MO-fevals), which is thus unaffected by this re-evaluation, although overall computation time will increase. This re-evaluation is not required for UHV-GOMEA, as it does not solve a dynamic problem. Finally, we use the same archive for Sofomore-GOMEA as described in Section 6.2.2.

6.4 MO-GOMEA

MO-GOMEA is a domination-based MOEA [38]. MO-GOMEA optimizes a population of MO-solutions that is aimed to approximate the Pareto front by balancing diversity and proximity. Besides the main population, MO-GOMEA maintains an elitist archive $\mathcal{E}$, as described in Section 6.2.2. In MO-GOMEA, MO-solutions are copied back from the archive into the population. It can therefore be roughly said that the aim of MO-GOMEA is to obtain an elitist archive that approximates the Pareto front as good as possible. We discuss the main characteristics of MO-GOMEA here, and refer the reader to [38] for a full description of the algorithm.

Again, as we assumed the multi-objective problem to be black-box, we use a full linkage model for MO-GOMEA. From a population of N_{mo} MO-solutions, truncation selection is performed based on domination rank, resulting in a selection of size τN. This selection is clustered into K_{mo} clusters, each of size $2\tau N/K_{mo}$. Each cluster models a part of the approximation front, and for this an objective-space based clustering method is used that guarantees overlapping clusters of equal size. For each cluster, a Gaussian distribution is estimated to sample new MO-solutions from. Similar to the single-objective GOMEA, MO-GOMEA only accepts improvements, meaning that offspring need to either dominate the parent, or be accepted into the elitist archive.

To align MO-GOMEA with the other algorithms, we set $N_{mo} = p \cdot N$ and $K_{mo} = 2p$ such that the overall number of MO-solutions in the populations is the same, and all sample distributions are estimated from the same number of MO-solutions. MO-GOMEA estimates its sample distributions in a similar fashion as the single-objective GOMEA which was used in UHV-GOMEA and Sofomore-GOMEA. This makes a comparison between these three approaches most fair. Finally, to be able to compare the limited-size $|\mathcal{S}_p| = p$ of UHV-GOMEA and Sofomore-GOMEA, we perform greedy Hypervolume Subset Selection (gHSS) [80] to select p solutions from the elitist archive $\mathcal{E}$.

6.4.1 A hybrid method

MO-GOMEA is expected to perform better initially, but to stagnate in terms of proximity to the Pareto set when the majority of MO-solutions in the population is non-dominated. We construct a simple hybrid approach where we initially run MO-GOMEA, which we terminate when it stagnates, i.e., when 90% of the MO-solutions in the population are non-dominated, or when the elitist archive target size is hit. We then switch to UHV-GOMEA-Lm starting from the elitist archive $\mathcal{E}$ that MO-GOMEA obtained so far. $\mathcal{E}$ is clustered into p clusters of equal size $2|\mathcal{E}|/p$ using the same clustering method that is used in MO-GOMEA. For this, the cluster means are initialized with gHSS, and distances are measured in decision space. If the cluster size is smaller than the desired population size, i.e., $2|\mathcal{E}|/p < N$, the remainder of the MO-solutions is sampled uniformly.

6.5 Experiments

Source code for the algorithms in this chapter is available online [140].

6.5.1 Experiment 1: Rate of convergence

In Section 6.1.1, we introduced the bi-sphere problem, which we use to demonstrate the rate of convergence for UHV-GOMEA with the three linkage models, Sofomore-GOMEA, and MO-GOMEA. The bi-sphere problem is a separable problem, defined on the entire real space $\mathcal{X} = \mathbb{R}^n$. Its Pareto set is a straight line between the origin and $\mathbf{e}_1$. Due to its separability, it can be solved with a diagonal variance matrix, and we therefore use a small population size $N = 31$. All algorithms are initialized from the same initial population (for the same random seed) on the domain $[-100, -50]^n$, away from the Pareto set. We use $n = 10$ MO-decision variables with $p = 9$ MO-solutions in a solution set, resulting in an IBMOP with 90 decision variables.

As performance indicator, we use the distance to the optimal hypervolume $\Delta\text{HV}_p = \text{HV}(\mathcal{A}_p^\star) - \text{HV}(\mathcal{A}_p)$, with $\mathcal{A}_p = A(\mathcal{S}_p)$, and $\mathcal{A}_p^\star$ the approximation set of size p with optimal hypervolume [7]. We determined $\text{HV}(\mathcal{A}_p^\star)$ empirically by solving lower-dimensional problem instances with a large computational budget.

As a second performance indicator, we use the Generational Distance (GD) [220],

$$\text{GD}(\mathcal{A}_p, \mathcal{A}^\star) = \frac{1}{|\mathcal{A}_p|} \sum_{\mathbf{x} \in \mathcal{A}_p} \min_{\mathbf{y} \in \mathcal{A}^\star} \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})\|.$$

Since a parametric expression of the Pareto set $\mathcal{A}^\star$ is available for the bi-sphere problem, we can compute the GD analytically. The GD measures the proximity

of $\mathcal{A}_p$ to the Pareto set, but does not take diversity into account. The GD is not Pareto compliant, but since we considered fixed-size approximation sets, it is a useful tool to measure proximity to the Pareto set. We are especially interested in the GD at the end of the optimization process, when all approximation sets are the same size and contain only non-dominated solutions. Since we compute the GD analytically, it does hold that any solution set containing only Pareto optimal solutions has GD $= 0$. Because of this property, the analytic GD is a good measure for proximity to the Pareto set.

As a third performance indicator, we count the non-dominated MO-solutions in $\mathcal{S}_p$, i.e., $|\mathcal{A}_p|$, and visualize the approximation fronts. In all experiments, we measure performance in terms of the number of MO-fevals. All experiments are repeated 30 times, and mean results are shown, unless mentioned otherwise. The hypervolume reference point is set to $r = (11, 11)$ in all experiments, which is far away from the Pareto front, thereby aiming that the endpoints of the Pareto front are in $\mathcal{A}_p^\star$. However, even by setting the reference point this far, this is not always achieved for all problems.

Results

Figure 6.3 shows that all hypervolume-based algorithms exhibit linear convergence in terms of $\log(\Delta\text{HV}_p)$ and $\log(\text{GD})$, albeit at different rates. Sofomore-GOMEA performs best, closely followed by UHV-GOMEA-Lm. Due to the small population size, large linkage elements are filtered from the linkage tree in UHV-GOMEA-Lt, which performs the same as UHV-GOMEA-Lm (up to randomness). UHV-GOMEA-Lf can still solve the problem, but it is inefficient. With full linkage, all MO-solutions are updated and evaluated simultaneously (albeit independently from the diagonal variance matrix, due to the small population size), and only then the corresponding hypervolume is computed. With marginal linkage however, MO-solutions are updated and evaluated one-by-one, and after each newly evaluated MO-solution, it is checked if this improved the corresponding hypervolume, which is beneficial here. In this scenario, where the optimization is initialized far from the Pareto set, the hypervolume-based algorithms are initially fully driven by the uncrowded distance towards the reference point. The uncrowded distance within UHV-GOMEA and Sofomore-GOMEA is effective in obtaining a set of non-dominated solutions, but is not as efficient as MO-GOMEA, which performs best initially in all three measures. The rate of convergence for the hypervolume-based algorithms is constant as soon as $|\mathcal{A}_9| = 9$ for UHV-GOMEA and Sofomore-GOMEA, which shows that linear convergence is due to the hypervolume optimization and not due to the uncrowded distance. As expected, MO-GOMEA

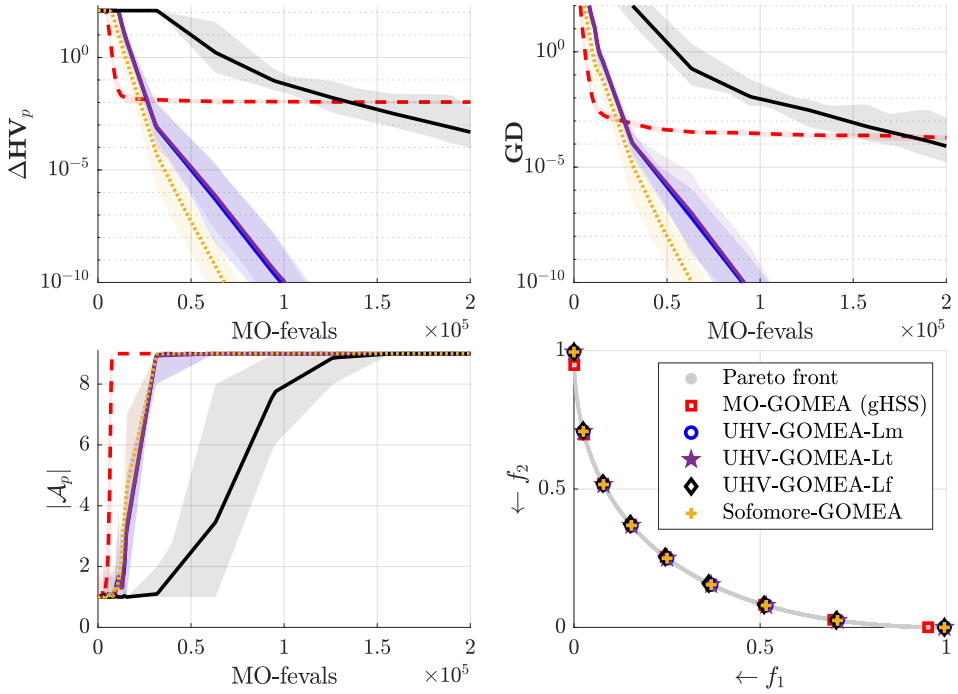

Figure 6.3: Convergence on the bi-sphere problem (with $n = 10$, $p = 9$, and $N = 31$). Lines show mean scores over 30 independent runs, and shaded areas show min/max scores. The objective space plot (bottom right) shows the result of a single run of each algorithm. The GD is computed analytically.

stagnates in terms of ΔHV_p, which happens at $\Delta\text{HV}_p \approx 0.01$. Its elitist archive contains a good distribution of solutions along the front, including solutions close to the endpoints of the Pareto set, but its approximation front has a slightly different distribution of solutions due to the gHSS, which could explain the stagnation in ΔHV_p. However, the analytic GD plot shows that these solutions do not converge to the Pareto set.

6.5.2 Experiment 2: Modeling dependencies

The bi-sphere problem could be solved efficiently with a small population size N and without dependency modeling. We now perform a number of experiments to investigate if there are scenarios where modeling the dependencies between MO-solutions is beneficial. For this, we use two additional benchmark problems, constructed from two well-known single-objective functions,

$$f_{\text{elli}}(\mathbf{x}) = \sum_{i=1}^{n} 10^{6\frac{i-1}{n-1}} x_i, \text{ and, } f_{\text{Rosenbrock}}(\mathbf{x}) = \sum_{i=1}^{n-1} (100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2).$$

These problems are, similar to the bi-sphere problem, defined for a scalable number of decision variables n and on $\mathcal{X} = \mathbb{R}^n$. From these, we construct two bi-objective optimization problems,

$$\begin{aligned} \mathbf{f}_{\text{sphere-rotatedElli}}(\mathbf{x}) &= [f_{\text{sphere}}(\mathbf{x}) \; ; \; f_{\text{elli}}(R\mathbf{x} - \mathbf{e}_1)] \,, \\ \mathbf{f}_{\text{sphere-Rosenbrock}}(\mathbf{x}) &= \left[\frac{1}{n} f_{\text{sphere}}(\mathbf{x}) \; ; \; \frac{1}{n-1} f_{\text{Rosenbrock}}(\mathbf{x})\right] , \end{aligned} \tag{6.5}$$

where R is a rotation matrix that defines a rotation around the origin of $\pi/4$ radians in all principal directions.

The sphere-rotatedElli problem has the same Pareto front as the bi-sphere problem (but a different, non-linear, Pareto set), and has one non-separable ill-conditioned objective. The sphere-Rosenbrock problem (also known as BD2s [31]) has pair-wise dependencies and a non-linear Pareto set. Especially the combination of an easy and a difficult objective can make it difficult to obtain an evenly spread approximation front. All three problems are scaled such that their Pareto fronts have endpoints at $(1, 0)$ and $(0, 1)$. The sphere-rotatedElli problem is initialized on the domain $[-100, -50]^n$, the sphere-Rosenbrock function on $[-5, 5]^n$.

Results: Linkage tree versus random linkage

We compared the linkage tree model, where linkage subsets are formed between neighboring solutions (see Section 6.2.2), with a model in which solutions are merged randomly. Results were compared for 30 runs on sphere-Rosenbrock, with a population size of $N = 200$, $n = 10$ and $p = 9$. The mean number of function evaluations that UHV-GOMEA-Lt required to obtain $\Delta\text{HV}_p < 10^{-10}$ was $1.28{\times}10^6$ (range: 1.02–1.66×10^6).

When we use the same UPGMA method of constructing the linkage tree, but with a random distance between linkage models, the required number of function evaluation increases significantly to 1.59×10^6 (range: 1.35–1.95×10^6), as tested with a Wilcoxon rank-sum test with $\alpha = 0.05$. This shows that the intuitively chosen construction approach for the linkage tree has added value, although better linkage models might of course still exist.

Results: Population size N

The effect of the population size parameter N on the required number of MO-fevals to reach a target accuracy of $\Delta\text{HV}_p < 10^{-5}$, and the corresponding SR is shown in Figure 6.4. The computational budget was set to 10^7 MO-fevals. Performance is rather predictable for $N \geq 200$, but for small population sizes, differences occur. UHV-GOMEA with Lt and Lm perform similar initially, as larger linkage subsets

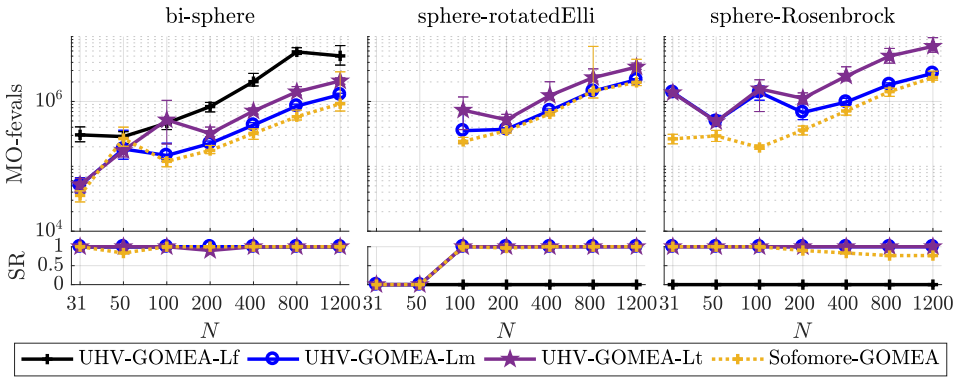

Figure 6.4: Effect of the population size parameter N on the hypervolume-based algorithms for various problems with $n = 10$ and $p = 9$. Success rate (SR) measures the fraction of runs in which the target accuracy ($\Delta\text{HV}_p < 10^{-5}$) was reached, out of 30 runs in total. The top row shows the number of MO-fevals required to reach this accuracy, for all successful runs.

are filtered out. When increasing the population size, UHV-GOMEA-Lt performs slightly worse, although the overhead seems to be a constant factor. This suggests that it is nor beneficial nor harmful to model dependencies in this scenario. UHV-GOMEA-Lf is not able to solve the sphere-rotatedElli, and is significantly slower for the bi-sphere problem. We therefore omit it from further experiments.

Sofomore-GOMEA performs similar to UHV-GOMEA-Lm for larger population sizes, but for smaller population sizes, it clearly outperforms all UHV-GOMEA variants on sphere-Rosenbrock. A relatively large population size is required to estimate a full covariance matrix, which is required to solve sphere-rotatedElli. For smaller population sizes, a diagonal variance matrix or regularized covariance matrix is estimated, which affects performance, making it hard to predict whether it is beneficial to increase or decrease the population size.

Interestingly, for large N, the success rate of Sofomore-GOMEA deteriorates on sphere-Rosenbrock, as it converges to a locally optimal distribution of solutions along the front, which can only be escaped if multiple solutions were to be updated simultaneously. This only happens for sphere-Rosenbrock due to the shape of its Pareto front, but also since the population is initialized close to the Pareto set. Since f_1 is easier to solve than f_2, solutions initially quickly converge towards one end the front, and the larger the population size, the faster this happens, resulting in little to no opportunity for each of the optimizers to adapt for the dynamic nature of its objective function.

Table 6.1: Success ratio (SR) to obtain $\Delta\text{HV}_p < 10^{-10}$, together with the MO-fevals per p of successful runs (± standard deviation) for the sphere-rotatedElli benchmark problem ($n = 3$), for $N = 50$ and $N = 100$. Bold scores are best obtained scores or those not statistically different from it.

		UHV-GOMEA-Lm		UHV-GOMEA-Lt		Sofomore-GOMEA	
	p	SR	MO-fevals/p	SR	MO-fevals/p	SR	MO-fevals/p
$N = 50$	3	***1.00***	3.3e+03±1.8e+02	***1.00***	5.2e+03±7.2e+02	***1.00***	**2.9e+03**±3.2e+02
	5	***1.00***	4.3e+03±2.1e+02	***1.00***	9.1e+03±1.7e+03	***1.00***	**3.6e+03**±4.1e+02
	9	***1.00***	**7.2e+03**±5.8e+02	***1.00***	1.6e+04±2.1e+03	***1.00***	7.5e+03±3.0e+02
	17	*0.73*	**2.4e+04**±7.5e+03	*0.90*	8.2e+04±1.9e+04	***1.00***	**2.2e+04**±7.0e+02
	33	*0.00*	-	***0.30***	**3.2e+05**±1.6e+05	*0.00*	-
	65	*0.00*	-	*0.00*	-	*0.00*	-
$N = 100$	3	***1.00***	6.5e+03±2.4e+02	***1.00***	8.5e+03±5.9e+02	***1.00***	**5.3e+03**±4.8e+02
	5	***1.00***	8.3e+03±2.2e+02	***1.00***	1.3e+04±9.9e+02	***1.00***	**6.1e+03**±4.2e+02
	9	***1.00***	**1.3e+04**±6.4e+02	***1.00***	2.4e+04±2.8e+03	***1.00***	1.4e+04±8.5e+02
	17	*0.63*	**3.5e+04**±7.3e+03	***1.00***	4.0e+04±3.5e+03	***1.00***	4.1e+04±1.7e+02
	33	*0.00*	-	***0.97***	**1.5e+05**±3.9e+04	*0.23*	**1.4e+05**±4.0e+03
	65	*0.00*	-	***0.97***	**8.3e+05**±2.4e+05	*0.00*	-

Results: Solution set size p

When the size p of the solution set $\mathcal{S}_p$ is set larger, resulting MO-solutions will be closer to each other on the front. It can thus be expected that for larger p, dependency modeling becomes more relevant. We let $p = 2^j + 1$ with $j \in \mathbb{N}$, and inspect how many runs obtain a high accuracy of $\Delta\text{HV}_p < 10^{-10}$. For this, we let each algorithm run with a large budget of 10^8 MO-fevals, or until it converged, i.e., when the standard deviation of the objective values in the population is less than 10^{-20}. At this point, the machine accuracy becomes an issue, and no further improvements can be obtained. We use the sphere-rotatedElli problem with $n = 3$, as it has dependencies, with the same Pareto front as the bi-sphere problem, so that we can find the empirical $\Delta\text{HV}(\mathcal{A}_p^\star)$ by solving this simpler problem with accuracy of 10^{-15}. Success rates and number of MO-fevals to reach the target accuracy are shown in Table 6.1. Differences were tested for statistical significance using the Wilcoxon rank-sum test at $\alpha = 0.05$.

For $N = 50$ and $p \leq 9$, all algorithms obtain the target accuracy in all runs, but for $p \geq 17$, performance deteriorates for UHV-GOMEA-Lm and Sofomore-GOMEA, while UHV-GOMEA-Lt can still solve the problem. Doubling the population size roughly doubles the required number of MO-fevals for all algorithms for $p \leq 9$. For UHV-GOMEA-Lt, increasing the population size allows it to solve problems with large p, clearly outperforming the other algorithms.

The number of function evaluations GOMEA requires to solve grey-box problems with independent decision variables scales logarithmically in the problem dimen-

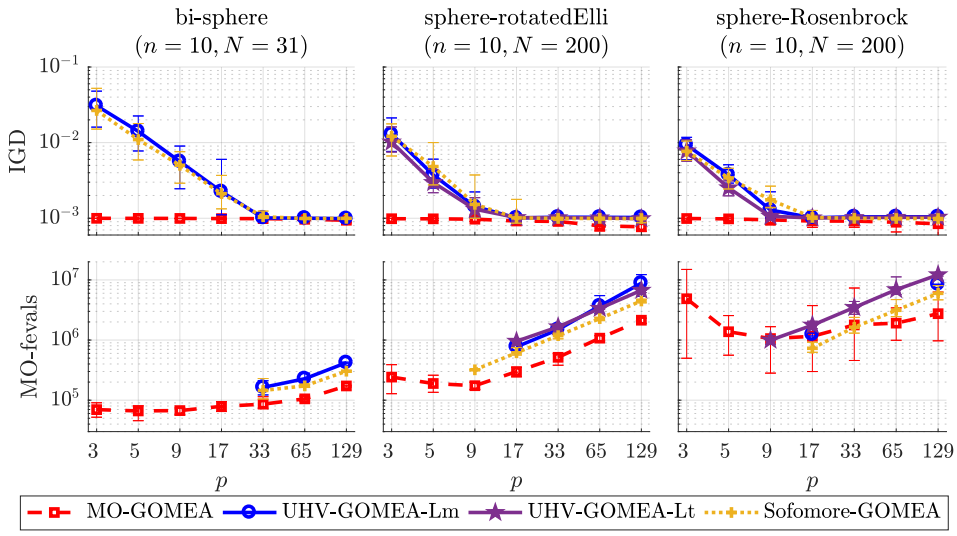

Figure 6.5: Inverted Generational Distance (IGD) and number of MO-fevals to reach IGD $< 10^{-3}$ in case of success.

sionality [37], which is for the IBMOP determined by p and n. Here, we see that the required number of MO-fevals grows super-linear in p. This indicates that the strength of the dependencies indeed increases with p, up to the point that UHV-GOMEA-Lm and Sofomore-GOMEA can no longer solve the problem, making dependency modeling essential.

6.5.3 Experiment 3: Elitist archive

In Experiment 1, the limited-size $\mathcal{S}_p$ obtained by the different algorithms was the basis of comparison together with a notion of proximity of each MO-solution to the Pareto set, which was favorable for the hypervolume-based algorithms. We now compare the elitist archives $\mathcal{E}$. All algorithms have the same archiving strategy, but only MO-GOMEA uses the archive as part of the optimization process. As performance indicator, we use the Inverted Generational Distance (IGD) [32],

$$\mathrm{IGD}(\mathcal{E}, \mathcal{A}^\star) = \frac{1}{|\mathcal{A}^\star|} \sum_{\mathbf{y} \in \mathcal{A}^\star} \min_{\mathbf{x} \in \mathcal{A}} \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{y})\|,$$

where we use a finite approximation of the Pareto set $\mathcal{A}^\star \approx \mathcal{A}^\star_{5000}$ of 5000 MO-solutions. The IGD measures both proximity and diversity, in contrast to the GD, which only measures proximity. We run each algorithm until a target accuracy of IGD $< 10^{-3}$ is obtained, or when it converged before.

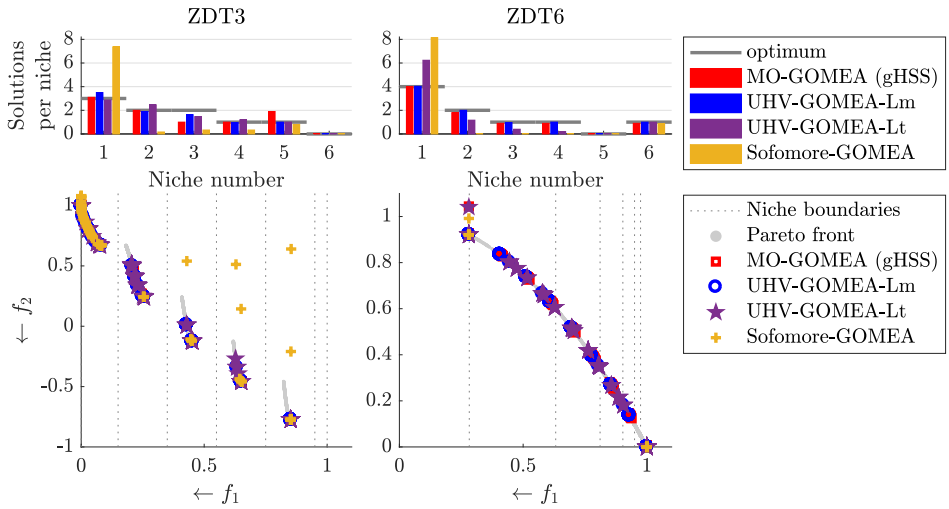

Figure 6.6: Distribution of solutions over different niches in the final solution sets $\mathcal{S}_p$ for the multimodal functions ZDT3 and ZDT6 ($n = 10$, $p = 9$, $N = 200$). Niche numbers correspond to the objective space regions separated by the dotted lines.

Results

Results are shown in Figure 6.5. Independent of the choice of p, MO-GOMEA was able to obtain the target IGD. Performance seems rather robust against the choice of p, which, in our experiment setup controls its population size $N_{mo} = p \cdot N$ and number of clusters $K_{mo} = 2p$. Furthermore, we see that MO-GOMEA outperforms the other algorithms in all cases by obtaining a better IGD with fewer MO-fevals, except for some runs of the sphere-Rosenbrock problem with $p = 17$ and $p = 33$. In terms of p, log(IGD) decreases linearly for the hypervolume-based algorithms, and for $p \geq 33$, all algorithms obtain the target IGD, for all three problems.

6.5.4 Experiment 4: Multimodal multi-objective problems

As the Sofomore framework performs a form of local search around a single approximation set, it is as expected that it performs well on unimodal functions. Real-world problems are often not that well-behaved. Therefore, we now include two multimodal problems from the ZDT problem set [54]. ZDT3 is characterized by a discontinuous Pareto front. ZDT6 has a concave front, and six local optima in f_1, resulting in a Pareto set consisting of six subsets that partially overlap.

The decision space of the ZDT problems is bounded to $[0, 1]^n$. We here use re-sampling as repair mechanism. We furthermore set $p = 9$ and therefore use $N = 200$, as this was found to be a good choice for all algorithms based on the results in Section 6.5.2.

Results

Figure 6.6 shows the distribution of MO-solutions per niche. UHV-GOMEA-Lm obtains the optimal hypervolume for all runs on ZDT6 and for 60% of the runs on ZDT3. MO-GOMEA obtains an elitist archive with MO-solutions in all niches, but gHSS does not result in the optimal distribution over niches. Sofomore-GOMEA obtains many more MO-solutions in the left-most niche (Niche 1). It quickly converges to the extremes of the Pareto front in the beginning of the optimization, but is not able to move MO-solutions along the front. For ZDT3, a good spread of MO-solutions within the left-most niche is obtained, but it is not able to move MO-solutions out of that local optimum to the other subsets of the Pareto front in most runs. When it did not have MO-solutions in one of the subsets of the front, it sometimes ended up with non-optimal MO-solutions within the niches that were obtained in an attempt to fill the resulting gaps in the front. For ZDT6, Sofomore-GOMEA shows the same behavior, but it is not able to move MO-solutions along the front due to the concavity of the front, and only obtains MO-solutions at the extremes of the Pareto front.

UHV-GOMEA-Lm obtains $\Delta\text{HV}_p < 10^{-5}$ in 60% of the runs for ZDT3 and 100% for ZDT6, while UHV-GOMEA-Lt obtains this accuracy 26% of the runs for ZDT3, and 0% for ZDT6. The other methods were not able to reach this accuracy in any of the runs.

6.5.5 Experiment 5: A hybrid approach

As seen in Figure 6.3, MO-GOMEA performs best initially, but stagnates in terms of proximity (e.g., GD) when the majority of MO-solutions in the population is non-dominated. We construct a simple hybrid approach where we initially run MO-GOMEA, which we terminate when it is expected to stagnate, i.e., when 90% of the MO-solutions in the population are non-dominated, or when the elitist archive target size is hit.

We then switch to UHV-GOMEA-Lm starting from the elitist archive $\mathcal{E}$ that MO-GOMEA obtained so far. $\mathcal{E}$ is clustered into p clusters of equal size $2|\mathcal{E}|/p$ using the same clustering method that is used in MO-GOMEA. For this, the cluster means are initialized with gHSS, and distances are measured in decision space. If the cluster size is smaller than the desired population size, i.e., $2|\mathcal{E}|/p < N$, the remainder of the MO-solutions is sampled uniformly.

Results

Convergence results for the discussed algorithms are shown in Figure 6.7. As intended, the hybrid approach terminates MO-GOMEA when stagnation occurs

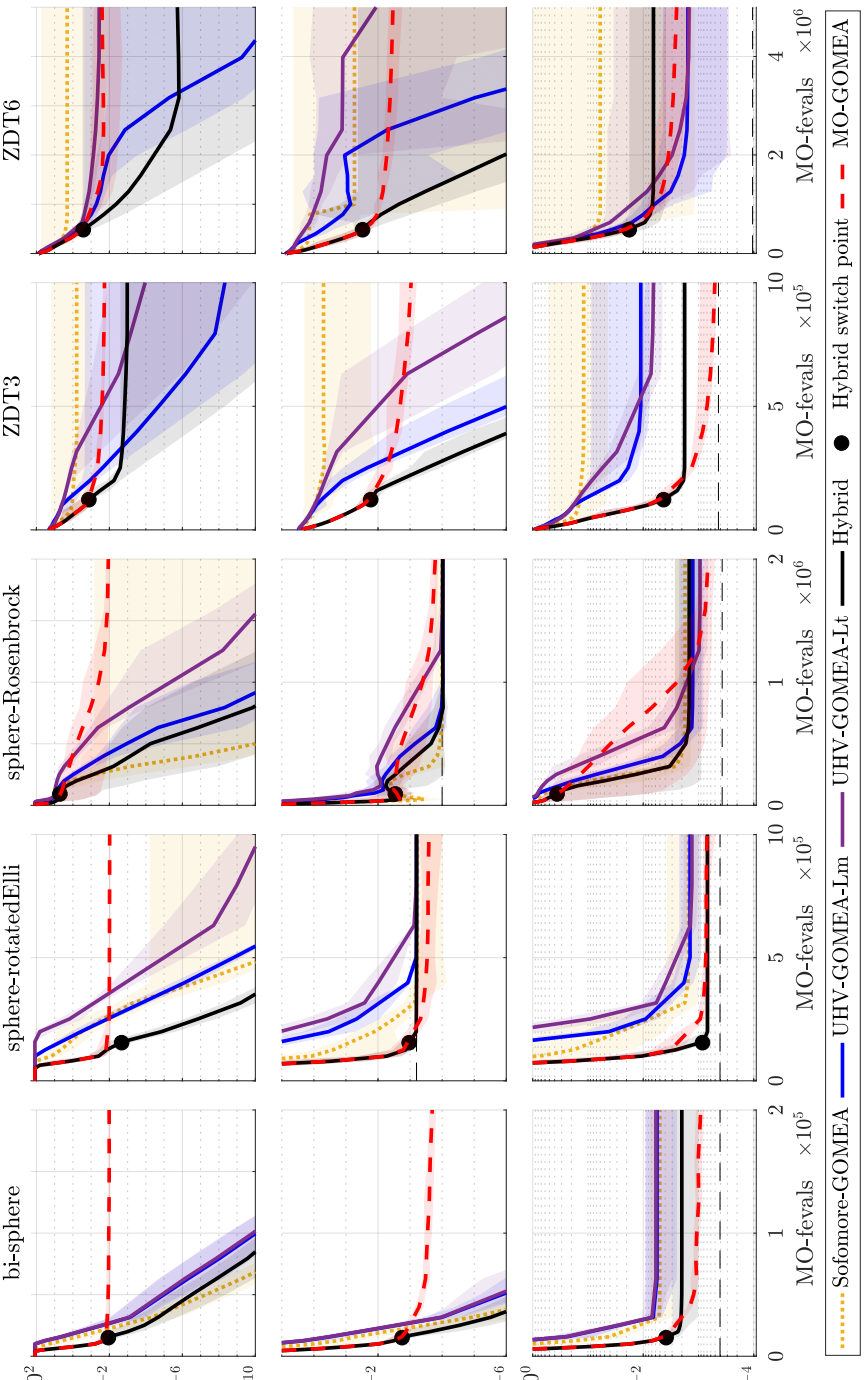

Figure 6.7: Mean scores (shaded areas are min/max) for the discussed algorithms. All problems are $n = 10$, $p = 9$, and run with $N = 200$, except the bi-sphere problem, which run with $N = 31$. Hybrid switch point shows the mean MO-fevals after which MO-GOMEA was terminated and UHV-GOMEA-Lm initialized in the hybrid approach. Dashed black lines shows the maximally achievable GD scores for problems where a finite approximation of the Pareto front is used, and the maximally achievable IGD with a limited archive size $|\mathcal{E}| \leq 1000$ for the IGD.

and UHV-GOMEA-Lm takes over from there (indicated by the black dot). The hybrid converges to the Pareto set in all cases, as shown by GD($\mathcal{A}_p$), and outperformed UHV-GOMEA-Lm, showing that a domination-based initialization is preferable over the uncrowded distance approach. In terms of the archive IGD, the hybrid approach also outperforms UHV-GOMEA-Lm for all problems but ZDT6, where none of the algorithms obtain the maximally achievable IGD.

6.5.6 Experiment 6: WFG benchmark

We benchmark the discussed methods on the commonly used WFG Benchmark [101]. This test suite consists of 9 benchmark functions with different properties. We consider the instances with $m = 2$ objectives, $k_{\text{WFG}} = 4$ position variables, and $l_{\text{WFG}} = 20$ distances variables, resulting in a total of $n = 24$ decision variables. The hypervolume reference point is set to $r = (11, 11)$. Of these problems, WFG1 is separable, but has a flat region in the decision space, which could cause stagnation. WFG2, WFG4, and WFG9 have one or more multimodal objectives. Problems WFG4–9 all have a concave front, WFG1 has a convex front, WFG2 has a disconnected convex front, and WFG3 has a linear front.

We solve these benchmark problems with $p = 9$ and a computational budget of 10^7 MO-fevals. Based on previous results, a population size of $N = 200$ was used for all algorithms. We include two versions of the hybrid method in this experiment. Both first run MO-GOMEA, and Hybrid-Lm then switches to UHV-GOMEA-Lm, while Hybrid-Lt switches to UHV-GOMEA-Lt. All experiments are repeated 30 times. Differences are tested for statistical significance (up to 4 decimals) by a Wilcoxon rank-sum test with $\alpha = 0.05$, pairwise to the best. Ranks (in brackets) are computed based on the mean hypervolume values. All statistics are computed per table.

Results

Results on the WFG benchmark are shown in Table 6.2. Problem WFG1, which has a plateau in its fitness landscape, is consistently solved better with MO-GOMEA. MO-GOMEA als outperforms the other methods on WFG2, which is multimodal, and has a discontinuous front, which prevents the hypervolume-based methods from obtaining the optimal distribution of solutions, as was shown before for the ZDT3 and ZDT6 problems. For the problems with a concave front, WFG4–9, the optimal hypervolume value is $\text{HV}_9^\star = 114.40\ldots$, which was obtained by many of the algorithms. Especially WFG3 which has optimal hypervolume $\text{HV}_9^\star = 116.50\ldots$, and WFG6 and WFG7 seem to be relatively easy, with none of the methods

Table 6.2: Results on the WFG Benchmark with $p = 9$ MO-solutions, resulting in an $n = 216$ dimensional optimization problem. A computational budget of 10^7 MO-fevals was used. Hypervolume values HV_p are shown (mean, $\pm$ standard deviation (rank)), computed with reference point $r = [11, 11]$. gHSS was used for MO-GOMEA. Bold numbers are best scores per problems, or those not statistically different from it. Bottom row shows the mean rank, together with the overall rank in brackets.

Problem	Sofomore-GOMEA	UHV-GOMEA-Lm	UHV-GOMEA-Lt
1	96.94 ± 1.74 (5)	93.98 ± 0.88 (6)	99.22 ± 2.43 (2)
2	110.16 ± 0.03 (5)	110.14 ± 0.01 (6)	110.21 ± 0.40 (2)
3	116.50 ± 0.00 (3)	**116.50** ± 0.01 (5)	**116.50** ± 0.00 (1)
4	112.81 ± 0.60 (5)	112.66 ± 0.64 (6)	113.23 ± 0.41 (2)
5	**112.27** ± 0.13 (1)	**112.08** ± 0.38 (6)	**112.22** ± 0.00 (2)
6	**114.40** ± 0.00 (1)	114.39 ± 0.02 (4)	114.39 ± 0.02 (3)
7	**114.40** ± 0.00 (3)	**114.40** ± 0.00 (2)	**114.40** ± 0.00 (5)
8	111.74 ± 0.20 (3)	111.47 ± 0.29 (5)	111.51 ± 0.23 (4)
9	111.44 ± 0.20 (6)	111.49 ± 0.05 (4)	111.48 ± 0.04 (5)
mean rank	3.56 (3)	4.89 (6)	2.89 (2)

Problem	MO-GOMEA	Hybrid-Lm	Hybrid-Lt
1	**103.52** ± 1.72 (1)	98.00 ± 0.88 (3)	98.00 ± 0.88 (4)
2	**112.28** ± 3.32 (1)	110.18 ± 0.00 (4)	110.18 ± 0.00 (3)
3	116.34 ± 0.04 (6)	**116.50** ± 0.00 (4)	**116.50** ± 0.00 (2)
4	113.19 ± 0.63 (3)	112.97 ± 0.56 (4)	**113.50** ± 0.55 (1)
5	112.17 ± 0.04 (4)	**112.14** ± 0.28 (5)	**112.22** ± 0.00 (3)
6	114.17 ± 0.08 (6)	114.29 ± 0.38 (5)	114.39 ± 0.01 (2)
7	114.26 ± 0.04 (6)	**114.40** ± 0.00 (4)	**114.40** ± 0.00 (1)
8	111.11 ± 0.15 (6)	111.80 ± 0.04 (2)	**111.83** ± 0.01 (1)
9	**112.09** ± 0.57 (1)	111.57 ± 0.25 (2)	111.57 ± 0.25 (3)
mean rank	3.78 (5)	3.67 (4)	**2.22 (1)**

performing particularly worse than the others. Differences occur often at more than two decimals of accuracy, and are therefore not visible in the table.

This experiment setup is unfavorable for MO-GOMEA, as it is not aimed to obtain the optimal distribution of exactly $p = 9$ solutions along the front. Despite this bias, it is still very competitive, obtaining the best scores in three of the problems. Especially the difference between UHV-GOMEA-Lm and UHV-GOMEA-Lt is noteworthy, demonstrating that taking linkage into account greatly improves performance on these more difficult problems. Overall, Hybrid-Lt performs best, although differences between different algorithms are small.

6.6 Discussion

We formulated the Uncrowded Hypervolume (UHV) measure in this chapter, which was used to achieve population-based hypervolume-driven MO optimization using a single-objective problem formulation. We compared this problem formulation to the dynamic interleaved Sofomore framework, which is also hypervolume-based, and the MO problem formulation based on Pareto-dominance that is typically used in MOEAs. These three problem formulations were all solved with versions of GOMEA [37, 38], for a modern and fair comparison.

We showed that the hypervolume-based methods do not exhibit the stagnation (e.g., in GD) that occurs with domination-based MOEAs, and thereby confirm the results obtained in e.g. [195]. This clearly shows the superiority of hypervolume-based methods when a small number of high-quality solutions is required. However, domination-based MOEAs initially outperform the hypervolume-based methods, especially when the initial population is far away from the Pareto set. A simple hybrid approach, in which a reasonably good approximation set that is obtained with a domination-based MOEA is used as the initial population of a hypervolume-based algorithm, showed to improve performance compared to the use of both approaches separately. When a large approximation set is required, the difficulty of the hypervolume-based problems increases, and dependency modeling becomes beneficial or even essential for solving them. Additionally, in our experiments, the domination-based MOEA almost always achieves a better elitist archive containing a large number of non-dominated solutions (e.g., $|\mathcal{E}| = 1000$ was used here).

MO-GOMEA naturally has diversity-enhancing mechanisms, which might be beneficial for successfully optimizing multimodal problems. The single-objective GOMEA, used here to optimize the hypervolume-based problem formulation, was not particularly developed for this objective function. A better understanding of multi-objective fitness landscapes, of which a first attempt was made in [107], might be helpful to adapt GOMEA, or any other single-objective optimizer, for this specific optimization task.

A limitation of the hypervolume-based approach is that a reference point is required, for which a suitable choice could be unknown in a black-box setting. Additionally, the computational complexity of the hypervolume increases when the number of objectives increases. This makes the UHV expensive for multi-objective optimization problems with $m \geq 3$, although the population size in this application is small, and approximation methods could be used [12, 67]. On the other hand, our IBMOP formulation allows multi-objective problems to be solved with a single-objective optimizer, which provides opportunities to explore

techniques such as multi-objective multimodal optimization as in Chapter 5, that are well explored for single-objective optimization, but are still upcoming in multi-objective optimization.

6.7 Conclusion

We introduced a single-objective problem formulation for multi-objective optimization based on the Uncrowded Hypervolume (UHV). We showed that problems formulated as such can be efficiently solved with GOMEA by exploiting grey-box properties of this problem formulation. We compared the resulting approach with a version of GOMEA that is based on a classical domination-based selection (MO-GOMEA) and a version that is based on hypervolume optimization (Sofomore-GOMEA). We showed that hypervolume-based optimization can overcome the stagnation from which domination-based methods suffer after a while, and that these methods show convergence to the optimal hypervolume, and thereby to a subset of the Pareto set. However, when the multi-objective problem at hand has difficult landscape features such as multimodality or deceptiveness, the domination-based MO-GOMEA outperformed the hypervolume based methods.

When the desired approximation set size is small, hypervolume-based methods are generally preferable. When the desired approximation set size is large, domination-based methods obtain a better approximation faster. Additionally, in the latter case, the resulting single-objective optimization problem becomes difficult, and dependency modeling becomes essential to still be able to solve the hypervolume-problem up to high accuracy. Hybrid methods, such as the one proposed in this chapter, stand the best chance at achieving the overall best performance and being most generally applicable, which also provides a promising area of future research.

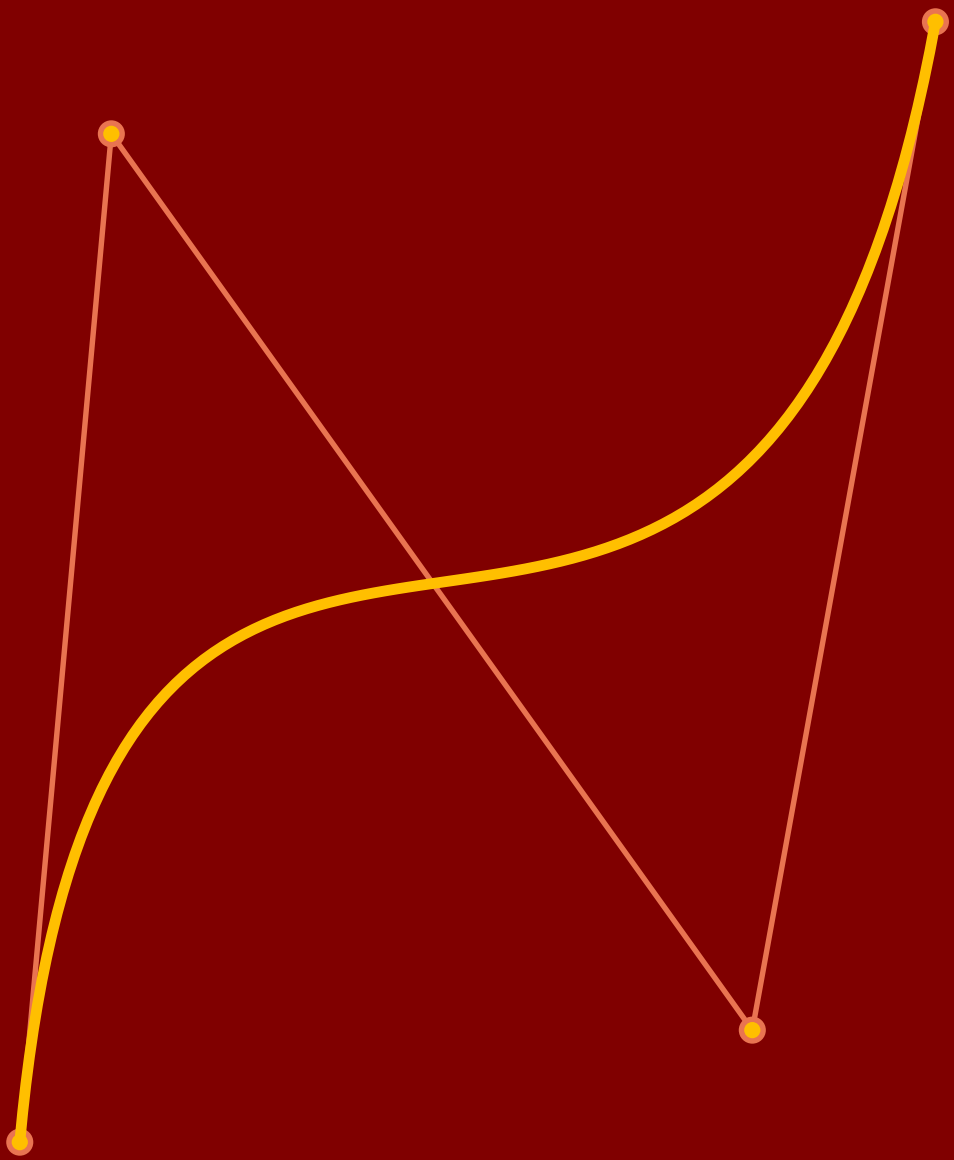

7

Ensuring smooth navigability of approximation sets

Research question 6:

How can smoothly navigable approximation sets be obtained efficiently and effectively via hypervolume-based bi-objective optimization, thereby making the selection of a desirable solution more intuitive for a decision maker?

Research question 7:

Are smoothly navigable approximation sets for bi-objective treatment planning of prostate brachytherapy obtainable without too much cost in obtained trade-offs and computation time?

This chapter is an adaptation of S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2020) *Ensuring Smoothly Navigable Approximation Sets by Bézier Curve Parameterizations in Evolutionary Bi-objective Optimization.* In Proceedings of the Parallel Problem Solving from Nature Conference, 215–228, Springer, Cham.

Abstract

The aim of bi-objective optimization is to obtain an approximation set of (near) Pareto optimal solutions. A decision maker then navigates this set to select a final desired solution, often using a visualization of the approximation front. The front provides a navigational ordering of solutions to traverse, but this ordering does not necessarily map to a smooth trajectory through decision space. This forces the decision maker to inspect the decision variables of each solution individually, potentially making navigation of the approximation set unintuitive. In this work, we aim to improve approximation set navigability by enforcing a form of smoothness or continuity between solutions in terms of their decision variables. Imposing smoothness as a restriction upon common domination-based multi-objective evolutionary algorithms is not straightforward. Therefore, we use the recently introduced Uncrowded Hypervolume (UHV) to reformulate the multi-objective optimization problem as a single-objective problem in which parameterized approximation sets are directly optimized. We study here the case of parameterizing approximation sets as smooth Bézier curves in decision space. We approach the resulting single-objective problem with the Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA), and we call the resulting algorithm BezEA. We analyze the behavior of BezEA and compare it to optimization of the UHV with GOMEA as well as the domination-based multi-objective GOMEA. We show that high-quality approximation sets can be obtained with BezEA, sometimes even outperforming the domination- and UHV-based algorithms, while smoothness of the navigation trajectory through decision space is guaranteed.

7.1 Introduction

The aim of multi-objective optimization is to obtain a set of solutions that is as close as possible to the set of Pareto-optimal solutions, with different trade-offs between the objective functions. A decision maker can then navigate the obtained set, called the *approximation set*, to select a desired solution. The decision maker often incorporates external factors in the selection process that are not taken into account in the optimization objectives. An inspection of the decision variables of individual solutions is therefore required to determine their desirability. To guide the selection in bi-objective optimization, a visualization of the *approximation front*

(i.e., the approximation set mapped to objective space) or trade-off curve can be used. The approximation front then intuitively implies a navigational order of solutions by traversing the front from one end to the other. However, solutions with similar objective values could still have completely different decision values. The decision values of all solutions then need to be inspected individually and carefully because they may not change predictably when the approximation front is traversed. This could make navigation of the approximation set unintuitive and uninsightful.

Population-based Multi-Objective Evolutionary Algorithms (MOEAs) have successfully been applied to real-world black-box optimization problems, for which the internal structure is unknown, or too complex to exploit efficiently by direct problem-specific design [38, 52, 218]. However, imposing a form of smoothness or continuity in terms of decision variables between solutions in the approximation set as a restriction upon the population of MOEAs is not straightforward. An underlying requirement to do so is that control over approximation sets as a whole is needed. However, typical dominance-based EAs use single-solution-based mechanics. Alternatively, multi-objective optimization problems can be formulated as a higher-dimensional single-objective optimization problem by using a quality indicator that assigns a fitness value to approximation sets. An interesting quality indicator is the hypervolume measure [219], as it is currently the only known Pareto-compliant indicator, meaning that an approximation set of given size with optimal hypervolume is a subset of the Pareto set [68, 112, 220]. However, the hypervolume measure has large drawbacks when used as quality indicator in indicator-based optimization, as it does not take dominated solutions into account. The uncrowded distance has been introduced to overcome this [63, 195], which then resulted in the Uncrowded Hypervolume (UHV) measure (Chapter 6). The UHV can be used directly as a quality indicator for indicator-based multi-objective optimization. To be able to optimize approximation sets in this approach, fixed-size approximation sets are parameterized by concatenating the decision variables of a fixed number of solutions [21, 142, 204]. A single-objective optimizer can then be used to directly optimize approximation sets. The resulting single-objective optimization problem is however rather high-dimensional. To efficiently solve it, the UHV-based Gene-pool Optimal Mixing Evolutionary Algorithm (UHV-GOMEA) (Chapter 6), exploits grey-box properties of the UHV problem by only updating a subset of the decision variables corresponding to one (or a few) multi-objective solutions.

In this chapter, we go beyond an unrestricted concatenation of the decision variables of solutions and we propose to model approximation sets as sets of points that lie on a Bézier curve [72] in decision space. Optimizing only the control

points of the Bézier curve, that define its curvature, enforces the decision variables of solutions in the approximation set to vary in a smooth, continuous fashion, thereby likely improving intuitive navigability of the approximation set. Previous work on parameterizations of the approximation set has been applied mainly in a post-processing step after optimization, or was performed in the objective space [23, 114, 152], but this does not aid in the navigability of the approximation set in decision space. Moreover, fitting a smooth curve through an already optimized set of solutions might result in a bad fit, resulting in a lower-quality approximation set. Additionally, we will show that specifying solutions as points on a Bézier curve directly enforces a form of diversity within the approximation set, which can actually aid in the optimization process, and furthermore reduces the problem dimensionality of the single-objective problem.

The remainder of this chapter is organized as follows. In Section 7.2, we introduce preliminaries on UHV-based multi-objective optimization. In Section 7.3, we define a measure for navigational smoothness of approximation sets. In Section 7.4, we introduce Bézier curves and the corresponding optimization problem formulation. Empirical benchmarking on a set of benchmark problems is performed in Section 7.5. Additionally, in Section 7.6, we demonstrate BezEA on a real-world optimization problem that arises in the treatment of prostate cancer with brachytherapy. Finally, we conclude in Section 7.7 with a discussion and outlook.

7.2 UHV-based multi-objective optimization

Let $\mathbf{f} : \mathcal{X} \to \mathbb{R}^m$ be a to-be-minimized m-dimensional vector function and $\mathcal{X} \subseteq \mathbb{R}^n$ be the n-dimensional (box-constrained) decision space. When the objectives in $\mathbf{f}$ are conflicting, no single optimal solution exists, but the optimum of $\mathbf{f}$ can be defined in terms of *Pareto optimality* [113]. We say that a solution $\mathbf{x} \in \mathcal{X}$ *weakly dominates* another solution $\mathbf{y} \in \mathcal{X}$, written as $\mathbf{x} \preceq \mathbf{y}$, if and only if $f_i(\mathbf{x}) \leq f_i(\mathbf{y})$ holds for all i. When the latter relation is furthermore strict (i.e., $f_i(\mathbf{x}) < f_i(\mathbf{y})$) for at least one i, we say that $\mathbf{x}$ *dominates* $\mathbf{y}$, written as $\mathbf{x} \prec \mathbf{y}$. A solution that is not dominated by any other solution in $\mathcal{X}$ is called *Pareto optimal*. The *Pareto set* $\mathcal{A}^\star$ is the set of all Pareto optimal solutions, i.e., $\mathcal{A}^\star = \{\mathbf{x} \in \mathcal{X} : \nexists \mathbf{y} \in \mathcal{X} : \mathbf{y} \prec \mathbf{x}\} \subset \mathcal{X}$. The image of the Pareto set under $\mathbf{f}$ is called the *Pareto front*, i.e., $\{\mathbf{f}(\mathbf{x}) : \mathbf{x} \in \mathcal{A}^\star\} \subset \mathbb{R}^m$.

The aim of multi-objective optimization is to approximate the Pareto set with a set of non-dominated solutions called an *approximation set* $\mathcal{A}$. Let $\mathcal{S} \subseteq \mathcal{X}$ be a solution set, that can contain dominated solutions and let $A : \wp(\mathcal{X}) \to \wp(\mathcal{X})$ be the approximation set given by $\mathcal{S}$, i.e., $A(\mathcal{S}) = \{\mathbf{x} \in \mathcal{S} : \nexists \mathbf{y} \in \mathcal{S} : \mathbf{y} \prec \mathbf{x}\}$, where $\wp(\mathcal{X})$

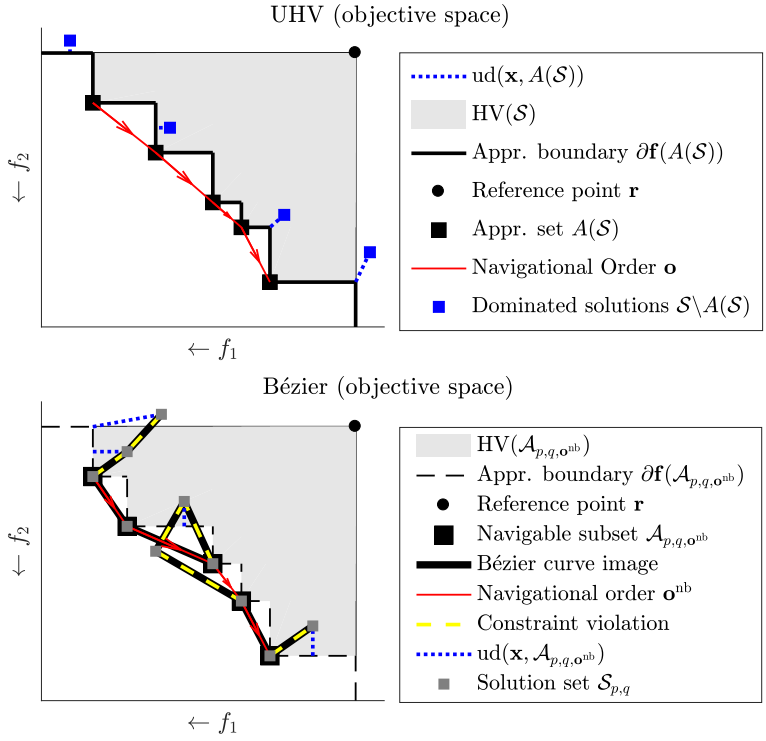

Figure 7.1: Illustration of the Uncrowded Hypervolume (UHV) (top) for a bi-objective minimization problem, and the Bézier parameterization (bottom).

is the powerset of $\mathcal{X}$. The *hypervolume* measure $\text{HV} : \wp(\mathcal{X}) \to \mathbb{R}$ [9, 220] measures the area or volume dominated by all solutions in the approximation set, bounded by a user-defined reference point $\mathbf{r} \in \mathbb{R}^m$, as shown in Figure 7.1. As the hypervolume ignores dominated solutions, we use the *uncrowded distance* to assign a quality value to dominated solutions [63]. The uncrowded distance $\text{ud}_\mathbf{f}(\mathbf{x}, \mathcal{A})$ measures the shortest Euclidean distance between $\mathbf{x}$ and the *approximation boundary* $\partial\mathbf{f}(\mathcal{A})$, when $\mathbf{x}$ is dominated by any solution in $\mathcal{A}$ or outside the region defined by $\mathbf{r}$, and is defined $\text{ud}_\mathbf{f}(\mathbf{x}, \mathcal{A}) = 0$ else (Figure 7.1). It is called the uncrowded distance as the shortest distance to $\partial\mathbf{f}(\mathcal{A})$ is obtained for a point on the boundary that is not in $\mathcal{A}$ itself. Combining the uncrowded distance with the hypervolume measure results in the *Uncrowded Hypervolume* (UHV) (Chapter 6),

$$\text{UHV}_\mathbf{f}(\mathcal{S}) = \text{HV}_\mathbf{f}(\mathcal{S}) - \frac{1}{|\mathcal{S}|} \sum_{\mathbf{x} \in \mathcal{S}} \text{ud}_\mathbf{f}(\mathbf{x}, A(\mathcal{S}))^m. \tag{7.1}$$

We use the subscript $\mathbf{f}$ to denote that the value of the UHV is computed with

respect to the multi-objective problem $\mathbf{f}$. To be able to optimize the UHV of a solution set, a parameterization of solution sets is required. Let $\phi \in \mathbb{R}^l$ be such a parameterization consisting of l decision variables, and let $S(\phi) = \{\mathbf{x}_1, \mathbf{x}_2, \ldots\}$ be an operator that transforms ϕ into its corresponding solution set. The resulting UHV-based optimization problem is then given by,

$$\begin{aligned} &\text{maximize} \quad \mathrm{UHV}_{\mathbf{f},S}(\phi) = \mathrm{HV}_{\mathbf{f}}(S(\phi)) - \frac{1}{|\mathcal{S}(\phi)|} \sum_{\mathbf{x} \in \mathcal{S}(\phi)} \mathrm{ud}_{\mathbf{f}}(\mathbf{x}, A(\mathcal{S}(\phi)))^m, \\ &\text{with} \quad \mathbf{f} : \mathcal{X} \subseteq \mathbb{R}^n \to \mathbb{R}^m, \quad S : \mathbb{R}^l \to \wp(\mathcal{X}), \quad \phi \in \mathbb{R}^l. \end{aligned} \tag{7.2}$$

In a parameterization that is commonly used, solution sets $\mathcal{S}_p$ of fixed size p are considered, and the decision variables of the solutions in $\mathcal{S}_p$ are simply concatenated, i.e., $\phi = [\mathbf{x}_1 \cdots \mathbf{x}_p] \in \mathbb{R}^{p \cdot n}$ [21, 142, 204]. Using this parameterization, the resulting single-objective optimization problem is $l = p \cdot n$ dimensional. In Chapter 6, GOMEA [37] was used to efficiently solve this problem by exploiting the *grey-box* (gb) property that not all solutions $\mathbf{x}_i$ have to be recomputed when only some decision variables change. The resulting algorithm, which we call UHVEA-gb here (and was called UHV-GOMEA-Lm in Chapter 6), greatly outperformed the mostly similar algorithm UHVEA-bb (called UHV-GOMEA-Lf there) but in which the UHV was considered to be a *black box* (bb). This problem parameterization however does not guarantee any degree of navigational smoothness of the approximation set, which is the key goal in this chapter.

7.3 A measure for navigational smoothness

We introduce a measure for the navigational smoothness of an approximation set. Let $\mathcal{S}_p = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p\}$ be an approximation set of size p. Furthermore, let the *navigation order* $\mathbf{o}$ be a permutation of (a subset of) $I = \{1, 2, \ldots, p\}$, representing the indices of the solutions in $\mathcal{S}_p$ that the decision maker assesses in the order the solutions are inspected. The *smoothness* $\mathrm{Sm}(\mathcal{S}_p, \mathbf{o})$ is then defined as,

$$\mathrm{Sm}(\mathcal{S}_p, \mathbf{o}) = \frac{1}{p-2} \sum_{i=2}^{p-1} \frac{\|\mathbf{x}_{\mathbf{o}_{i-1}} - \mathbf{x}_{\mathbf{o}_{i+1}}\|}{\|\mathbf{x}_{\mathbf{o}_{i-1}} - \mathbf{x}_{\mathbf{o}_i}\| + \|\mathbf{x}_{\mathbf{o}_i} - \mathbf{x}_{\mathbf{o}_{i+1}}\|}. \tag{7.3}$$

This smoothness measure measures the *detour length*, which can be understood as the extra distance traveled (in decision space) when going to another solution via an intermediate solution, compared to directly going there.

Throughout this chapter, we will consider a navigational order $\mathbf{o}$ for approx-

imation sets $\mathcal{A}$ such that $f_1(\mathbf{x}_{\mathbf{o}_i}) < f_1(\mathbf{x}_{\mathbf{o}_j})$ holds whenever $i < j$ holds, i.e., from left to right in the objective space plot Figure 7.1. We therefore simply write $\mathrm{Sm}(\mathcal{A}, \mathbf{o}) = \mathrm{Sm}(\mathcal{A})$ from now on. Note that $\mathrm{Sm}(\mathcal{A}) \in [0, 1]$, and only if all solutions are co-linear in decision space, $\mathrm{Sm}(\mathcal{A}) = 1$ holds. This we consider the ideal scenario, where the decision variables of solutions change perfectly predictably. This also implies that any other (continuous) non-linear curve is not considered to be perfectly smooth. Although one could argue for different definitions of smoothness, we will see later that this measure serves our purpose for distinguishing smoothly from non-smoothly navigable approximation sets.

7.4 Approximations sets as a Bézier curve

A Bézier curve $\mathbf{B}(t; \mathcal{C}_q)$ is a parametric curve that is commonly used in computer graphics and animations to model smooth curves and trajectories [72]. An n-dimensional Bézier curve is fully specified by an ordered set of $q \geq 2$ control points $\mathcal{C}_q = \{\mathbf{c}_1, \ldots, \mathbf{c}_q\}$ with $\mathbf{c}_j \in \mathbb{R}^n$. For $0 \leq t \leq 1$, with $\binom{q}{j}$ the binomial coefficients, $\mathbf{B}(t; \mathcal{C}_q)$ is given by,

$$\mathbf{B}(t; \mathcal{C}_q) = \sum_{j=1}^{q} b_{j-1,q-1}(t)\mathbf{c}_j, \quad \text{with,} \quad b_{j,q}(t) := \binom{q}{j}(1-t)^{q-j}t^j, \tag{7.4}$$

Note that if all control points are points in a box-constrained decision space, i.e., $\forall \mathbf{c}_i \in \mathcal{X}$, then $\mathbf{B}(t; \mathcal{C}_q) \subset \mathcal{X}$ holds. The first and last control points are always the end points of the Bézier curve, while intermediate control points do not generally lie on the curve. For certain choices of the control points, a Bézier curve with $q \geq 4$ may intersect itself or have a cusp, as illustrated in Figure 7.2.

We parameterize a solution set $\mathcal{S}_p = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ of fixed size p using an n-dimensional Bézier curve $\mathbf{B}(t; \mathcal{C}_q)$ with q control points. On this curve, p points $\mathbf{x}_i = \mathbf{B}\left((i-1)/(p-1); \mathcal{C}_q\right)$ are selected, evenly spread in the domain of t. The parameterized solution set $S_{p,q}(\phi) = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_p\}$ is then given by,

$$S_{p,q}(\phi) = \left\{\mathbf{B}\left(\frac{0}{p-1}; \mathcal{C}_q\right), \mathbf{B}\left(\frac{1}{p-1}; \mathcal{C}_q\right), \ldots, \mathbf{B}\left(\frac{p-1}{p-1}; \mathcal{C}_q\right)\right\},$$

with $\phi = [\mathbf{c}_1 \cdots \mathbf{c}_q] \in \mathbb{R}^{q \cdot n}$. Note that inverting the order of control points does not affect the Bézier curve. To avoid this symmetry in the parameterization, we standardize the curve direction throughout optimization. After a change of the curve, we check if $f_1(\mathbf{c}_1) < f_1(\mathbf{c}_q)$ holds. If not, the order of the control points is simply inverted.

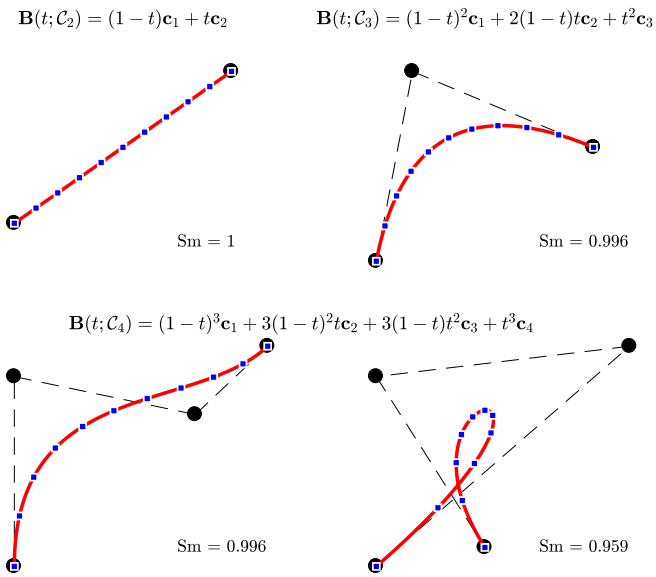

Figure 7.2: Illustration of Bézier curves (red) in decision space with different control points (black). Blue points correspond to $p = 10$ evenly spread values of t, and the smoothness (Sm) of these p points is given, computed based on $\mathbf{o}^{\text{bez}}$.

7.4.1 A navigational order for Bézier parameterizations

Solution sets $\mathcal{S}_{p,q} = S_{p,q}(\phi)$ parameterized by a Bézier curve introduce an intrinsic order $\mathbf{o}^{\text{bez}}$ of solutions by following the curve from $t = 0$ to $t = 1$. Even though the solutions in $\mathcal{S}_{p,q}$ now lie on a smooth curve in decision space, it might very well be that some of these solutions dominate others. We define a navigational-Bézier (nb) order $\mathbf{o}^{\text{nb}}$ for a solution set $\mathcal{S}_{p,q}$ that follows the order of solutions $\mathbf{o}^{\text{bez}}$ along the Bézier curve, but also aligns with the left-to-right ordering described in Section 7.3. Pseudo code for $\mathbf{o}^{\text{nb}}$ is given in Algorithm 7.1, and an example is given in Figure 7.1. The navigational order $\mathbf{o}^{\text{nb}}$ starts from the solution with best f_1-value and continues to follow the Bézier curve (i.e., in the order $\mathbf{o}^{\text{bez}}$) until the solution with best f_2-value is reached, only improving in f_2 (and thereby worsening in f_1) along the way, and skipping solutions that violate this property. Let $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}} = A^{\text{nb}}(\mathcal{S}_{p,q}, \mathbf{o}^{\text{bez}})$ be the resulting subset of $\mathcal{S}_{p,q}$ pertaining to exactly the solution indices as specified in $\mathbf{o}^{\text{nb}}$, and note that this is an approximation set.

7.4.2 Unfolding the Bézier curve (in objective space)

Smoothly navigable approximation sets can now be obtained by maximizing the hypervolume of $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}}$. To maximize the number of navigable solutions $|\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}}| = |\mathbf{o}^{\text{nb}}|$, we need to unfold the Bézier curve in objective space. For

Algorithm 7.1: Navigational order for Bézier parameterizations

function: $[\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}}, (\mathbf{o}^{\text{nb}})] = \text{A}^{\text{nb}}(\mathcal{S}_{p,q}, \mathbf{o}^{\text{bez}})$
input : Bézier solution set $\mathcal{S}_{p,q} = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ with intrinsic ordering $\mathbf{o}^{\text{bez}}$
output : Approximation (sub)set $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}}$, (navigational order $\mathbf{o}^{\text{nb}}$),

$\eta = \arg\min_{i \in \{1,\ldots,p\}} f_1(\mathbf{x}_{\mathbf{o}_i^{\text{bez}}})$;
$\mathbf{o}^{\text{nb}} = [\mathbf{o}_\eta^{\text{bez}}]$ and $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}} = \{\mathbf{x}_{\mathbf{o}_\eta^{\text{bez}}}\}$;
for $j = \eta, \ldots, p$ **do**
  if $\mathbf{x}_{\mathbf{o}_j^{\text{bez}}} \in A(\mathcal{S}_{p,q})$ **and** $f_2(\mathbf{x}_{\mathbf{o}_j^{\text{bez}}}) < f_2(\mathbf{x}_{\mathbf{o}_{\text{end}}^{\text{nb}}})$ **then**
    $\mathbf{o}^{\text{nb}} = [\mathbf{o}^{\text{nb}} \,;\, \mathbf{o}_j^{\text{bez}}]$ and $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}} = \mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}} \cup \{\mathbf{x}_{\mathbf{o}_j^{\text{bez}}}\}$; `// here` $\mathbf{o}_{\text{end}}^{\text{nb}} = \mathbf{o}_j^{\text{bez}}$

Algorithm 7.2: Bézier constraint violation function

function: $C(\mathcal{S}_{p,q}, \mathbf{o}^{\text{bez}}) \geq 0$
input : Bézier solution set $\mathcal{S}_{p,q} = \{\mathbf{x}_1, \ldots, \mathbf{x}_p\}$ with intrinsic ordering $\mathbf{o}^{\text{bez}}$
output : Constraint value $C \geq 0$

$[\mathcal{A}, \mathbf{o}^{\text{nb}}] = \text{A}^{\text{nb}}(\mathcal{S}_{p,q}, \mathbf{o}^{\text{bez}})$; `// See Algorithm 7.1`
$C = \frac{1}{|\mathcal{S}_{p,q}|} \sum_{\mathbf{x} \in \mathcal{S}_{p,q}} \text{ud}_{\mathbf{f}}(\mathbf{x}, \mathcal{A}^m)$; `// Uncrowded distance (ud), see (7.1)`
for $j = 1, \ldots, |\mathcal{S}_{p,q}| - 1$ **do**
  if $\mathbf{o}_j^{\text{bez}} \notin \mathbf{o}^{nb}$ *or* $\mathbf{o}_{j+1}^{\text{bez}} \notin \mathbf{o}^{nb}$ **then**
    $C = C + \|\mathbf{f}(\mathbf{x}_{\mathbf{o}_j^{\text{bez}}}) - \mathbf{f}(\mathbf{x}_{\mathbf{o}_{j+1}^{\text{bez}}})\|$; `// Euclidean distance in` $\mathbb{R}^m$

this, we introduce a constraint violation function $C(\mathcal{S}_{p,q}, \mathbf{o}^{\text{nb}}) \geq 0$, as given in Algorithm 7.2 and illustrated in Figure 7.1. It is composed of two parts. The first part is similar to the uncrowded distance term in Eqn. (7.1), but the approximation boundary is now given by $\mathcal{A}_{p,q,\mathbf{o}^{\text{nb}}}$. The second part aims to pull solutions that are not in $\mathcal{S}_{p,q,\mathbf{o}^{\text{nb}}}$ towards neighboring solutions on the Bézier curve.

7.4.3 Bézier parameterization + GOMEA = BezEA

The resulting Bézier curve optimization problem is given by,

$$\begin{aligned} \text{maximize} \quad & \text{HV}_{\mathbf{f}, S_{p,q}}(\phi) = \text{HV}_{\mathbf{f}}(A^{\text{nb}}(S_{p,q}(\phi))), \\ \text{with} \quad & C(S_{p,q}(\phi), \mathbf{o}^{\text{nb}}(\phi)) = 0, \\ & \mathbf{f} : \mathcal{X} \subseteq \mathbb{R}^n \to \mathbb{R}^m, \quad S_{p,q} : \mathbb{R}^{q \cdot n} \to \wp(\mathcal{X}), \quad \phi \in \mathbb{R}^{q \cdot n}. \end{aligned} \tag{7.5}$$

We use *constraint domination* to handle constraint violations [51]. With constraint domination, the fitness of a solution is computed regardless of its feasibility. When comparing two solutions, if both are infeasible (i.e., $C > 0$), the solution with

the smallest amount of constraint violation is preferred. If only one solution is infeasible, the solution that is feasible is preferred. Finally, if both solutions are feasible (i.e., $C = 0$), the original ranking based on fitness is used.

Bézier curves have no local control property, meaning that a change of a control point affects all solutions on the curve. Partial evaluations can therefore no longer be exploited with this parameterization, and we thus solve this problem with the black-box version of GOMEA. Analogous to the UHV naming, we brand the resulting algorithm Bézier-GOMEA-bb, which we abbreviate to BezEA. A detailed description of GOMEA can be found in [37], and a description of UHV-GOMEA in Chapter 6.

7.5 Numerical experiments

We compare BezEA with UHVEA-gb and UHVEA-bb. Source code for these algorithms is available online [140]. These methods use a different hypervolume-based representation of the multi-objective problem, but use very similar variation and selection mechanisms, making the comparison between these methods most fair. We use the guideline setting for the population size N of GOMEA with full linkage models in a black-box setting [31], which for separable problems yields $N = \lfloor 10\sqrt{l} \rfloor$ and for non-separable problems $N = 17 + \lfloor 3l^{1.5} \rfloor$. BezEA solves a single-objective problem of $l = qn$ decision variables. UHVEA-bb solves a single objective problem of $l = pn$ decision variables. UHVEA-gb solves the same problem by not considering all pn decision variables simultaneously, but by updating only subsets of $l = n$ decision variables, on which we base the population size guideline for UHVEA-gb.

We furthermore compare to the domination-based MO-GOMEA [38]. In MO-GOMEA, a population of N_{mo} solutions is aimed to approximate the Pareto front by implicitly balancing diversity and proximity. Truncation selection is then performed based on domination rank. This selection is clustered into K_{mo} overlapping clusters that model different parts of the approximation front. A Gaussian distribution is then fitted on each cluster, using very similar update rules as the single-objective GOMEA. This therefore allows for a most fair comparison to BezEA and UHVEA. MO-GOMEA maintains an elitist archive, aimed to contain 1000 solutions. For a fair comparison to the hypervolume-based methods that obtain an approximation set of at most p solutions, we reduce the obtained elitist archive of MO-GOMEA to p solutions using greedy Hypervolume Subset Selection (gHSS) [80], which we denote by MO-GOMEA*. As described in Chapter 6, to align

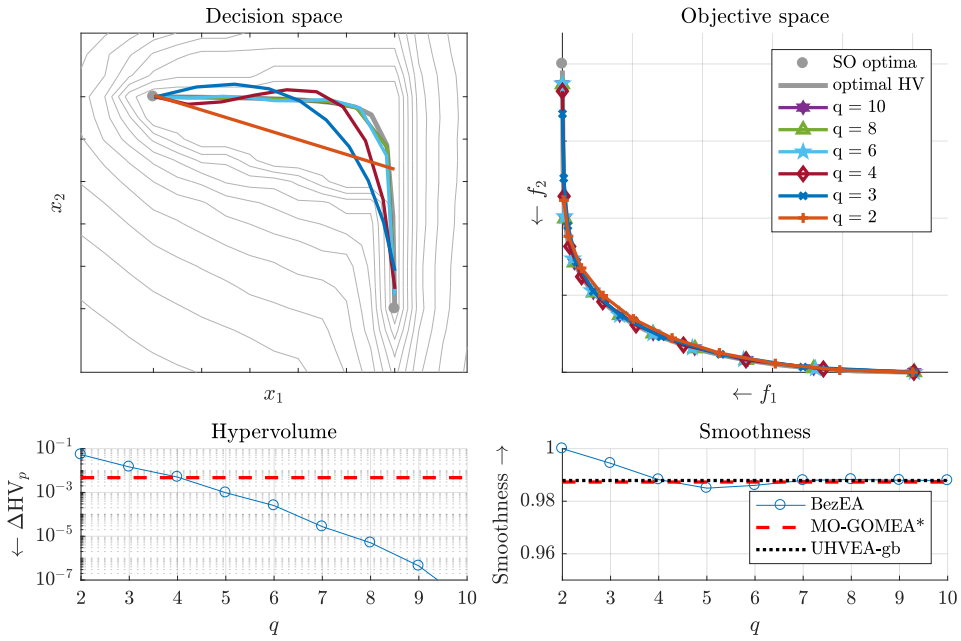

Figure 7.3: Bézier curve approximations of the Pareto set of the curvePS problem (left), obtained with BezEA. Contour lines show domination ranks, the corresponding approximation fronts (middle), and $\Delta\mathrm{HV}_{10}$ together with smoothness (right).

MO-GOMEA with the other algorithms, we set $N_{mo} = p \cdot N$ and $K_{mo} = 2p$ such that the overall number of solutions in the populations is the same, and all sample distributions are estimated from the same number of solutions. As performance measure, we use $\Delta\mathrm{HV}_p = \mathrm{HV}_p^\star - \mathrm{HV}(\mathcal{A}_p)$, which is the distance to the maximal hypervolume $\mathrm{HV}_p^\star$ obtainable with a solution set of at most p solutions.

7.5.1 Increasing q

We illustrate how increasing the number of control points q of the Bézier curve improves achievable accuracy of BezEA (with $q = \{2, \ldots, 10\}$ and $p = 10$) in case the Pareto set is non-linear. For this, we construct a simple two-dimensional problem *curvePS*, with objective functions $f_1^{\mathrm{curvePS}}(\mathbf{x}) = (x_1 - 1)^2 + 0.01x_2^2$ and $f_2^{\mathrm{curvePS}}(\mathbf{x}) = x_1^2 + (x_2 - 1)^2$. A large computational budget was used to show maximally achievable hypervolume, and standard deviations are therefore too small to be visible.

Results

Results are shown in Figure 7.3. A larger q results in a better approximation of the leftmost endpoint of the Pareto front (second subfigure), thereby improving $\Delta\mathrm{HV}_p$

(third subfigure), but slightly lowering smoothness (fourth subfigure), as the Bézier curve deviates from a straight line. MO-GOMEA*, UHVEA-gb, and BezEA for large q all obtain a very similar smoothness. As MO-GOMEA* does not explicitly optimize the hypervolume of its approximation set, it obtains a slightly different distribution of solutions, which results in a lower hypervolume. Additionally, MO-GOMEA* does not converge to the Pareto set due to the finite population size and infinitely large Pareto set, as described in more detail in Chapter 6. Even though this is a fundamental limitation of domination-based MOEAs, this level of accuracy is often acceptable in practice.

7.5.2 Comparison with UHV optimization

Next, we demonstrate the behavior of BezEA compared to UHVEA on the simple *bi-sphere* problem, which is composed of two single-objective sphere problems, $f_{\text{sphere}}(\mathbf{x}) = \sum_{i=1}^{n} x_i^2$, of which one is translated, $f_1^{\text{bi-sphere}}(\mathbf{x}) = f_{\text{sphere}}(\mathbf{x})$, and $f_2^{\text{bi-sphere}} = f_{\text{sphere}}(\mathbf{x} - \mathbf{e}_1)$, where $\mathbf{e}_i$ is the i^{th} unit vector. We set $n = 10$, and initialize all algorithms in $[-5, 5]^n$. This is a separable problem and we therefore use the univariate population size guideline (i.e., $N = \sqrt{l}$). We consider the cases $p = \{10, 100\}$. The computational budget is set to $2p \cdot 10^4$ evaluations of the multi-objective problem given by $\mathbf{f}$ (MO-fevals). When the desired number of solutions p along the front is large, neighboring solutions are nearby each other on the approximation front. This introduces a dependency between these solutions, which needs to be taken into account in the optimization process to be able to effectively solve the problem, as was shown in Chapter 6.

Results

Results are shown in Figure 7.4. This problem is unimodal with a linear Pareto set, and the smoothness of (a subset) of the Pareto set is therefore 1.0. As UHVEA-gb converges to a subset of the Pareto set (see Chapter 6), it ultimately obtains a smoothness of 1.0, even though its smoothness is initially lower. MO-GOMEA* does not converge to the Pareto set, and its smoothness stagnates close to 1.0 when $p = 10$, but stagnates around 0.7 when $p = 100$. BezEA with $q = 2$ has per construction a perfect smoothness of 1.0, and for $q = 3$ and $q = 4$, the obtained smoothness is close to 1. With $q = 5$ control points, BezEA does not converge within the given budget, resulting in a lower smoothness within the computational budget. UHVEA-gb furthermore shows a better convergence rate, which could be because UHVEA-gb can exploit partial evaluations, while this is not possible with BezEA. However, UHVEA-bb, which also does not perform partial evaluations, is unable to solve the problem for $p = 100$. This difference between BezEA and

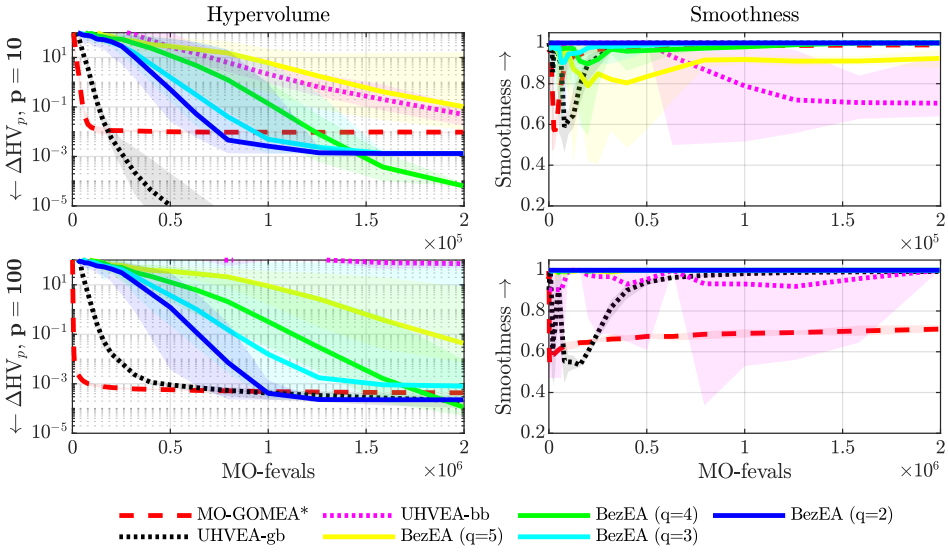

Figure 7.4: Comparison of UHVEA with BezEA and MO-GOMEA* on the bi-sphere problem with $n = 10$ and $p = 10$ (top row) and $p = 100$ (bottom row). Left two subfigures show mean scores, and the shaded areas represent min/max scores, obtained over 10 runs. Objective and decision space subfigures show results of a single run. Solutions in the decision space projection are sorted based on their f_0-value, from best to worst.

UHVEA-bb could be attributed to the lower degree of freedom that BezEA has due to the rather fixed distribution of solutions. This distribution does however not exactly correspond to the distribution of $\text{HV}_p^\star$. This is why a stagnation in terms of hypervolume convergence can be observed for small values of q. The solutions of BezEA are equidistantly distributed along the curve in terms of t. By doing so, intermediate control points can be used to adapt the distribution of solutions (when $q > 2$). This is why BezEA with $q = 4$ can obtain a better ΔHV_p than BezEA with $q = 2$, even though the Pareto set is linear. For $p = 100$, BezEA obtains a better ΔHV_p than UHVEA-gb, which can be explained by the increased problem complexity when the desired number of solutions along the front is large. Increasing the population size N of UHVEA-gb would (at least partially) overcome this, but we aimed here to show that BezEA does not suffer from this increased complexity as its problem dimensionality depends on q, not p.

7.5.3 WFG benchmark

We benchmark BezEA, UHVEA, and MO-GOMEA on the nine commonly used WFG functions [101]. We consider bi-objective WFG problems with $n = 24$ decision variables of which $k_{\text{WFG}} = 4$ are WFG-position variables. We furthermore

Table 7.1: Obtained hypervolume HV_p (mean ± standard deviation (rank)) and mean navigational smoothness (Sm) for the 9 WFG problems with $p = 9$ solutions. Bold are best scores per problems, or those not statistically different from it.

#	MO-GOMEA*		UHVEA-gb		BezEA ($q = 2$)		BezEA ($q = 3$)	
	HV_9	Sm	HV_9	Sm	HV_9	Sm	HV_9	Sm
1	**97.60** ± 0.7 (1)	*0.76*	93.62 ± 1.7 (2)	*0.67*	90.35 ± 1.1 (4)	*1.00*	90.37 ± 1.2 (3)	*0.99*
2	110.09 ± 0.0 (2)	*0.86*	**110.38** ± 1.0 (1)	*0.66*	97.74 ± 0.0 (4)	*1.00*	97.85 ± 0.0 (3)	*0.98*
3	116.11 ± 0.1 (4)	*0.93*	116.42 ± 0.1 (3)	*0.71*	**116.50** ± 0.0 (1)	*1.00*	**116.50** ± 0.0 (2)	*1.00*
4	111.88 ± 0.8 (3)	*0.75*	**112.37** ± 0.7 (1)	*0.69*	111.59 ± 1.3 (4)	*1.00*	**112.19** ± 1.3 (2)	*0.98*
5	112.03 ± 0.1 (3)	*0.66*	111.86 ± 0.3 (4)	*0.63*	112.17 ± 0.0 (2)	*1.00*	**112.19** ± 0.0 (1)	*1.00*
6	113.86 ± 0.3 (3)	*0.88*	114.23 ± 0.2 (2)	*0.72*	**114.34** ± 0.1 (1)	*1.00*	113.02 ± 0.3 (4)	*0.99*
7	114.06 ± 0.1 (4)	*0.94*	114.32 ± 0.1 (3)	*0.66*	114.37 ± 0.0 (2)	*1.00*	**114.38** ± 0.0 (1)	*1.00*
8	110.70 ± 0.2 (4)	*0.79*	**111.24** ± 0.3 (1)	*0.67*	111.07 ± 0.1 (3)	*1.00*	111.14 ± 0.0 (2)	*1.00*
9	**111.70** ± 0.5 (1)	*0.68*	111.46 ± 0.1 (2)	*0.68*	110.19 ± 0.7 (3)	*1.00*	109.36 ± 2.9 (4)	*0.98*

set $p = 9$ and a computational budget of 10^7 MO-fevals. In Chapter 6, a population size of $N = 200$ was shown to work well for UHVEA, which we use here also for BezEA. We perform 30 runs, and a pair-wise Wilcoxon rank-sum test with $\alpha = 0.05$ is used to test whether differences with the best obtained result are statistically significant (up to 4 decimals). Ranks (in brackets) are computed based on the mean hypervolume values.

Results

Results are given in Table 7.1.WFG1 is problematic, as none of the algorithms have an explicit mechanism to deal with its flat region. WFG2 has a disconnected Pareto front. MO-GOMEA* and UHVEA-gb both obtain solutions in multiple subsets, while BezEA obtains all solutions in a single connected subset, and spreads out well there. The linear front of WFG3 corresponds to the equidistant distribution of solutions along the Bézier curve, and BezEA outperforms the other methods there. Increasing q generally increases performance of BezEA, except for WFG6 and WFG9. Both these problems are non-separable, and require a larger population size than the currently used $N = 200$ to be properly solved. However, the guideline for non-separable problems results in a population size that is too large to be of practical relevance here. In terms of smoothness, BezEA with $q = 3$ is able to obtain a smoothness close to 1, while simultaneously obtaining the best HV_9 for 4/9 problems. MO-GOMEA* obtains a mean smoothness of 0.81 while UHVEA-gb obtains the worst mean smoothness (0.68). To illustrate the obtained smoothness a parallel coordinate plot for WFG7 is given in Figure 7.5. This figure shows a clear pattern in decision variable values along the front (in the order $\mathbf{o}$) for BezEA. This pattern is not obvious for the other two methods, while they achieve only a slightly lower hypervolume, and a lower smoothness.

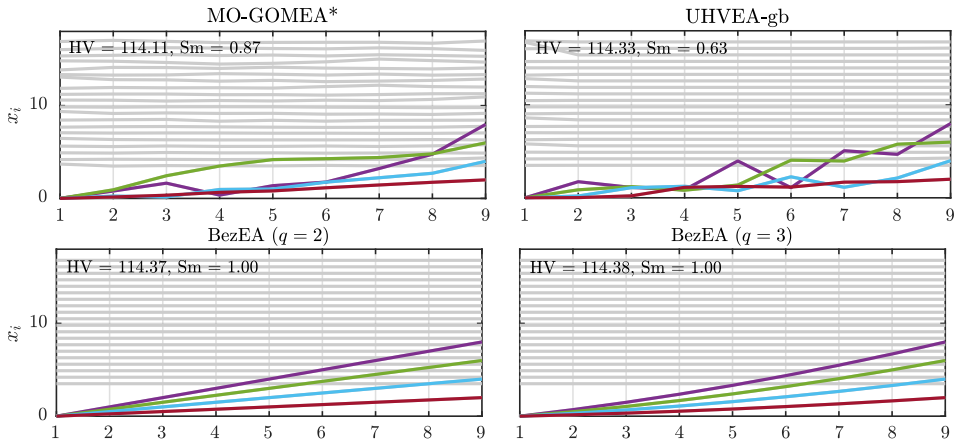

Figure 7.5: Parallel coordinate plots of the decision variables x_i for WFG7. In color the $k_{\text{WFG}} = 4$ position-type decision variables, in grey the remaining decision variables.

7.6 Real-world application: brachytherapy

In this section, we demonstrate the use of BezEA to solve a real-world bi-objective optimization problem that arises in the treatment of prostate cancer with brachytherapy [131]. Brachytherapy is a form of internal radiation therapy. In brachytherapy for prostate cancer, catheters are temporarily placed in, or close to, the prostate. Through these catheters, a radioactive source can be moved, which can be stopped at predefined *dwell positions*. The longer the source dwells at a certain position, the more the surrounding tissue is irradiated. Treatment planning is the process of determining these *dwell times*, such that the tumor is irradiated as much as possible, while surrounding healthy tissue is spared as much as possible. These conflicting objectives make brachytherapy treatment planning inherently a multi-objective optimization problem.

In the Amsterdam University Medical Centers, location Academic Medical Center, treatment planning is performed based on Magnetic Resonance Imaging (MRI), from which a 3D model of the patient is constructed, as described in more detail in Chapter 2. For this, a radiation oncologist and radiation treatment technologist manually delineate the structures on the MRI scans that are of importance for the treatment planning: the location of the catheters, and thereby the location of the dwell positions; the tumor or the *target volumes*, which are the volumes that need to be treated; and other important organs and structures that need to be spared, which are referred to as the Organs At Risk (OARs). An example of such a 3D model of the patient is shown in Figure 1.2. Using the

3D model, the radiation *dose distribution* can be simulated [180, 202]. The dose distribution can be projected as a heatmap on top of the MRI scans (as shown in Figure 1.3), which is used to visually determine the quality of the treatment plan. Additionally, a number of planning criteria have been formulated that measure the quality of different aspects of the dose distribution. We make use of the bi-objective planning model, as given in Chapter 3, Equation (3.1), to combine the planning criteria into two conflicting objectives: the Least Coverage Index (LCI) and the Least Sparing Index (LSI). These two objectives capture the most important trade-off in treatment planning: the trade-off between coverage of the target volumes and sparing of the OARs. Since only two objectives are used, the approximation set can be visualized as an insightful trade-off curve. A radiation oncologist selects one or more treatment plans with desirable trade-offs from this curve. These selected plans are then visually inspected. This visual inspection however takes time, and a limited number of plans can therefore be inspected in clinical practice. Hence, the selection of plans to study further, should be as intuitive as possible. In particular, if plans along the trade-off curve would vary smoothly in terms of the underlying dwell times, there would be an intuitive or sensible variation in the properties of these plans. This is expected to make the selection of plans to inspect as well as the actual visual inspection of plans with similar trade-offs more intuitive and user friendly.

7.6.1 Problem definition

Let us formally define the optimization problem. Denote a treatment plan by a set of dwell times $\mathbf{t} \in \mathbb{R}^n_{\geq 0}$. Typically, depending on the patient and the number of implanted catheters, a few hundred dwell times need to be optimized. The 3D dose distribution is modeled by n_d randomly sampled *dose calculation points* in the target volumes and OARs, and can be formulated as a vector $\mathbf{d} \in \mathbb{R}^{n_d}$. The computation of the dose at the location of the dose calculation points consists of a large matrix-vector multiplication $\mathbf{d} = R\mathbf{t}$, where the *dose-rate matrix* $R^{n_d \times n}$ can be precomputed before optimization. We set the number of *dose calculation points* $|\mathbf{d}|$ to $n_d = |\mathbf{d}| = 2 \cdot 10^4$, and fix the random seed used to sample the dose calculation points for all runs to make the results of different algorithms better comparable [36].

We formulate the to-be-maximized objective functions of the bi-objective model as $\mathbf{f}^{\text{brachy}}(\mathbf{t}) = [\text{LCI}(R\mathbf{t}) \; ; \; \text{LSI}(R\mathbf{t})]$, where we recall that both these objectives are worst-case objectives, making this optimization problem a bi-objective maximin optimization problem. The unit of the LCI is percentage of volume, the unit of the LSI is dose percentage. The LCI and LSI are non-linear and non-separable

functions, and both have a computational complexity of $\mathcal{O}(n_d \log n_d)$. To allow for a more detailed analysis of the results, we consider three of the planning criteria that are part of the LSI in more detail. The volume index $V_{200\%}^{\text{prostate}}$ measures the fraction of the prostate volume that receives 200% of the prescribed dose (13 Gy). This should not be more than 20% of the prostate volume, which we can write as, $\Delta V_{200\%}^{\text{prostate}} = 0.2 - V_{200\%}^{\text{prostate}}$. The larger the value of $\Delta V_{200\%}^{\text{prostate}}$, the better, and the planning criterion is satisfied if $\Delta V_{200\%}^{\text{prostate}} \geq 0$. We furthermore consider the dose indices $D_{1\text{cm}^3}^{\text{bladder}}$ and $D_{1\text{cm}^3}^{\text{rectum}}$, which measure the lowest dose in the most irradiated 1 cm^3 of respectively the bladder and the rectum. We express the dose indices as a fraction of the prescribed dose (13 Gy). The corresponding planning criteria are $\Delta D_{1\text{cm}^3}^{\text{bladder}} = 0.86 - D_{1\text{cm}^3}^{\text{bladder}} \geq 0$ and $\Delta D_{1\text{cm}^3}^{\text{rectum}} = 0.78 - D_{1\text{cm}^3}^{\text{rectum}} \geq 0$.

To focus optimization on a clinically relevant part of the objective space, the bi-objective planning model makes use of a reference point, which we set here to $r = (-0.04, -0.2)$. Based on this reference point, a constraint function is defined,

$$C(R\mathbf{t}) = \max\{-\text{LCI}(R\mathbf{t}) - 0.04, 0\} + \max\{-\text{LSI}(R\mathbf{t}) - 0.2, 0\}. \tag{7.6}$$

Constraint domination is used to handle constraint violations (when $C(R\mathbf{t}) > 0$) [51]. When a dose distribution given by $R\mathbf{t}$ satisfies the constraint (i.e., $C(R\mathbf{t}) = 0$), it satisfies $\text{LCI}(R\mathbf{t}) \geq -0.04$ and $\text{LSI}(R\mathbf{t}) \geq -0.2$. The reference point used here is set tighter than in [36], as results in Chapter 2 indicated that plans with $\text{LCI} < -0.04$ were not of clinical interest. This constraint is not necessary when solving the bi-objective planning problem with UHVEA or BezEA, as simply setting the hypervolume reference point to r is sufficient to guide the search. For the computation of the objective values, it can be exploited that the required matrix-vector multiplication can be performed on a Graphics Processing Unit (GPU) [36]. In the same work, a form of exponential weighting of DVIs was incorporated in the LCI and LSI, so that components of the LCI and LSI that are not the worst, are still optimized, which we use here as well.

7.6.2 A linkage model for UHVEA and MO-GOMEA

When only a few dwell times change, the dose in the dose calculation points can be quickly updated, as only a small subset of the matrix-vector multiplication has to be performed. This is a property that MO-GOMEA can exploit to be able to solve this rather high-dimensional problem with a smaller population size, and thereby in less time [131]. Which subsets of dwell times (i.e., decision variables) are changed simultaneously is captured in a linkage model [38]. To construct the linkage model, hierarchical clustering is used to iteratively cluster dwell times together based on

the distance between the corresponding dwell positions [36]. The result is a linkage model that contains overlapping differently-sized linkage subsets of dwell times.

After each change of the dwell times of a plan, the LCI and LSI of the plan need to be recomputed, even if only a single dwell time changes. This gives a (substantial) constant computational overhead, which makes it inefficient to consider small linkage subsets. Linkage subsets of fewer than five dwell times are therefore removed from the linkage model [36].

To use UHVEA-gb to solve this real-world bi-objective problem, it is reformulated as a $p{\cdot}n$-dimensional single-objective problem, where the dwell times of p plans are concatenated. For each of these plans, the same linkage model is constructed as in MO-GOMEA. A single linkage model for UHVEA-gb is then constructed by taking the union of all p linkage models. Note that treatment plans are thus updated independently (i.e., the maximum number of decision variables that are changed simultaneously is n).

7.6.3 Bézier-specific exploitable properties

In BezEA, solution sets $\mathcal{S}_{p,q} = \{\mathbf{t}_1, \ldots, \mathbf{t}_p\}$ of p solutions are parameterized as points on a Bézier curve with q control points. We can exploit the linearity of the dose distribution computation to reduce computation time. Intuitively, the dose distribution specified by a solution $\mathbf{t}_i$ is a linear interpolation between the dose distributions that correspond to the control points $\mathbf{c}_j$. To evaluate an entire solution set, p dose distributions $\mathbf{d}_i = R\mathbf{t}_i$ need to be computed, from which the objective values $\mathbf{f}^{\text{brachy}}(\mathbf{t}_i) = [\text{LCI}(\mathbf{d}_i)\ ;\ \text{LSI}(\mathbf{d}_i)]$ can be computed. However, since p is generally larger than q, we use that,

$$\mathbf{t}_i = \mathbf{B}\left(\frac{i-1}{p-1}; \mathcal{C}_q\right) = \sum_{j=1}^{q} b_{j-1,q-1}\left(\frac{i-1}{p-1}\right) \cdot \mathbf{c}_j := \sum_{j=1}^{q} b_{j,q}^{i,p} \cdot \mathbf{c}_j.$$

This gives the following expression for the p dose distribution computations in terms of the q control points,

$$\mathbf{d}_i = R\mathbf{t}_i = R\left(\sum_{j=1}^{q} b_{j,q}^{i,p}\mathbf{c}_j\right) = \sum_{j=1}^{q} b_{j,q}^{i,p}(R\mathbf{c}_j).$$

Since the product $R\mathbf{c}_j$ is independent of i, the required number of matrix-vector multiplications reduces hereby from p to q. Note again that the LCI and LSI are non-separable and still need to be computed p times.

Instead of directly optimizing the dwell times, $\mathbf{x} = \sqrt{\mathbf{t}}$ is optimized with UHVEA and BezEA, so that the search space is unbounded, which gives more freedom to

BezEA to fit a curve close to $t_i = 0$, and also makes it easier for UHVEA to sample solutions with dwell times close to zero, without having to worry about boundary handling.

A linkage model for BezEA

For BezEA, when one decision variable of a single control point changes, the corresponding dwell time changes in all p plans. It is therefore not efficient to use the same linkage model as in UHVEA, as UHVEA had the possibility to update solutions independently. Instead, first, all corresponding decision variables of all q control points are clustered together. This results in n clusters each of size q. The same UPGMA clustering algorithm as used in MO-GOMEA is then used to construct a linkage model by iteratively merging the n clusters of size q. In line with MO-GOMEA and UHVEA-gb, the lower bound on the cluster size is set to $5q$, such that always at least 5 dwell times of a plan are changed simultaneously. By this construction, the maximum cluster size is qn (i.e., all of the decision variables).

7.6.4 Experimental setup

It was found that MO-GOMEA performs well on this problem with $N_{\text{mo}} = 96$ and $K_{\text{mo}} = 5$ [36]. Using the relations $K_{\text{mo}} = 2p$ and $N_{\text{mo}} = pN$ as presented before, we deduce from this a population size of $N = 38$ for UHVEA-gb, and $N = 38q$ for BezEA. We compare BezEA with linear approximation sets ($q = 2$) to UHVEA-gb and MO-GOMEA.

As it is in clinical practice only feasible time-wise to inspect a limited set of different plans, we aim the search for solution sets of size $p = 10$. For a fair and insightful comparison of MO-GOMEA with BezEA and UHVEA-gb, we again apply gHSS [80] to reduce the obtained elitist archive (of up to 1250 plans) to an approximation set of p plans for MO-GOMEA, which we denote by MO-GOMEA*.

Since all methods exploit problem-specific properties differently, we perform the analysis with respect to run time instead of MO-fevals. All methods are implemented in C++ and are run on the same Central Processing Unit (CPU) with a time limit of 2 hours, which corresponds to roughly 30 seconds on a GPU. The dose computation is dominates the overall runtime, and is particularly well-suited to be implemented on a GPU. All methods are therefore expected to benefit similarly from a GPU implementation.

We consider three patients here, with respectively $n = \{200, 218, 195\}$ dwell times to optimize. To get insight in the stochastic behavior of the algorithms, we repeat all experiments 10 times and report mean and min/max performance.

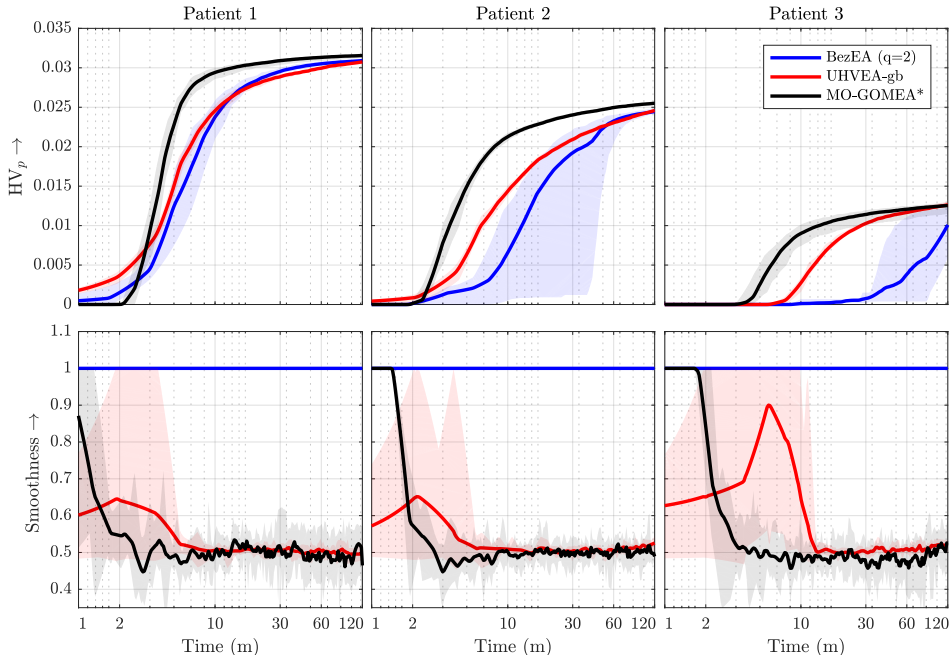

Figure 7.6: Obtained hypervolume (HV_p) and smoothness for three patients with $p = 10$ plans over time in minutes (m). Mean values over 10 runs are shown and the shaded areas show min/max performance. Note that BezEA has a constant smoothness of 1.

7.6.5 Results

Mean HV_p and smoothness results are shown in Figure 7.6. In terms of hypervolume, all methods obtain rather similar values at the end of the run, however, MO-GOMEA* obtains these values faster. Patient 3 was found to be more difficult to optimize. All methods obtain a lower hypervolume for this patient. The difference in obtained hypervolume between UHVEA-gb and BezEA is initially large for this patient. BezEA ultimately obtains similar hypervolume values in most runs, but some runs are not yet converged, which can be overcome by allowing for longer run times. BezEA shows initially more variance in obtained hypervolume (as indicated by the shaded min-max performance), compared to the other two methods. In this phase, BezEA aims to satisfy the constraints, i.e., it aims to find an approximation set in which no solution is dominated. As soon as a solution set is obtained that does not violate any of the constraints introduced in the Bézier problem formulation, the rate of convergence is however rather constant. At that point, the search is driven by hypervolume maximization, similar as in UHVEA-gb, which shows little to no variance in the obtained hypervolume values. This

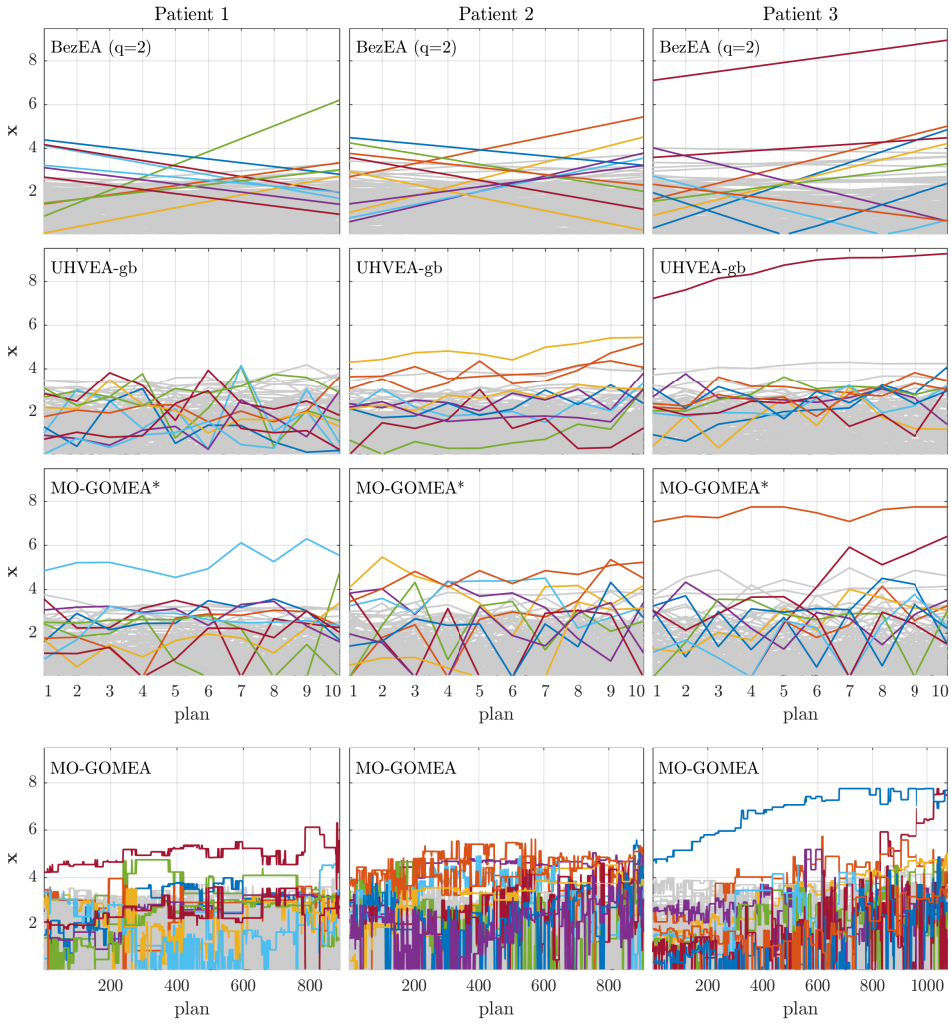

Figure 7.7: Parallel coordinate plot of the final approximation set of a randomly selected run of each method. The plans are sorted and numbered according to the navigational order, i.e., from left to right on the approximation front. Top three rows show the approximation set of size $p = 10$ for BezEA, UHVEA-gb, and MO-GOMEA*. The bottom row shows the elitist archive of MO-GOMEA (with up to 1250 plans). On the horizontal axis, obtained plans are shown, ordered from worst to best LCI. The decision variables ($x_i = \sqrt{t_i}$, unit $\sqrt{s}$) of different plans are connected by a line. The 10 decision variables with the largest standard deviation across plans are highlighted in color, and the others are shown in grey.

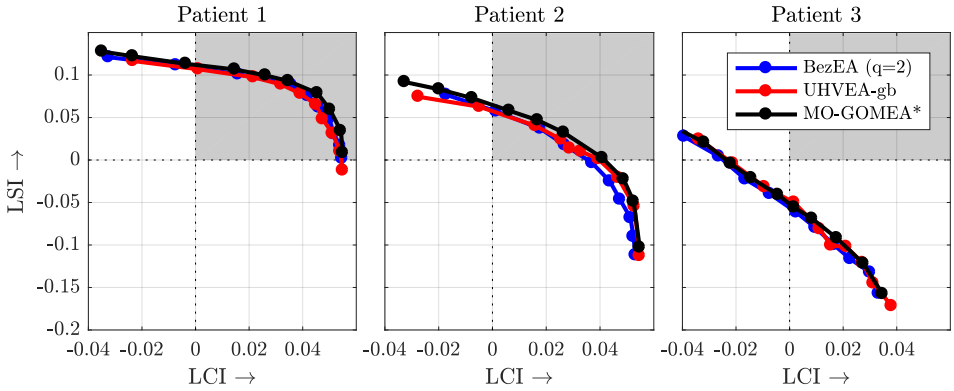

Figure 7.8: Approximation fronts (objective space) obtained with different algorithms, of which the decision values are shown in the parallel coordinate plots in Figure 7.7. Plans in the shaded area (with LCI> 0 and LSI> 0) are particularly interesting from a clinical point of view. However, for some patients (here, Patient 3), it is impossible to obtain plans with these properties, for example due to geometrical reasons.

suggests that the constrained problem formulation of BezEA could be improved in order to better guide the search towards the feasible domain.

BezEA with $q = 2$ has by definition a perfect smoothness of 1.0. The smoothness of both MO-GOMEA* and UHVEA-gb fluctuate around a value of 0.5, which corresponds to a zig-zag pattern in the parallel coordinate plot in Figure 7.7. This parallel coordinate plot shows the obtained decision values. A clear difference can be observed between the approximation set obtained by BezEA and approximation sets obtained by the other methods. Especially MO-GOMEA* shows a clear zig-zag pattern, indicating that plans that are next to each other on the approximation front can have very different decision values, and thereby potentially very different dose distributions.

Figure 7.8 interestingly shows that quite similar looking approximation fronts were obtained by the different methods, even though the corresponding decision values are very different from each other. This suggests that this real-world problem is either highly multimodal, or has (many) small plateaus. The objective space visualization is an important tool for decision making in clinical practice, as it gives high-level insight in the patient-specific trade-off between the LCI and the LSI in the maximally achievable plan quality. The obtained fronts are visually very similar, which would suggest that it likely matters little which of the three fronts would be used in clinical practice to select a single preferred treatment plan. This is especially the case considering also the uncertainties that are at play in the real world setting [202]. Furthermore, as mentioned before, especially plans with

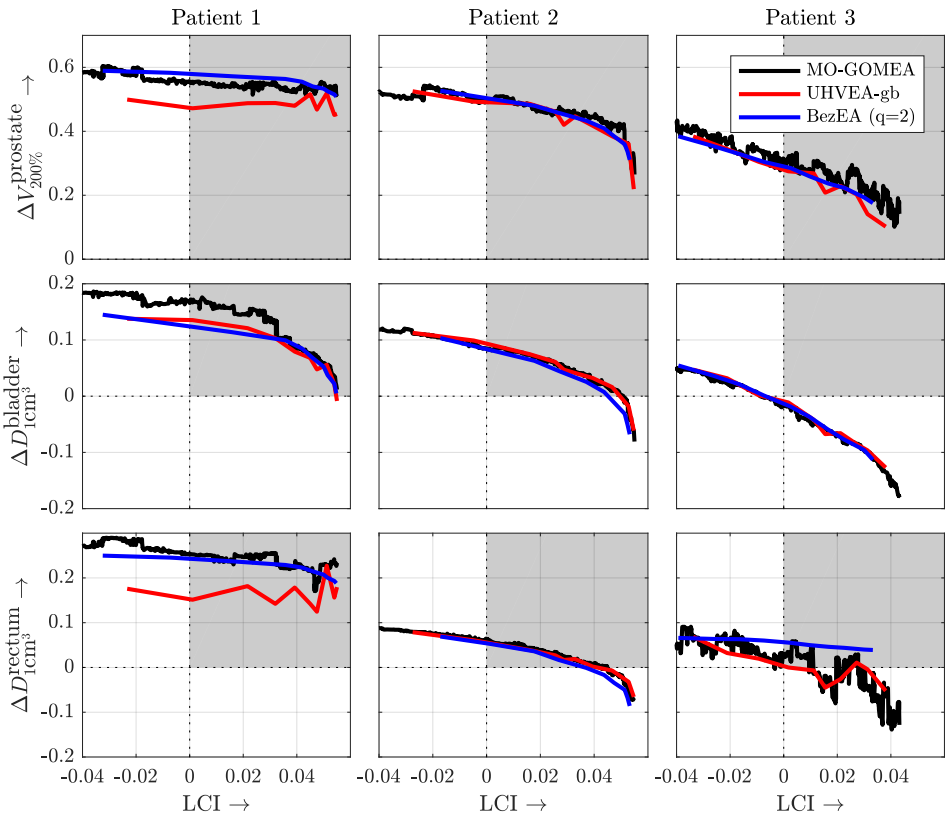

Figure 7.9: Different DVIs, as used for clinical decision making, corresponding to the approximation sets in Figure 7.8. For each plan, the LCI value is taken as before, but on the vertical axis, for each figure, a single DVI-based planning criterion is shown. Plans in the shaded area satisfy all coverage planning aims, and the plotted sparing criterion.

LCI > 0 and LSI > 0 are of clinical interest, if obtainable. The middle parts of the obtained fronts are rather similar, which are generally the plans of clinical interest, although BezEA obtains plans that are of slightly lower objective values.

Figure 7.9 shows that values for the planning criteria obtained with BezEA also fluctuate less than those obtained with UHVEA-gb or MO-GOMEA*. For Patient 2, all methods obtain similar results, and for Patient 1, BezEA is slightly better than UHVEA-gb. For Patient 3, BezEA is unable to obtain plans at the right end of the front, but it obtains a slightly better rectum DVI $\Delta D^{\text{rectum}}_{1\text{cm}^3}$. Especially in terms of smoothness for Patient 3, it is clear that BezEA obtains less variation between values for the planning criteria, which suggests that navigation of the approximation set would be more intuitive.

7.7 Discussion and outlook

In this chapter, we parameterized approximation sets as smooth Bézier curves in decision space, thereby explicitly enforcing a form of smoothness between decision variables of neighboring solutions when the approximation front is traversed, aimed at improving its navigability. We used an UHV-based MO problem formulation that directly allows for the optimization of parameterized approximation sets. Solving this Bézier problem formulation with GOMEA (BezEA), was shown to eventually be competitive in obtained hypervolume compared to UHV-based optimization and domination-based MOEAs, while smoothness is guaranteed. We showed that approximation sets obtained with BezEA show a smoother pattern in terms of decision variables when traversing the approximation front on a set of benchmark problems, which suggests that this approach will lead to more intuitive approximation set navigability for real-world optimization problems.

We limited the experiments on the real-world prostate brachytherapy problem to linear Bézier curves, with $q = 2$ control points, which puts the largest restriction on the obtained solution set. By increasing the number of control points, the small loss in plan quality can be even further reduced, but this will come at additional computational cost. Domain experts will need to be consulted to verify that plans obtained with BezEA are indeed of acceptable quality from a clinical point of view, or that an increase in the number of control points is needed, and whether navigation of smooth approximation sets is indeed more intuitive and faster.

We chose to fix the solution set size p for BezEA during and after optimization, but since a parametric expression of the approximation set is available, it is straightforward to construct a large approximation set after optimization. This could be exploited to increase the performance of BezEA, as it currently shows computational overhead even on the simple bi-sphere problem in terms of multi-objective function evaluations compared to UHVEA.

It is rather straightforward to enforce restrictions upon approximation sets in hypervolume-based multi-objective optimization. Here, we enforced a form of smooth navigability on approximation sets. One could furthermore enforce that the obtained approximation sets reside in a single mode. This could result in a very interesting approach for multimodal optimization of multi-objective optimization problems. When an approximation set spreads out over multiple modes, it is generally not smooth, as one might jump back and forth between different modes when the front is traversed. To improve smoothness in that case, the approximation set can be restricted to only a single mode. This means that BezEA can in principle be used in multimodal optimization. However, it is not very efficient

in its current form. One of the reasons for this is the usage of a hard constraint to ensure that solutions on the Bézier curve do not dominate each other in objective space. The constraint could cause approximation sets that reside in multiple niches to be 'stuck'. Using a niching framework with restarts, such as HillVallEA, can circumvent this problem, but we observed in preliminary results that this approach is rather inefficient. Overcoming this, by using a different parameterization, or a different constraint definition, would make BezEA a very interesting approach for multimodal optimization of multi-objective optimization problems.

The results of BezEA on the bi-objective brachytherapy treatment planning problem show that smoothly navigable approximation sets can be obtained at little to no loss in plan quality in terms of the objective values and at not too much extra cost in terms of running time. This suggests that BezEA is a good alternative to solve the problem while also ensuring navigational smoothness. As such, BezEA may eventually be preferable in a clinical setting.

The smoothness measure introduced in this chapter is a measure for entire solution sets $\mathcal{S}_p$, and not for individual solutions $\mathbf{x}$. It can therefore not be added directly as an additional objective to the original multi-objective problem $\mathbf{f}(\mathbf{x})$. We chose in this chapter to introduce a parameterization of approximation sets that directly enforces smoothness. Alternatively, smoothness could also be added as a second objective to the UHV-based problem formulation. This then results in the pn-dimensional bi-objective optimization problem, given by $h(\mathcal{S}_p) = [\mathrm{UHV}_{\mathbf{f}}(\mathcal{S}_p)\ ;\ \mathrm{Sm}(\mathcal{S}_p)]$. This problem can then be solved with a domination-based MOEA, or even by again formulating it as a (much) higher-dimensional UHV-based single-objective problem. The advantage of included smoothness as a separate objective in the optimization process is that no predetermined parameterization is required. However, whether this approach can be efficient, even when grey-box properties such as partial evaluations are exploited, remains future work.

Finally, the problems considered in this chapter were limited to problems involving two objectives. We believe that the presented results are sufficiently promising that extending this work to problems with $m > 2$ objectives can be said to be an interesting future avenue of research. The Pareto front of non-degenerate problems with m objectives is an $(m-1)$-dimensional manifold. Instead of a one-dimensional Bézier curve, the Pareto set can then be modeled by an $(m-1)$-dimensional Bézier simplex [114]. For problems with more than two objectives, the approximation front (or manifold) does however not directly imply a navigational order. How navigation is performed in case of more than two objectives is therefore likely to be problem dependent, and how navigational smoothness is experienced should therefore be evaluated with end-users.

f

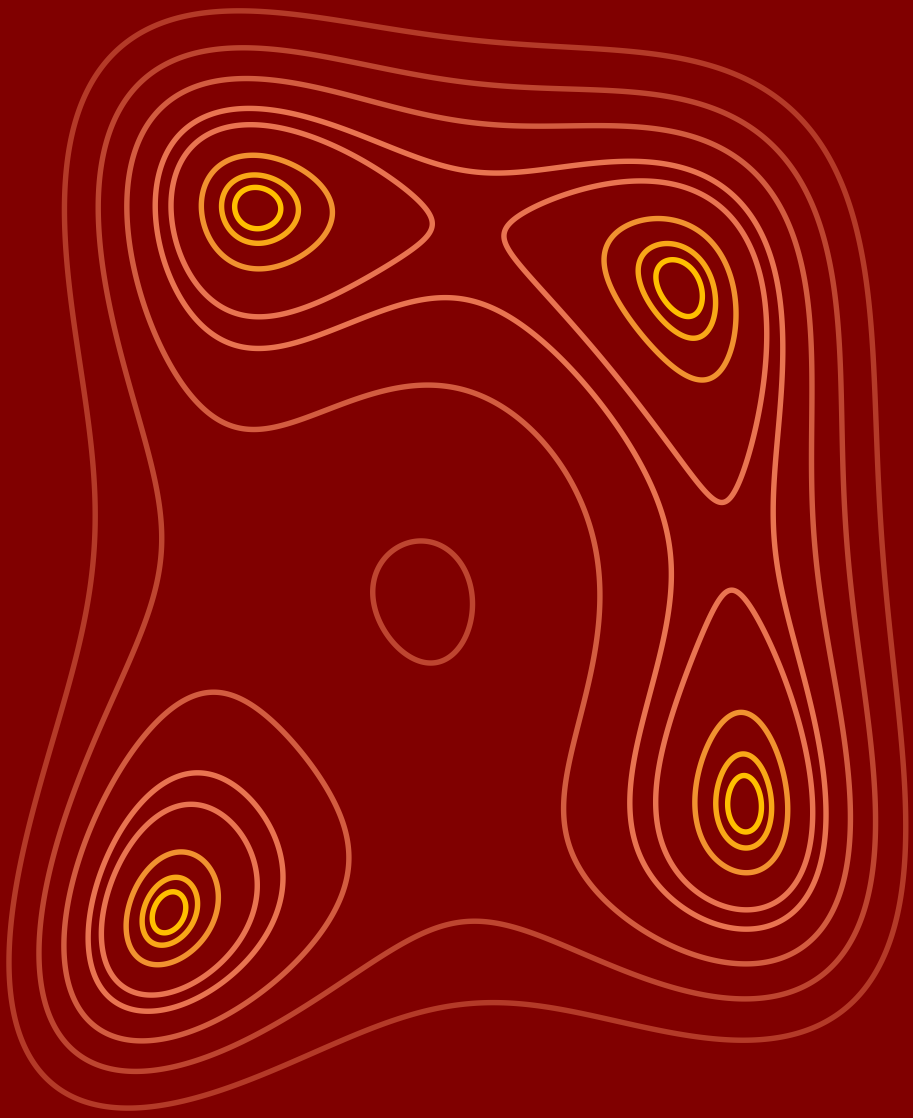

8

General discussion

Abstract

In this chapter, the main findings of the research in this thesis are discussed. In the following section, we first provide answers to the research questions. Afterwards, we address limitations, implications, and future research directions for each of the three pillars: multi-objective optimization, multimodal optimization, brachytherapy treatment planning, and the overall topic of this thesis: model-based evolutionary algorithms.

8.1 Answers to research questions

8.1.1 Research question 1

Can clinically acceptable prostate brachytherapy treatment plans be obtained by optimizing the bi-objective planning model with MO-GOMEA? *(Chapter 2)*

To answer research question 1, clinical planning sessions were recorded on video, in which a planner manually fine-tuned an initial treatment plan that was obtained by HIPO. An analysis of these videos showed that the bi-objective planning model was easily configured based on the clinical protocol and represented our clinical practice well. It was furthermore demonstrated in a retrospective observer study that the bi-objective method automatically generates plans of high-quality, as resulting plans were preferred over the clinically used plan in 98% of the cases. Additionally, according to all observers, who are all brachytherapists, clinically acceptable treatment plans were obtained for all patients. The observers furthermore highly appreciated the possibility to compare multiple high-quality plans and considered it insightful. On multiple occasions in this study the observers mentioned they were interested in possible effects of further manual fine-tuning of the selected plans. This desire to explore whether it would be possible to further improve these plans suggests that there are apparently other clinical desirabilities or considerations of a plan that are not taken into account in the current objectives of the bi-objective planning model. An example of these is the dose to normal tissue, which is not included in the current clinical protocol. Another widely discussed topic is dose homogeneity, and closely related, the variation between dwell times of consecutive dwell positions in individual catheters [13].

With the considered form of bi-objective treatment planning, clinically acceptable plans were obtained for all patients without further fine tuning, suggesting that these clinical desirabilities could be *nice to have*, and are not essential for high-quality treatment plans. Alternatively, it could be that these clinical desirabilities can be satisfied without worsening the current planning criteria too much. This was the case for the only patient in which the clinically used plan was preferred over the presented automatically generated plans. When clinical desirabilities are not taking into account in the optimization objectives, multiple plans can still be presented to the observer, as done in bi-objective treatment planning. When a sufficient number of diverse plans is presented, one of these plans could be expected to satisfy these clinical desirabilities or at least, come closer in this regard than another plan. In the used study setup, only five automatically generated plans were however presented to the observers as we had not yet developed more sophisticated plan selection software. Considering more plans could therefore result in even more desirable plans. However, it is then required that plan selection is intuitive so that many plans can be quickly inspected.

8.1.2 Research question 2

Can time-consuming manual treatment planning for prostate brachytherapy be overcome by automatic tuning of the penalty weights of clinically-available inverse planning methods? *(Chapter 3)*

To answer research question 2, we used the multi-objective evolutionary algorithm MAMaLGaM [181] to automatically tune the parameters of the inverse panning methods Inverse Planning Simulated Annealing (IPSA) [124] and Hybrid Inverse Plan Optimization (HIPO) [118]. To allow for an insightful comparison, automatic tuning was aimed at maximizing the objectives of the bi-objective planning model. By doing so, a set of diverse plans could be obtained with an inverse planning method, which allowed an insightful comparison of its maximally achievable plan quality to other planning methods, by comparing their obtained trade-off curves. Automatic tuning of both IPSA and HIPO was found to result in a set of diverse plans with a wide range of trade-offs. By tuning HIPO, better trade-offs were obtained than by tuning IPSA. For most patients, automatic tuning of HIPO resulted in plans close to quality of the plans obtained by optimizing the bi-objective planning model directly with MO-GOMEA, which gave the best results. This suggest that automatic tuning of the penalty weights of HIPO results in

treatment plans of sufficient quality to be used in clinical practice, and could thereby overcome manual treatment planning. The proposed tuning approach is however computationally expensive, and a significant speed up is required to make it suitable for clinical practice. Modern hardware, such as a Graphics Processing Unit (GPU) could allow for such speed up, as was also found to be very beneficial to speed up MO-GOMEA when optimizing the bi-objective planning model [36].

To avoid having to perform the computationally-intensive automatic parameter tuning during treatment planning, it can be used to generate a class solution (standard penalty weights). A class solution is used in clinical practice to generate a treatment plan that is used as starting point for manual treatment planning. Starting treatment planning with a better class solution could reduce manual treatment planning time, or entirely overcome it. Three class solutions were constructed from the automatically tuned plans, representing three different trade-offs: one with more focus on sparing, one focused more on coverage, and one in the middle. As a result of this research, the class solutions obtained for HIPO were introduced in clinical practice using a two-step procedure. First, the middle class solution was used to generate a treatment plan. If this plan, after inspection of the dose distribution, was found to be delivering too little dose to the target volumes, the coverage class solution was used. Else, the sparing class solution could be used. The resulting plan from this two-step procedure was found to be of better quality than the class solution that was previously used in clinical practice, but manual adaptations of the treatment plan were still required to obtain clinically acceptable treatment plans. As the starting point for manual optimization is better, this is expected to require less manual effort.

These class solutions have since then been used clinically for five patients. It was found that manual fine tuning was still required to obtain desirable plans. The main reason for further manual tuning was an undesirably large variation of dwell times corresponding to dwell positions within the same catheter. The amount of allowable variation can be controlled in HIPO by the Dwell Time Gradient Restriction (DTGR). In our study, the DTGR was disabled, as it is in conflict with all planning criteria in the clinical protocol [13]. Enabling the DTGR can reduce thus the variation in dwell times at some cost in terms of the planning criteria. This effect was however not further investigated, as optimizing the bi-objective planning model with MO-GOMEA results in plans with better values for the DVIs, in fewer computation time, and was therefore considered to be a more promising future direction.

8.1.3 Research question 3

How can model-based evolutionary algorithms be adapted to perform efficient and effective (single-objective) multimodal optimization, without making assumptions on the number of modes, or on their shape or size? *(Chapter 4)*

In Chapter 4, we introduced a two-phase niching framework to make model-based evolutionary algorithms suitable for multimodal optimization. In the first phase, a clustering method is used to cluster the search space into niches, and in the second phase, each of the located niches is explored with a core search algorithm such as AMaLGaM or GOMEA. We introduced two clustering methods in this thesis to cluster the search space into niches. The first of these methods is Hierarchical Gaussian Mixture Learning (HGML) [143]. In HGML, a Gaussian mixture model is fitted to the fitness landscape by clustering the population. Hierarchical clustering was employed to prevent the need for any assumptions on the number of modes in the fitness landscape. We showed that HGML outperforms other methods when there is a good match between the fitness landscape and a Gaussian mixture model. However, when the match is not that good, the performance of HGML was found to deteriorate rapidly. This is in line with what one can expect from a model-based approach. However, there are a few steps in HGML that could be further improved. HGML is based on a greedy clustering order that does not always result in the desired merge order, and therefore does not always obtain the desired Gaussian mixture model. To overcome this, we introduced a second niching method, called Hill-valley Clustering (HVC). In HVC, clusters are formed by pairwise checking whether two solutions belong to the same niche with the *hill-valley test* [197]. The hill-valley evolutionary algorithm (HillVallEA), a combination of HVC and the model-based evolutionary algorithm AMaLGaM-univariate as core search algorithm, was shown to outperform all competitors on the niching competition for multimodal optimization at the Genetic and Evolutionary Computation Conference in both 2018 and 2019 [125]. The problems used in the niching competition are rather low dimensional, ranging from 1 to 20 decision variables. While HillVallEA is expected to scale well to higher-dimensional problems, this has not been thoroughly investigated. Additionally, in practice, it rarely happens that exactly the same (global) optimal value is obtainable in multiple modes. While no assumptions are made in HillVallEA on the value of the global optimum, a relative threshold needs to be set by the user that distinguishes low-quality optima from high-quality optima

based on the best objective value found so far. As this threshold is expressed in terms of objective values, it can often be set based on domain knowledge by determining how much quality the decision maker is willing to trade-off for diversity.

Finally, HillVallEA was found to be unable to fully solve all benchmark problems, even with a large computational budget. This could be due to the use of a two-phase niching framework. As the decision space is only initially clustered into niches in HillVallEA, a very large number of initially sampled solutions is needed to locate small niches. To overcome this, one could perform clustering also later on in the search process. It is however not straightforward to do this efficiently and effectively, and no other method so far has been able to fully solve all these problems.

8.1.4 Research question 4

How can niching via hill-valley clustering be applied to multi-objective optimization problems to allow for efficient and effective multi-objective multimodal optimization, without making assumptions on the number of modes, or on their shape or size? *(Chapter 5)*

We introduced hill-valley clustering for clustering multi-objective decision spaces into niches in Chapter 5. To do so, single-objective hill-valley clustering is performed once for each objective. The result is multiple (different) clusterings of the population. These are then reduced to a single cluster set by taking intersections. We combined hill-valley clustering with MAMaLGaM into the multimodal MOEA MO-HillVallEA. We empirically showed that MO-HillVallEA outperforms MAMaLGaM and other MOEAs in multimodal optimization on a set of multimodal benchmark functions. Furthermore, and perhaps most importantly, we showed that MO-HillVallEA is capable of obtaining, maintaining, and improving multiple local Pareto sets simultaneously over time. MOEAs already maintain approximation sets of multiple solutions, and to successfully aim these methods at multimodal optimization, a balance needs to be found between diversity within a solution set (i.e., different trade-offs), and among solution sets (i.e., different modes). It was shown that improving diversity in parameter space resulted in a loss of diversity in objective space.

Multimodal optimization for multi-objective optimization problems is typically benchmarked on problems with a small number of relatively similar modes, as was done in Chapter 5. We demonstrated that MO-HillVallEA performs rather

well on this type of problems. In case there is a limited number of modes, the population of the MOEA can be divided over all located modes. All these modes can then be explored in parallel, which is how MO-HillVallEA performs the search. However, when there are many modes, this approach is inefficient or even ineffective, as there are too few solutions in each mode to be able to properly model the fitness landscape. To make MO-HillVallEA more efficient on such problems, a regularization of some sort is required that limits the number of modes that is explored in parallel, for exampling using a racing procedure [24].

8.1.5 Research question 5

How can evolutionary multi-objective optimization approaches be designed where the individuals in the population represent entire approximation sets and fitness is directly based on the hypervolume indicator or extensions thereof so as to ensure convergence to a subset of the Pareto set?
(Chapter 6)

We used the uncrowded distance [195] to extend the hypervolume indicator into the Uncrowded Hypervolume (UHV). The UHV takes dominated solutions into account and can therefore be used directly as an indicator in indicator-based multi-objective optimization. In UHV-based multi-objective optimization, multi-objective optimization problems are formulated as high-dimensional single-objective optimization problems, in which the UHV of a solution set of fixed size is maximized. As the resulting problem is high-dimensional, we approached it with GOMEA [37], which was explicitly developed to exploit weak dependencies between subsets of decision variables, thereby allowing these problems to be solved efficiently with a relatively small population size. Additionally, partial evaluations could be exploited to reduce the number of evaluations of the multi-objective problem.

A marginal linkage model does not model dependencies between variables pertaining to different solutions. When the solution sets that are being optimized are of small size, the dependencies between decision variables pertaining to different solutions are weak. In that case, UHV-GOMEA-Lm, which uses a marginal linkage model, can be used to efficiently solve the problem. However, especially when the desired number of solutions on the front is large, the dependencies between decision variables of solutions that are nearby each other increase in strength, making it less efficient or even impossible to solve it with UHV-GOMEA-Lm. We overcame this by introducing a linkage-tree (Lt) model for UHV optimization, which, combined

with UHV-GOMEA, is called UHV-GOMEA-Lt. Results obtained with UHV-GOMEA-Lt show that this limitation was overcome, at some additional cost of fitting a larger model.

An (efficient) implementation of the UHV has currently only been derived for bi-objective optimization problems. The hypervolume computation becomes rather expensive for $m \geq 3$ objectives. Currently, practical algorithms that can compute the hypervolume indicator $\mathcal{O}((m-1)N \log N)$ are available [69]. An efficient computation of the uncrowded distance in these efficient hypervolume computation algorithms would increase the applicability of the UHV-based approach to multi-objective optimization with more than two objectives.

8.1.6 Research question 6

How can smoothly navigable approximation sets be obtained efficiently and effectively via hypervolume-based bi-objective optimization, thereby making the selection of a desirable solution more intuitive for a decision maker?
(Chapter 7)

In Chapter 7, we parameterized approximation sets as smooth Bézier curves in decision space. We thereby explicitly enforced a form of smoothness between decision variables of neighboring solutions when the trade-off curve is traversed. For this, we adapted the parameterization based on the UHV that was introduced in Chapter 6. Solving this Bézier problem formulation with GOMEA, which we refer to as BezEA for short, was shown to be competitive to UHV-based optimization and domination-based MOEAs on various benchmark problems, while smoothness is guaranteed. We showed that approximation sets obtained with BezEA show a more smooth pattern in terms of decision variables when traversing the trade-off curve. This suggests that this approach will lead to a more intuitive and smoothly navigable approximation set for real-world optimization problems. To the best of our knowledge, this is the first attempt to explicitly quantify and improve navigational smoothness of approximation sets.

The Bézier problem formulation is a constrained UHV-based optimization problem. The constraint is required to prevent having dominated solutions in the approximation set. The result is however that the initial part of the optimization process, when the aim is to obtain a solution set that does not violate the constraint, can be rather expensive. This is the reason that this Bézier problem formulation is inefficient for multi-objective multimodal optimization. A problem formulation

that more naturally directs the search to a single mode could overcome this, but it remains a challenge to determine such a formulation.

8.1.7 Research question 7

> Are smoothly navigable approximation sets for bi-objective treatment planning of prostate brachytherapy obtainable without too much cost in obtained trade-offs and computation time? *(Chapter 7)*

In Chapter 7, we applied BezEA to the bi-objective treatment planning problem. We showed that BezEA can efficiently solve the bi-objective treatment planning problem by exploiting that the dose distributions in the approximation set can be quickly interpolated when a Bézier parameterization is used. By doing so, BezEA can obtain smoothly navigable approximation sets at little to no loss in quality in terms of obtained trade-offs with the same computational budget as MO-GOMEA and UHVEA-GOMEA-Lm (also called UHVEA-gb). Differences in plan quality, in terms of the objectives of the bi-objective planning mode are small. This suggests that BezEA is a good alternative to MO-GOMEA while navigational smoothness is guaranteed. When traversing the trade-off curve obtained by BezEA, both the dwell times and dose-volume indices of treatment plans change in a smooth fashion. This indicates that the smoothness in terms of the dose distributions will also be better when the BezEA trade-off curve is traversed, compared to trade-off curves obtained by UHVEA-GOMEA-Lm and MO-GOMEA. However, domain experts will need to be consulted to verify that the navigation is indeed more smooth and thereby results in a faster and more intuitive plan selection.

The presented approach is furthermore rather flexible in that it allows the user to enforce restrictions upon approximation sets. Instead of enforcing a form of smooth navigability on approximation sets, as done in this work, one could enforce that the obtained approximation sets reside in a single mode. This could result in a very interesting approach for multimodal multi-objective optimization problems. When an approximation set spreads out over multiple modes, it is generally not smooth, as one might jump back and forth between different modes when the front is traversed. To improve smoothness in that case, the approximation set can be restricted to only a single mode. This means that BezEA can in principle be used for multimodal optimization of multi-objective problems. However, it is not very efficient in its current form. One of the reasons for this is the usage of a hard constraint to ensure that solutions on the Bézier curve do not dominate each other in objective space.

This hard constraint could cause approximation sets that reside in multiple niches to be 'stuck'. Using a niching framework with restarts, such as HillVallEA, can circumvent this problem, but we observed in preliminary results that this approach is rather inefficient. Overcoming this, by using a different parameterization, or a different constraint definition, would make BezEA an interesting approach for multimodal optimization of multi-objective problems.

8.2 On model-based evolutionary algorithms

A difficulty in the design and comparison of black-box optimization algorithms, such as evolutionary algorithms, is nicely captured in the no-free-lunch theorem for search and optimization [208]. This theorem states that there can be no search algorithm that outperforms *all* other algorithms on *all* problems [93]. An important implication of the no-free-lunch theorem is that when you improve the performance of an algorithm for one problem, it will worsen for another problem. The overall *best* algorithm does thus not exist, and research should therefore not focus on this either.

Instead, when the aim of research is to design and develop widely applicable black-box optimization algorithms, it should focus on algorithms that have a decent out-of-the-box performance on a set of reasonable optimization problems and are easy to customize to the problem at hand. Customization of an algorithm to the problem at hand is in practice often performed by tuning the hyper parameters of the algorithm, aimed to improve performance in terms of runtime and achievable solution quality. Alternatively, adaptive approaches can be used to overcome manual parameter tuning, such as the increasing population size scheme [37] that is used in HillVallEA, but better performance on specific problems can generally be obtained by problem-specific hyper parameter tuning. It was for example found that tuning the hyper parameters of GOMEA on the brachytherapy treatment planning problem on a training set of a few patients could improve its performance [35]. The new configuration of GOMEA could then be used to optimize treatment plans for future patients, because we may assume that the structure of the optimization problems of these future patients is very similar to the training set.

Metaheuristics, among them evolutionary algorithms, are aimed to quickly obtain reasonably good solutions. In doing so, they are often engineered in such a way that it is no longer possible to establish a sufficient mathematical foundation in which one can prove convergence to optimality for a class of problems. The lack of a mathematical foundation for evaluating and comparing algorithmic performance

is typically overcome by empirically benchmarking algorithms via a set of problems with known properties. To this end, many high-quality benchmark sets have been developed [54, 84, 85, 101]. Good benchmark sets contain a variety of functions, each with clear and understandable features and characteristics that help the user to understand why an algorithm might or might not work, and help to overcome the lack of theoretical results, especially for real-world optimization problems.

Many metaheuristics are being developed, that are motivated by a nature-inspired metaphor [65]. Being inspired by other fields, such as biology, can lead to very appealing methods, of which neural networks, that 'mimic' the brain, are a well-known example. However, a notable issue in the field of metaheuristics is the tendency for authors to use terminology that is derived from the domain of inspiration, rather than the broader domains of metaheuristics and optimization [128]. This makes it difficult to understand how these methods work, and even more important, why they work. It can furthermore be hard to determine as a reader whether such algorithms are simply another flavor of an existing algorithm, hidden behind a novel nature-inspired metaphor, or actually contribute to the general knowledge and understanding of black-box optimization. Model-based algorithms partially overcome the lack of a theoretical foundation, by explicitly modeling certain problem features, such that it can be measured or tested to what extend the problem at hand fits the model. In this thesis, we for example explicitly incorporated multimodality into the optimization model. When using a model to guide the search, one can easily verify whether the problem at hand fits the model, which can give a good indication about performance during optimization. However, in a black-box scenario, it is unknown beforehand if these problem features will be available, and whether it is therefore a sensible choice to apply a certain model-based evolutionary algorithm.

To this extent, exploratory landscape analysis (ELA) [153] is a related and relevant research area, in which an analysis of the optimization problem to be solved is performed by spending a certain number of function evaluations and computing various characteristics. Based on this analysis, the user gets an impression of the features that are present in the problem at hand, such as the skewness of the objective function, its degree of convexity, or the estimated number of peaks. This can help practitioners to select the right algorithm for their problem, and help researchers to design algorithms that can handle these features well. Automatic algorithm selection [25] based on ELA has furthermore been developed, in which problem features are being correlated to algorithmic performance. Algorithm selection highlights another interesting research avenue for algorithm design, in which (black-box) optimization algorithms are designed

for a small niche of problems that exhibits certain features, and for which there is no or only few well-performing algorithms. ELA could guide the design of problem sets towards sets with a wide diversity of problems that mimic features of real-world problems. This can help to minimize the gap between performance of algorithms on benchmark sets and on real-world problems. Due to a lack of a diverse set of benchmark problems, fundamental limitations of algorithms could remain hidden [165].

8.3 On single-objective multimodal optimization

Multimodal optimization is commonly formulated as the problem of obtaining all global optima of a (single-objective) optimization problem [125]. In numerical real-world optimization, it is however rarely the case that multiple solutions have exactly the same globally best objective value. It would therefore be sensible to broaden this definition and aim to obtain many, if not all, high-quality locally optimal solutions, which should be distinguishable from low-quality local optima by some (problem-specific) threshold. However, not all high-quality solutions are local optima, as local optimality is a property of a solution with respect to its neighbors in decision space, and is not related to the fitness value of that solution itself. Especially when the fitness landscape has a plateau, or a very flat slope, many *diverse* solutions can be considered to be of high quality while still belonging to the same mode. Multimodal optimization algorithms do not currently search for these solutions. It is rather strange to only consider two high-quality solutions desirable if there are lower-quality solutions in between [160]. A future research direction for multimodal optimization, or rather an extension of it, can therefore be to reformulate the problem as the search for a diverse set of high-quality solutions, in which diversity can be understood as a distance in decision space for example. Formulating the optimization aim as the search for diverse high-quality solutions furthermore allows to deal better with noisy functions, for which local optima are not properly defined, or with problems with too many high-quality local optima to reasonably aim for in the optimization process. Dealing with both diversity and quality is a topic that is covered to great extent in multi-objective optimization, and could inspire future research directions for multimodal optimization.

8.4 On multi-objective optimization

Multi-objective optimization has great practical value, which can be attributed to a key idea in economics,"There ain't such thing as a free lunch" [150]. This idea, after which the no-free-lunch theorem is named, suggests that it is impossible to get something for nothing, and a trade-off always has to be made. In order to make insightful decisions, it helps to know which good options are available to choose from. To find those options, multi-objective optimization can be employed.

Domination-based Multi-Objective Evolutionary Algorithms (MOEAs) are classically state of the art in multi-objective optimization. These algorithms particularly excel in quickly obtaining reasonable good solutions, as we observed in Chapter 6. In domination-based MOEAs, the population guides the search and simultaneously functions as approximation set. Especially in the beginning of the optimization process, when most (or all) solutions are of low quality, the population as a whole can be used as a guide towards high-fitness regions in decision space. However, later on in the search process, when most (or all) solutions are of high fitness, these solutions typically need to drift apart to be able to improve diversity. Using the population as a whole to estimate a single (unimodal) search distribution to guide individual solutions is then no longer sufficient to find improvements. The MOEAs we considered in this work, MO-GOMEA and MaMaLGaM, divide the population into (overlapping) clusters, and estimate multiple search distributions, to be able to focus on different sub-parts of the search space. However, as long as a cluster is aimed to find improvements for multiple solutions, it will stagnate at some point. In UHV-GOMEA, this stagnation does not occur, as each solution is individually modeled. An interesting future research direction is to combine the advantages of hyper-volume based optimization in the final part of the search, with the rapid initial gains obtained by domination-based methods. We introduced a first simple hybrid approach to do this in Chapter 6, which was shown to already outperform both individual methods on many problems. It is however a very crude approach, and a more fundamental integration of both concepts is key to next-generation multi-objective optimization algorithms.

The aim of multi-objective optimization is to find a good approximation of the Pareto front. This aim in itself is two-sided, as getting close to the front conflicts with getting a good spread of solutions along the front, given a limited amount of time [33]. Domination-based MOEAs do generally not explicitly formulate these two aims. For example, in case of MAMaLGaM or MO-GOMEA, an adaptive objective-space discretization scheme [134] is used to enhance diversity in case a user-defined number of non-dominated solutions has been obtained. It is in this

case rather unclear how proximity is traded-off with diversity. In Chapter 6, this is made explicit via the definition of the UHV. An explicitly formulated aim allows to investigate whether algorithms show convergence to optimality, and at what rate. While convergence to optimality is not essential for successful real-world optimization, it is a powerful indication that an algorithm is fundamentally capable of solving a well-defined problem.

We showed that when the number of solutions along the front is small, the dependencies between solutions can be ignored to efficiently solve the problem. However, when many solutions are desired along the front, optimization will stagnate when these dependencies are ignored. The problem of evenly distributing solutions along the Pareto front or surface resembles an optimal scattering problem, also known as a packing problem, which has some remarkable properties that are not present in common benchmark problems [29]. These packing problems could give insight in how dependencies between decision variables of neighboring solutions along the front should be handled. We showed in Chapter 6 that handling these dependencies is essential for finding large solution sets with high accuracy.

8.4.1 Gradient-based multi-objective optimization

When gradient information of a problem is available, it is almost always beneficial to exploit this in the search process. However, even when the gradients of the individual objective functions are available, it is not straightforward to combine them such that an approximation set with diverse solutions can be obtained [28, 171, 188]. The UHV problem formulation however allows for a straightforward use of the gradients of the multiple objective functions to perform gradient-based multi-objective optimization, as we recently found [58]. Such a method converges linearly to an approximation set with optimal hypervolume, when using out-of-the-box first-order gradient-descent methods. This also creates new possibilities for hybrid optimization methods that combine gradient-information to speed up search, and use a population to avoid getting stuck in a local optimum.

8.4.2 Multimodality in multi-objective optimization

In general, multi-objective fitness landscapes are not as well understood as their single objective variants, as it is extremely challenging to imagine interaction between multiple decision variables and multiple objectives simultaneously [107, 108]. Research and analysis of multi-objective problems is therefore often restricted to properties of the Pareto set and Pareto front. Recent work has shown that multimodality might even help multi-objective optimization, in contrast to its impact on single-objective optimization, for which it complicates the search [108].

In single-objective optimization, solutions can be strictly ordered based on their fitness value. A solution is then locally optimal when it is strictly better than all solutions in some neighborhood. In multi-objective optimization, this strict ordering of solutions no longer exists, as solutions can be mutually non-dominated. In that case, one could distinguish two types of local optimality: *strong* local optimality for solutions that dominate all solutions in their neighborhood, and *weak* local optimality for solutions that are not dominated by any solution in their neighborhood. In most continuous (bi-objective) optimization problems we considered in this work, the Pareto set is a continuous curve in decision space, implying that strong (locally) optimal solutions do not exist for these problems, and weak locally optimal solutions should be the aim of optimization in practice. In general one might however question whether a single-solution based view of these landscapes is most insightful. The aim of multi-objective optimization is to obtain an approximation set of diverse high-quality solutions. To that extent, a multimodal approach to multi-objective optimization might be better defined in terms of (locally) optimal approximation sets, as we did heuristically in Chapter 5.

The BezEA approach we discussed in Chapter 7 does not explicitly perform multimodal optimization, but approximation sets are generally only smoothly navigable when they reside in a single mode. Even though BezEA is not directly usable for multimodal optimization, it could be used as inspiration for future research.

Problems that have more than three objectives are classified as many-objective optimization problems. Domination-based MOEAs are typically ineffective for solving many-objective optimization problems. This is mainly due to the large-dimensional objective space, which causes most solutions in the population to be non-dominated. In that case, selection pressure disappears, resulting in an unguided search, as we discussed in Chapter 6. Several approaches to overcome this limitation have been proposed. One approach is to relax the definition of domination by considering epsilon-dominance [217]. Alternatively, one can consider indicator-based optimization approaches, based on the R2-indicator [83] or the hypervolume indicator, as we did in Chapter 6. However, it can already be difficult for a decision maker to intuitively select the most desirable solution in case the problem at hand has two objectives, and the Pareto front is thus easily visualized as a trade-off curve (Chapter 7). The Pareto front of a many-objective optimization problem is a high-dimensional manifold, and can thus no longer be naturally visualized. One can then rely on projections, but an overview of the entire approximation set at one glance is no longer available. It is thus questionable whether it is useful, from a practical point of view, to aim optimization of many-

objective optimization problems towards obtaining approximation sets as a whole.

From an algorithm-design perspective, there are interesting challenges in multi- and many-objective multimodal optimization. However, one should keep in mind that simply presenting a decision maker with more (good) options to choose from is not always advantageous. Information overload, also known as "infobesity" or "infoxication", can make decision making particularly difficult, and even result in poorer decisions. When we humans are overloaded with options, we tend to avoid the problem by postponing decision making, we make poorer decisions and are less satisfied with the chosen solution [187]. Insightful and easy decision making should therefore be of key importance in the design and development of metaheuristics for obtaining diverse high-quality solutions.

In practice, instead of considering many-objective optimization problems, or multimodal multi-objective optimization, it might therefore be more interesting to consider interactive optimization approaches [185, 193], goal programming [45], or the worst-case grouping that we considered in our work to reduce the number of objectives to be optimized. In these approaches, the decision maker guides the search, and the search informs the decision maker at the same time, thereby integrating the optimization and decision making processes by presenting few options at a time. Additionally, approaches such as BezEA can simplify decision making as not all solutions have to be inspected due to the imposed smoothness, allowing the decision maker to quickly inspect properties of an entire approximation set by inspecting only one or a few solutions.

8.5 On brachytherapy treatment planning

The difficulty in validating a novel (brachytherapy) treatment planning approach is that it is unknown a priori which plan would turn out to be the best for the current patient. Even though there has been a lot of work in correlating specific quality indices of treatment plans to different aspects of the treatment outcome [26, 92, 100, 148, 190], there is still a lot of uncertainty. Additionally, it is not directly clear how to combine all this information properly into an optimization approach. For the bi-objective treatment planning method, we developed an approach to translate treatment protocols, consisting of multiple planning criteria, directly into two objective functions [38, 131, 147]. To verify whether this would lead to desirable plans, we let physicians determine whether the resulting plans are clinically acceptable to treat the patient with. This is similar to how treatment planning currently works in the clinic, where, among others, the treating physician

approves the final treatment plan. The physician can base his or her desires for a treatment plan on guidelines or the planning criteria in the clinical protocol, on patient-specific information, or on experience, general knowledge, or intuition.

A desirable and rather straightforward extension of bi-objective treatment planning, as employed in this work for high-dose-rate prostate brachytherapy, is to apply it to other tumor sites and different treatment protocols. The step that is required to extend this treatment planning approach to other tumor sites, or other treatment protocols, is the formulation of different objective functions. The framework, as used in this work, in which clinical planning criteria are translated into two objectives, can help in doing so to great extent [131]. For prostate brachytherapy, we showed that clinically acceptable plans could be obtained by configuring the objectives directly based on the planning criteria. This might however not be the case for all other treatment sites or treatment protocols. To overcome this, additional planning criteria need to be determined, for example to take dose homogeneity or normal tissue dose into account.

In [42], an automatic treatment planning approach for HDR prostate brachytherapy has been developed, by making use of a so-called wish-list that contains the optimization protocol, hard constraints and planning objectives. After the optimization of each objective function, a new constraint is added to the wish-list to ensure that the obtained function value is maintained while minimizing objectives with lower priority. The concept of a wish-list, in which lower-priority optimization aims can be included without worsening the planning criteria can be a valuable extension to the objectives framework of the bi-objective planning model used in this work. It allows for example the inclusion of a dose-homogeneity-related planning criterion, which is in conflict with both coverage- and sparing-related planning criteria. Via this wish-list approach, a single plan is however obtained, and the same wish-list is used for all patients, although achievable objective function values are adapted during optimization. It is therefore important that this wish-list gives a desirable plan for all patients. The bi-objective planning model overcomes this by presenting a planner with a diverse set of plans, from which a desirable plan can be selected. In this selection process, the planner can incorporate details about the patient that are not known to the planning method, without having to rerun the optimization.

8.5.1 Gradient-based bi-objective treatment planning

Population-based gradient methods could have an application in bi-objective treatment planning for brachytherapy as well. The gradients of the DVIs in the bi-objective treatment planning model are not analytically available, but different

approximations have been derived [82]. An interesting short-term research avenue is therefore to combine this method with the bi-objective treatment planning model, which could greatly reduce treatment computation time. While the current approach is fast enough for clinical practice, a further speed up could allow for model extensions, such as finding multiple approximation fronts in a multimodal fashion.

8.5.2 Data-driven treatment planning

An alternative approach to construct the optimization objectives, that is not (yet) common in brachytherapy, but more commonly applied in external beam radiation therapy (EBRT), is to leverage a database of previously used treatment plans [210]. From the database, plans of patients that were successfully treated can be used to automatically construct the objective function(s) [209]. Plans of patients that suffered from severe adverse effects could be used to furthermore determine undesirable aspects of treatment plans. To do so, it is important that (long-term) adverse effects are properly monitored, as this information is essential to determine whether used plans are of high quality, and which aspects are desirable.

Brachytherapy is as not wide-spread used as EBRT, which can be attributed, among other factors, to the fact that this treatment modality is rather labor-intensive due to the catheter placement being performed manually by a physician [71]. The lack of sufficient historical data made the development of data-driven treatment planning methods a less promising research avenue. Sharing more data among medical institutes can greatly help in improving treatment planning. Even though regulations are of crucial importance to protect the privacy of individuals, they do complicate data exchanges. Approaches such as distributed learning approaches could make knowledge sharing easier [57]. In distributed learning, the data remains in the hands of the clinic itself, and instead, the model is sent there to be locally trained. Only the relevant information, such as aspiration values for planning criteria, would be extracted and shared in such a scenario.

8.5.3 Beyond dwell time optimization

We specifically focused on dwell time optimization, which we considered as a separate part of the treatment planning process. However, in practice, it is part of a chain of steps that forms the entire treatment. The treatment starts with the placement of catheters in the operating room. These catheters are currently placed manually, although automatic catheter placement methods are being investigated [214]. In our clinic, these catheters are placed according to a pre-plan that is made based on ultrasound imaging. Catheter placement itself can be formulated

as an optimization problem. Traditionally, catheters were placed based on simple heuristics, such as equal spacing of catheters, often applied via the use of a template. In work done in parallel to the work in this thesis, it has been investigated whether the positions of the catheters can be optimized simultaneously with the dwell times [200, 203]. Although the computational effort required to solve the resulting optimization problem is significantly larger, the use of GPU based computations can make this approach clinically feasible. A difficulty of the analysis of catheter position optimization is however that it is hard or impossible to place catheters exactly as specified in the pre-plan. Additionally, in our clinic, the positioning of the patient is different during catheter placement and treatment, which would need to be taken into account in the treatment planning process.

After the catheters have been placed, at our institute, MRI scans are made of the patient in treatment position, which show both the catheter positions and the important structures. On these scans, important structures are delineated. This is currently performed manually, and is rather time consuming, but automatic image segmentation software based on machine learning approaches is an active research topic [46]. However, these methods are not yet clinically available in the treatment planning software of most vendors. Additionally, due to the limited image resolution and non-static patient anatomy, these delineations or segmentations inhibit some intrinsic geometric uncertainty [207]. This uncertainty can be taken into account in the optimization process via a technique known as probabilistic planning [76, 196]. In probabilistic planning, voxels belong to a certain structure with a certain probability. This could be a first step towards robust bi-objective treatment planning, in which treatment plans are made robust against variations that are likely to occur [207]. An approach to robust bi-objective treatment planning has been developed in [199, 201], in which different uncertainties are directly taken into account in the optimization process by optimizing a worst-case scenario. Robust optimization is furthermore a very interesting technique in which the uncertainties of manual catheter placement can be taken into account when determining a pre-plan.

8.5.4 Extensions to other radiotherapy modalities

In the optimization of the bi-objective planning model with MO-GOMEA, it is exploited that radiation dose is cumulative. This implies that dwell times can be updated independently without having to recompute the entire dose distribution. That is, an efficient update can be performed when only a few dwell times change. Bi-objective treatment planning with MO-GOMEA can easily be extended to other

radiotherapy modalities that also satisfy this property. One such example is low-dose-rate brachytherapy, in which radioactive seeds are permanently implanted in the tumor. For low-dose-rate brachytherapy, the to-be-optimized variables are the seed positions, and not the corresponding times. Seed position optimization of low-dose-rate brachytherapy shares many properties with catheter position optimization of high-dose-rate prostate brachytherapy, which can also be performed efficiently in a bi-objective fashion with MO-GOMEA [200, 203].

Treatment planning for EBRT is, in itself, a very different problem. The number of decision variables that needs to be optimized in EBRT treatment planning is typically at least a magnitude larger than the number of dwell positions that are optimized in brachytherapy treatment planning. In EBRT, a more precise and computationally expensive dose model is furthermore used [179], which is required as lower doses are applied, and dose spreads out more to other regions of the body compared to high-dose-rate brachytherapy. The goal of treatment planning for EBRT is however the same as it is for brachytherapy, as both aim to treat the patient by maximizing dose to the tumor while minimizing dose to healthy organs. The idea of presenting multiple plans, and an intuitive visualization of these trade-offs, is therefore a very interesting avenue for multi-objective optimization in EBRT treatment planning.

8.6 Implications for society

We live in an age of *algorithms* [1]. More and more tasks that were previously deemed to require human intelligence can now be performed by machines or computers, controlled by algorithms. The term *artificial intelligence* or *AI* is often used to describe these algorithms, among which evolutionary algorithms, that mimic cognitive functions we associate with the human mind, such as *learning* and *problem solving* [186]. Ironically, as algorithms become increasingly capable, tasks that can be performed with artificial intelligence are often no longer considered to require artificial intelligence, but merely a computational method: a phenomenon known as the *AI effect* [151]. It is therefore joked that "AI is whatever hasn't been done yet" [136]. By this effect, AI can thus do what we think is impossible without human intellect, which fuels discussions about the ethics of creating artificial beings endowed with human-like intelligence. Especially in popular culture, AI is often depicted as a machine taking over the world, such as in "I, Robot" [5] and "Ex Machina" [73]. To take away this frightening and unknown side of AI, it helps to understand why it solves problems in a certain way. In this thesis, instead

of letting the algorithm decide upon which solution is most desirable, we let the algorithm find a number of diverse high-quality solutions. Human intelligence, and external information, can then be employed to select the most desirable solution. Especially when the problems being solved have a large impact, such as in medicine, it is extremely important that the right solutions are chosen. The methods and approaches we explore in this thesis exploit computational power to do the computationally demanding search to find diverse high-quality solutions, while it remains up to the domain experts to use their expertise to select the best solution for the problem at hand. This is key in the development and application of high-impact software solutions.

A direct impact of this work is the clinical use of bi-objective treatment planning in the Amsterdam University Medical Centers, location Academic Medical Center, where it is referred to as *BRachytherapy via artificially Intelligent GOMEA Heuristic based Treatment planning* (BRIGHT). As of spring 2020, prostate cancer patients treated with HDR brachytherapy are treated with plans based on BRIGHT. Figure 1.4 is the trade-off curve obtained for the first patient treated with BRIGHT. A follow up project has furthermore been funded to extend this treatment planning approach and to apply it in the treatment of cervical cancer.

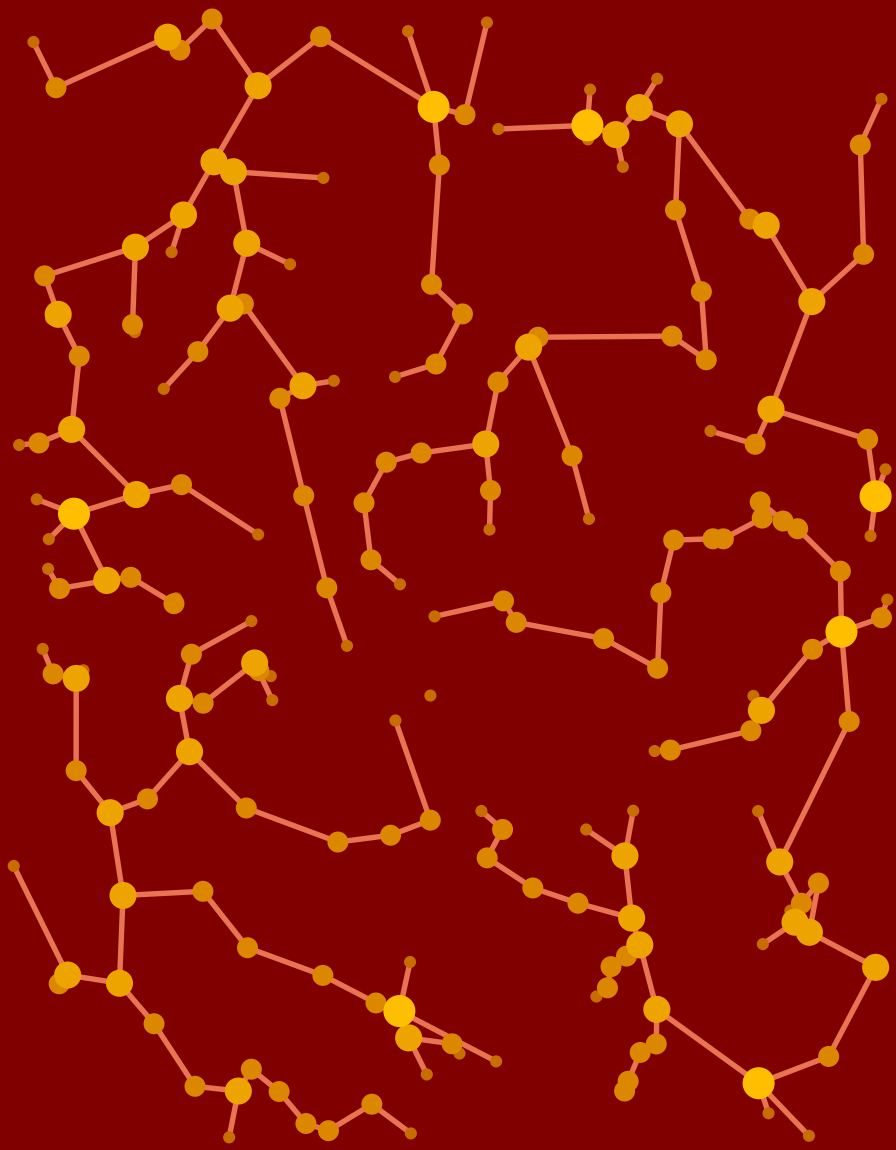

Bibliography

[1] Abiteboul, S., and Dowek, G. *The Age of Algorithms*. Cambridge University Press, 2020.

[2] Ahrari, A., Deb, K., and Preuss, M. Benchmarking covariance matrix self adaption evolution strategy with repelling subpopulations for GECCO 2017 competition on multimodal optimization. *COIN Report 2017014* (2017).

[3] Ahrari, A., Deb, K., and Preuss, M. Multimodal Optimization by Covariance Matrix Self-Adaptation Evolution Strategy with Repelling Subpopulations. *Evolutionary Computation 25*, 3 (2017), 439–471.

[4] Alterovitz, R., Lessard, E., Pouliot, J., Hsu, I.-C. J., O'Brien, J. F., and Goldberg, K. Optimization of HDR brachytherapy dose distributions using linear programming with penalty costs. *Medical Physics 33*, 11 (2006), 4012–4019.

[5] Asimov, I. *I, Robot*. Gnome Press, 1950.

[6] Audet, C., and Hare, W. *Derivative-Free and Blackbox Optimization*. Springer International Publishing, 2017.

[7] Auger, A., Bader, J., Brockhoff, D., and Zitzler, E. Theory of the hypervolume indicator. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2009), ACM Press, pp. 87–102.

[8] Auger, A., Bader, J., Brockhoff, D., and Zitzler, E. Theory of the hypervolume indicator: Optimal μ-distributions and the choice of the reference point. In *Proceedings of the Tenth ACM SIGEVO Workshop on Foundations of Genetic Algorithms - FOGA* (2009), ACM Press, pp. 87—102.

[9] Auger, A., and Hansen, N. A restart CMA evolution strategy with increasing population size. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (2005), IEEE Press, pp. 1769–1776.

[10] Bachand, F., Martin, A. G., Beaulieu, L., Harel, F., and Vigneault, É. An eight-year experience of HDR brachytherapy boost for localized prostate cancer: Biopsy and PSA outcome. *International Journal of Radiation Oncology Biology Physics 73*, 3 (2009), 679–684.

[11] Bäck, T. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, Inc., USA, 1996.

[12] Bader, J., and Zitzler, E. Hype: An algorithm for fast hypervolume-based many-objective optimization. *Evolutionary Computation 19*, 1 (2011), 45–76.

[13] Balvert, M., Gorissen, B. L., den Hertog, D., and Hoffmann, A. L. Dwell time modulation restrictions do not necessarily improve treatment plan quality for prostate HDR brachytherapy. *Physics in Medicine & Biology 60*, 2 (2015), 537–548.

[14] Bangma, C. H. *Leerboek Urologie*. Bohn Stafleu van Loghum, 2013.

[15] Bangma, C. H., Valdagni, R., Carroll, P. R., Van Poppel, H., Klotz, L., and Hugosson, J. Active surveillance for low-risk prostate cancer: developments to date. *European Urology 67*, 4 (2015), 646–648.

[16] Baumann, M., and Petersen, C. TCP and NTCP: A basic introduction. *Rays - International Journal of Radiological Sciences 30*, 2 (2005), 99–104.

[17] Bélanger, C., Cui, S., Ma, Y., Després, P., Adam M Cunha, J., and Beaulieu, L. A GPU-based multi-criteria optimization algorithm for HDR brachytherapy. *Physics in Medicine & Biology 64*, 10 (2019), 105005.

[18] Berghammer, R., Friedrich, T., and Neumann, F. Set-based multi-objective optimization, indicators and deteriorative cycles. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2010), ACM Press, pp. 495–502.

[19] Berghammer, R., Friedrich, T., and Neumann, F. Convergence of set-based multi-objective optimization, indicators and deteriorative cycles. *Theoretical Computer Science 456* (2012), 2–17.

[20] Beume, N., Fonseca, C., López-Ibáñez, M., Paquete, L., and Vahrenhold, J. On the complexity of computing the hypervolume indicator. *IEEE Transactions on Evolutionary Computation 13*, 5 (2009), 1075–1082.

[21] Beume, N., Naujoks, B., and Emmerich, M. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research 181*, 3 (2007), 1653–1669.

[22] Beyer, H. G., and Sendhoff, B. Covariance Matrix Adaptation Revisited - the CMSA Evolution Strategy. In *Parallel Problem Solving from Nature - PPSN* (2008), vol. 5199, Springer, pp. 123–132.

[23] Bhardwaj, P., Dasgupta, B., and Deb, K. Modelling the Pareto-optimal set using B-spline basis functions for continuous multi-objective optimization problems. *Engineering Optimization 46*, 7 (2014), 912–938.

[24] Birattari, M., Stützle, T., Paquete, L., and Varrentrapp, K. A racing algorithm for configuring metaheuristics. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2002), ACM Press, pp. 11–18.

[25] Bischl, B., Mersmann, O., Trautmann, H., and Preuss, M. Algorithm selection based on exploratory landscape analysis and cost-sensitive learning. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2012), ACM Press, pp. 313–320.

[26] Boersma, L. J., Van Den Brink, M., Bruce, A. M., Shouman, T., Gras, L., Te Velde, A., and Lebesque, J. V. Estimation of the incidence of late bladder and rectum complications after high-dose (70-78 Gy) conformal radiotherapy for prostate cancer, using dose-volume histograms. *International Journal of Radiation Oncology Biology Physics 41*, 1 (1998), 83–92.

[27] Bosman, P. A. N. The anticipated mean shift and cluster registration in mixture-based EDAs for multi-objective optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2010), ACM Press, pp. 351–358.

[28] Bosman, P. A. N. On gradients and hybrid evolutionary algorithms for real-valued multiobjective optimization. *IEEE Transactions on Evolutionary Computation 16*, 1 (2011), 51–69.

[29] Bosman, P. A. N., and Gallagher, M. The importance of implementation details and parameter settings in black-box optimization: a case study on Gaussian estimation-of-distribution algorithms and circles-in-a-square packing problems. *Soft Computing 22* (2018), 1209–1223.

[30] Bosman, P. A. N., and Grahl, J. Matching inductive search bias and problem structure in continuous estimation-of-distribution algorithms. *European Journal of Operational Research 185*, 3 (2008), 1246–1264.

[31] Bosman, P. A. N., Grahl, J., and Thierens, D. Benchmarking parameter-free AMaLGaM on functions with and without noise. *Evolutionary Computation 21*, 3 (2013), 445–469.

[32] Bosman, P. A. N., and Thierens, D. Multi-objective optimization with diversity preserving mixture-based iterated density estimation evolutionary algorithms. *International Journal of Approximate Reasoning 31*, 3 (2002), 259–289.

[33] Bosman, P. A. N., and Thierens, D. The balance between proximity and diversity in multiobjective evolutionary algorithms. *IEEE Transactions on Evolutionary Computation 7*, 2 (2003), 174–188.

[34] Bouter, A., Alderliesten, T., Bel, A., Witteveen, C., and Bosman, P. A. N. Large-scale parallelization of partial evaluations in evolutionary algorithms for real-world problems. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2018), ACM Press, pp. 1199–1206.

[35] Bouter, A., Alderliesten, T., Pieters, B. R., Bel, A., Niatsetski, Y., and Bosman, P. A. N. Bi-objective optimization of dosimetric indices for HDR prostate brachytherapy within 30 seconds. *Radiotherapy & Oncology 133*, Supplement 1 (2019), S199–S200.

[36] Bouter, A., Alderliesten, T., Pieters, B. R., Bel, A., Niatsetski, Y., and Bosman, P. A. N. GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy. *Medical Physics 46*, 3 (2019), 3776–3787.

[37] Bouter, A., Alderliesten, T., Witteveen, C., and Bosman, P. A. N. Exploiting linkage information in real-valued optimization with the real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2017), ACM Press, pp. 705–712.

[38] Bouter, A., Luong, N. H., Alderliesten, T., Witteveen, C., and Bosman, P. A. N. The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2017), ACM Press, pp. 537–544.

[39] Boyd, S., and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.

[40] Brahme, A. Individualizing cancer treatment: Biological optimization models in treatment planning and delivery. *International Journal of Radiation Oncology Biology Physics 49*, 2 (2001), 327–337.

[41] Bray, F., Lortet-Tieulent, J., Ferlay, J., Forman, D., and Auvinen, A. Prostate cancer incidence and mortality trends in 37 European countries. *European Journal of Cancer 46* (2010), 3040–3052.

[42] Breedveld, S., Bennan, A. B. A., Aluwini, S., Schaart, D. R., Kolkman-Deurloo, I. K. K., and Heijmen, B. J. M. Fast automated multi-criteria planning for HDR brachytherapy explored for prostate cancer. *Physics in Medicine & Biology 64*, 20 (2019), 205002.

[43] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing 16*, 5 (1995), 1190–1208.

[44] Catalona, W. J. Prostate cancer screening. *The Medical clinics of North America 102*, 2 (2018), 199–214.

[45] Charnes, A., Cooper, W., and Ferguson, R. Optimal estimation of executive compensation by linear programming. *Management Science 1* (1955), 138–151.

[46] Clark, T., Zhang, J., Baig, S., Wong, A., Haider, M. A., and Khalvati, F. Fully automated segmentation of prostate whole gland and transition zone in diffusion-weighted MRI using convolutional neural networks. *Journal of Medical Imaging 4*, 4 (2017), 041307.

[47] Craft, D. L., Hong, T. S., Shih, H. A., and Bortfeld, T. R. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *International Journal of Radiation Oncology Biology Physics 82*, 1 (2012), e83–e90.

[48] Cui, S., Després, P., and Beaulieu, L. A multi-criteria optimization approach for HDR prostate brachytherapy: I. Pareto surface approximation. *Physics in Medicine & Biology 63* (2018), 205004.

[49] Cui, S., Després, P., and Beaulieu, L. A multi-criteria optimization approach for HDR prostate brachytherapy: II. Benchmark against clinical plans. *Physics in Medicine & Biology 63* (2018), 205005.

[50] De Jong, K. A. An analysis of the behaviour of a class of genetic algorithms. *PhD Thesis, University of Michigan* (2015).

[51] Deb, K. An efficient constraint handling method for genetic algorithms. *Computer Methods in Applied Mechanics and Engineering 186*, 2 (2000), 311–338.

[52] Deb, K. *Multi-objective Optimization*. Wiley, Chichester, UK, 2001.

[53] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation 6*, 2 (2002), 182–197.

[54] Deb, K., Thiele, L., and Zitzler, E. Scalable multi-objective optimization test problems. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (2002), vol. 1, IEEE Press, pp. 825–830.

[55] Deb, K., and Tiwari, S. Omni-optimizer: A procedure for single and multi-objective optimization. In *Evolutionary Multi-Criterion Optimization - EMO* (2005), Springer, pp. 47–61.

[56] Deist, T. M., and Gorissen, B. L. High-dose-rate prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing. *Physics in Medicine & Biology 61*, 3 (2016), 1155–1170.

[57] Deist, T. M., Jochems, A., van Soest, J., Nalbantov, G., Oberije, C., Walsh, S., Eble, M., Bulens, P., Coucke, P., Dries, W., Dekker, A., and Lambin, P. Infrastructure and distributed learning methodology for privacy-preserving multi-centric rapid learning health care: euroCAT. *Clinical and Translational Radiation Oncology 4* (2017), 24–31.

[58] Deist, T. M., Maree, S. C., Alderliesten, T., and Bosman, P. A. N. Multi-objective Optimization by Uncrowded Hypervolume Gradient Ascent. In *Parallel Problem Solving from Nature - PPSN* (2020), Springer, pp. 186–200.

[59] Dinkla, A. M., Pieters, B. R., Koedooder, C., Van Wieringen, N., Van Der Laarse, R., van der Grient, J. N., Rasch, C. R., Koning, C. C., and Bel, A. Improved tumour control probability with MRI-based prostate brachytherapy treatment planning. *Acta Oncologica 53*, 3 (2013), 658–665.

[60] Dinkla, A. M., van der Laarse, R., Kaljouw, E., Pieters, B. R., Koedooder, K., van Wieringen, N., and Bel, A. A comparison of inverse optimization algorithms for HDR/PDR prostate brachytherapy treatment planning. *Brachytherapy 14*, 2 (2015), 279–288.

[61] Dinkla, A. M., van der Laarse, R., Koedooder, C., Kok, P. H., van Wieringen, N., Pieters, B. R., and Bel, A. Novel tools for stepping source brachytherapy treatment planning: enhanced geometrical optimization and interactive inverse planning. *Medical Physics 42*, 1 (2015), 348–353.

[62] Emmerich, M., Beume, N., and Naujoks, B. An EMO algorithm using the hypervolume measure as selection criterion. In *Conference on Evolutionary Multi-Criterion Optimization - EMO* (2005), vol. 3410 of LNCS, Springer, pp. 62–76.

[63] Emmerich, M., Deutz, A., and Beume, N. Gradient-based/evolutionary relay hybrid for computing Pareto front approximations maximizing the S-metric. In *International Workshop on Hybrid Metaheuristics* (2007), vol. 4771 of LNCS, Springer, pp. 140–156.

[64] Fatyga, M., Williamson, J. F., Dogan, N., Todor, D., Siebers, J. V., George, R., Barani, I., and Hagan, M. A comparison of HDR brachytherapy and IMRT techniques for dose escalation in prostate cancer: A radiobiological modeling study. *Medical Physics 36*, 9 (2009), 3995–4006.

[65] Fausto, F., Reyna-Orta, A., Cuevas, E., Andrade, Á. G., and Perez-Cisneros, M. From ants to whales: metaheuristics for all tastes. *Artificial Intelligence Review 53*, 1 (2020), 753–810.

[66] Ferlay, J., Colombet, M., Soerjomataram, I. Dyba, T., Randi, G., Bettio, M., Gavin, A., Visser, O., and Bray, F. Cancer incidence and mortality patterns in Europe: Estimates for 40 countries and 25 major cancers in 2018. *European Journal of Cancer 103* (2018), 356–387.

[67] Fieldsend, J. E. Efficient real-time hypervolume estimation with monotonically reducing error. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2019), ACM Press, pp. 532–540.

[68] Fleischer, M. The measure of Pareto optima. applications to multi-objective metaheuristics. In *Conference on Evolutionary Multi-Criterion Optimization - EMO* (2003), vol. 2632, Springer, pp. 519–533.

[69] Fonseca, C., Paquete, L., and López-Ibáñez, M. An improved dimension-sweep algorithm for the hypervolume indicator. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (2006), IEEE Press, pp. 1157–1163.

[70] Foster, W., Cunha, J. A. M., Hsu, I. C., Weinberg, V., Krishnamurthy, D., and Pouliot, J. Dosimetric impact of interfraction catheter movement in high-dose rate prostate brachytherapy. *International Journal of Radiation Oncology Biology Physics 80*, 1 (2011), 85–90.

[71] Fricker, J. Brachytherapy: halting the spiral of decline. *Cancerworld Spring* (2019), 4–15.

[72] Gallier, J. *Curves and Surfaces in Geometric Modeling: Theory and Algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.

[73] Garland, P. *Ex Machina*. Universal Studios, United Kingdom, 2015.

[74] Glover, F., and Laguna, M. Tabu search. In *Handbook of Combinatorial Optimization*. Springer, 1998, pp. 2093–2229.

[75] Goldberg, D., and Richardson, J. Genetic algorithms with sharing for mulitmodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Application* - (1987), Lauwrence Erlbaum Associates, pp. 41–49.

[76] Gordon, J., Sayah, N., Weiss, E., and Siebers, J. Coverage optimized planning: probabilistic treatment planning based on dose coverage histogram criteria. *Medical Physics 37*, 2 (2010), 550–563.

[77] Gorissen, B. L., den Hertog, D., and Hoffmann, A. L. Mixed integer programming improves comprehensibility and plan quality in inverse optimization of prostate HDR brachytherapy. *Physics in Medicine & Biology 58*, 4 (2013), 1041–1057.

[78] Gronau, I., and Moran, S. Optimal implementations of upgma and other common clustering algorithms. *Information Processing Letters 104*, 6 (2007), 205–210.

[79] Guedea, F., Venselaar, J., Hoskin, P., Hellebust, T. P., Peiffert, D., Londres, B., Ventura, M., Mazeron, J. J., Limbergen, E. V., Pötter, R., and Kovacs, G. Patterns of care for brachytherapy in Europe: Updated results. *Radiotherapy & Oncology 97*, 3 (2010), 514–520.

[80] Guerreiro, A., Fonseca, C., and Paquete, L. Greedy hypervolume subset selection in low dimensions. *Evolutionary Computation 24*, 3 (2016), 521–544.

[81] Gupta, M. R., and Chen, Y. Theory and use of the EM algorithm. *Foundations and Trends in Signal Processing 4*, 3 (2010), 223–296.

[82] Guthier, C., Damato, A. L., Viswanathan, A. N., Hesser, J. W., and Cormack, R. A. A fast multi-target inverse treatment planning strategy optimizing dosimetric measures for high-dose-rate (HDR) brachytherapy. *Medical Physics 44*, 6 (2017), 4452–4462.

[83] Hansen, M., and Jaszkiewicz, A. Evaluating the quality of approximations of the non-dominated set. Tech. rep., IMM Technical report IMM-REP-1998-7, 1998.

[84] Hansen, N., Auger, A., Mersmann, O., Tusar, T., and Brockhoff, D. COCO: A platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* (2020), 1–31.

[85] Hansen, N., Auger, A., Ros, R., Finck, S., and Pošík, P. Comparing results of 31 algorithms from the black-box optimization benchmarking BBOB-2009. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2010), ACM Press, pp. 1689–1696.

[86] Hansen, N., and Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *IEEE Computational Intelligence Magazine 9*, 2 (2001), 159–195.

[87] Harik, G. R., and Lobo, F. G. A parameter-less genetic algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (1999), ACM Press, pp. 258–265.

[88] Hart, W. E., Krasnogor, N., and Smith, J. E. *Memetic Evolutionary Algorithms*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005, pp. 3–27.

[89] Hastie, T., Tibshirani, R., and Friedman, J. H. *The Elements of Statistical Learning: Data mining, Inference, and Prediction*. Springer, New York, NY, USA, 2009.

[90] Hauschild, M., and Pelikan, M. An introduction and survey of estimation of distribution algorithms. *Swarm and Evolutionary Computation 1*, 3 (2011), 111–128.

[91] Higham, N. Cholesky factorization. *Wiey Interdisciplinary Reviews: Computational Statistics 1*, 2 (2008), 251–254.

[92] Ho, A. Y., Burri, R. J., Cesaretti, J. A., Stone, N. N., and Stock, R. G. Radiation dose predicts for biochemical control in intermediate-risk prostate cancer patients treated with low-dose-rate brachytherapy. *International Journal of Radiation Oncology Biology Physics 75*, 1 (2009), 16–22.

[93] Ho, Y.-C., and Pepyne, D. L. Simple explanation of the no-free-lunch theorem and its implications. *Journal of Optimization Theory and Applications 115*, 3 (2002), 549–570.

[94] Holm, Å., Larsson, T., and Tedgren, Å. C. Impact of using linear optimization models in dose planning for HDR brachytherapy. *Medical Physics 39*, 2 (2012), 1021–1028.

[95] Holm, Å., Larsson, T., and Tedgren, Å. C. A linear programming model for optimizing HDR brachytherapy dose distributions with respect to mean dose in the DVH-tail. *Medical Physics 40*, 8 (2013), 081705.

[96] Hoskin, P. J., Colombo, A., Henry, A., Niehoff, P., Paulsen Hellebust, T., Siebert, F. A., and Kovacs, G. GEC/ESTRO recommendations on high dose rate afterloading brachytherapy for localised prostate cancer: An update. *Radiotherapy & Oncology 107*, 3 (2013), 325–332.

[97] Hoskin, P. J., Rojas, A. M., Bownes, P. J., Lowe, G. J., Ostler, P. J., and Bryant, L. Randomised trial of external beam radiotherapy alone or combined with high-dose-rate brachytherapy boost for localised prostate cancer. *Radiotherapy & Oncology 103*, 2 (2012), 217–222.

[98] Hou, Z., Li, G., and Bai, S. High dose versus conventional dose in external beam radiotherapy of prostate cancer: a meta-analysis of long-term follow-up. *Journal of Cancer Research and Clinical Oncology 141*, 6 (2015), 1063–1071.

[99] Hsu, I. C., Bae, K., Shinohara, K., Pouliot, J., Purdy, J., Ibbott, G., Speight, J., Vigneault, E., Ivker, R., and Sandler, H. Phase II trial of combined high-dose-rate brachytherapy and external beam radiotherapy for adenocarcinoma of the prostate: Preliminary results of RTOG 0321. *International Journal of Radiation Oncology Biology Physics 78*, 3 (2010), 751–758.

[100] Hsu, I. C., Hunt, D., Straube, W., Pouliot, J., Cunha, A., Krishnamurthy, D., and Sandler, H. Dosimetric analysis of radiation therapy oncology group 0321: The importance of urethral dose. *Practical Radiation Oncology 4*, 1 (2014), 27–34.

[101] Huband, S., Barone, L., While, L., and Hingston, P. A scalable multi-objective test problem toolkit. In *International Conference on Evolutionary Multi-Criterion Optimization - EMO* (2005), Springer, pp. 280–294.

[102] Igel, C., Hansen, N., and Roth, S. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation 15*, 1 (2007), 1–28.

[103] Ishibuchi, H., Hitotsuyanagi, Y., Tsukamoto, N., and Nojima, Y. Many-objective test problems to visually examine the behavior of multiobjective evolution in a decision space. In *Parallel Problem Solving from Nature - PPSN* (2010), Springer, pp. 91–100.

[104] Ishibuchi, H., Masuda, H., Tanigaki, Y., and Nojima, Y. Modified distance calculation in generational distance and inverted generational distance. In *International Conference on Evolutionary Multi-Criterion Optimization - EMO* (2015), Springer, pp. 110–125.

[105] Karger, C. P. Biological models in treatment planning. In *New Technologies in Radiation Oncology*. Springer, 2006, pp. 221–235.

[106] Kennedy, J. *Swarm Intelligence*. Springer US, Boston, MA, 2006, pp. 187–219.

[107] Kerschke, P., and Grimme, C. An expedition to multimodal multi-objective optimization landscapes. In *9th International Conference on Evolutionary Multi-Criterion Optimization - EMO* (2017), Springer-Verlag, pp. 329–343.

[108] Kerschke, P., Wang, H., Preuss, M., Grimme, C., Deutz, A. H., Trautmann, H., and Emmerisch, M. T. M. Search dynamics on multimodal multiobjective problems. *Evolutionary Computation 27*, 4 (2018), 577–609.

[109] Kim, Y., Hsu, I. C. J., and Pouliot, J. Measurement of craniocaudal catheter displacement between fractions in computed tomography-based high dose rate brachytherapy of prostate cancer. *Journal of Applied Clinical Medical Physics 8*, 4 (2007), 1–13.

[110] Kingma, D., and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[111] Klotz, L. Active surveillance for prostate cancer: Overview and update. *Current Treatment Options in Oncology 14*, 1 (2013), 97–108.

[112] Knowles, J. Local-search and hybrid evolutionary algorithms for Pareto optimization. Tech. rep., PhD thesis, University of Reading, 2002.

[113] Knowles, J., Thiele, L., and Zitzler, E. A tutorial on the performance assessment of stochastic multiobjective optimization. Tech. rep., Computer Engineering and Networks Laboratory (TIK), ETH Zurich – TIK Report 214, 2006.

[114] Kobayashi, K., Hamada, N., Sannai, A., Tanaka, A., Bannai, K., and M., S. Bezier simplex fitting: Describing Pareto fronts of simplicial problems with small samples in multi-objective optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2019), vol. 33, pp. 2304–2313.

[115] Kolkman-Deurloo, I. K. K., Deleye, X. G. J., Jansen, P. P., and Koper, P. C. M. Anatomy based inverse planning in HDR prostate brachytherapy. *Radiotherapy & Oncology 73*, 1 (2004), 73–77.

[116] Kovalchuk, N., Furutani, K. M., MacDonald, O. K., and Pisansky, T. M. Dosimetric effect of interfractional needle displacement in prostate high-dose-rate brachytherapy. *Brachytherapy 11*, 2 (2012), 111–118.

[117] Kuban, D. A., Tucker, S. L., Dong, L., Starkschall, G., Huang, E. H., Cheung, M. R., Lee, A. K., and Pollack, A. Long-term results of the M. D. Anderson randomized dose-escalation trial for prostate cancer. *International Journal of Radiation Oncology Biology Physics 70*, 1 (2008), 67–74.

[118] Lahanas, M., and Baltas, D. Are dose calculations during dose optimization in brachytherapy necessary? *Medical Physics 30*, 9 (2003), 2368–2375.

[119] Lahanas, M., Baltas, D., and Giannouli, S. Global convergence analysis of fast multiobjective gradient-based dose optimization algorithms for high-dose-rate brachytherapy. *Physics in Medicine & Biology 48*, 5 (2003), 599–617.

[120] Lahanas, M., Baltas, D., and Zamboglou, N. A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy. *Physics in Medicine & Biology 48*, 3 (2003), 399–415.

[121] Larrañaga, P., and Lozano, J. A. *Estimation of distribution algorithms: A new tool for evolutionary computation*, vol. 2. Springer Science & Business Media, 2001.

[122] Lay, D. *Linear Algebra and its Applications*. Addison Wesley, New York, NY, USA, 1993.

[123] Ledoit, O., and Wolf, M. Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics 40*, 2 (2012), 1024–1060.

[124] Lessard, E., and Pouliot, J. Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function. *Medical Physics 28*, 5 (2001), 773–779.

[125] Li, X., Engelbrecht, A., and Epitropakis, M. G. Benchmark functions for CEC 2013 special session and competition on niching methods for multimodal function optimization. Tech. rep., Evolutionary Computation and Machine Learning Group, RMIT University, Australia, 2013.

[126] Li, X., Epitropakis, M. G., Deb, K., and Engelbrecht, A. Seeking Multiple Solutions: An Updated Survey on Niching Methods and Their Applications. *IEEE Transactions on Evolutionary Computation 21*, 4 (2017), 518–538.

[127] Liang, J. J., Yue, C. T., and Qu, B. Y. Multimodal multi-objective optimization: a preliminary study. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (2016), IEEE Press, pp. 2454–2461.

[128] Lones, M. A. Mitigating metaphors: A comprehensible guide to recent nature-inspired algorithms. *SN Computer Science 1*, 1 (2020), 49.

[129] Luke, S. *Essentials of Metaheuristics*. Lulu, 2013.

[130] Luong, N., Alderliesten, T., Pieters, B., Bel, A., Niatsetski, Y., and Bosman, P. A. N. Fast and insightful bi-objective optimization for prostate cancer treatment planning with high-dose-rate brachytherapy. *Applied Soft Computing 84* (2019), 105681.

[131] Luong, N. H., Alderliesten, T., Bel, A., Niatsetski, Y., and Bosman, P. A. N. Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment. *Swarm and Evolutionary Computation 40* (2018), 37–52.

[132] Luong, N. H., Alderliesten, T., and Bosman, P. A. N. Improving the performance of MO-RV-GOMEA on problems with many objectives using Tchebycheff scalarizations. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2018), ACM Press, pp. 705–712.

[133] Luong, N. H., Alderliesten, T., Pieters, B. R., Bel, A., Niatsetski, Y., and Bosman, P. A. N. Fast and insightful bi-objective HDR prostate brachytherapy planning. *Radiotherapy & Oncology 127*, Supplement 1 (2018), S130.

[134] Luong, N. H., and Bosman, P. A. N. Elitist archiving for multi-objective evolutionary algorithms: To adapt or not to adapt. In *Parallel Problem Solving from Nature - PPSN* (2012), Springer, pp. 72–81.

[135] Luong, N. H., Bouter, A., Van Der Meer, M. C., Niatsetski, Y., Witteveen, C., Bel, A., Alderliesten, T., and Bosman, P. A. N. Efficient, effective, and insightful tackling of the high-dose-rate brachytherapy treatment planning problem for prostate cancer using evolutionary multi-objective optimization algorithms. In *Proceedings of the Medical Applications of Genetic and Evolutionary Computation MedGEC Workshop at the Genetic and Evolutionary Computation Conference - GECCO* (2017), ACM Press, pp. 1372–1379.

[136] Maloof, M. Artificial intelligence: An introduction. Tech. rep., Department of Computer Science, Georgetown University, Washington, USA, 2017.

[137] Maree, S. C. Correcting non-positive definite correlation matrices. Tech. rep., TU Delft, 2012.

[138] Maree, S. C. *HillVallEA C++ source code on Github*, 2019. https://github.com/scmaree/HillVallEA.

[139] Maree, S. C. *HillVallEA C++ source code on Github*, 2019. https://github.com/scmaree/MOHillVallEA.

[140] Maree, S. C. *Uncrowded-hypervolume multi-objective optimization C++ source code on Github*, 2019. https://github.com/scmaree/uncrowded-hypervolume.

[141] Maree, S. C., Alderliesten, T., and Bosman, P. A. N. Benchmarking HillVallEA for the GECCO 2019 competition on multimodal optimization. *preprint arXiv:1907.10988* (2019).

[142] Maree, S. C., Alderliesten, T., and Bosman, P. A. N. Uncrowded hypervolume-based multi-objective optimization with gene-pool optimal mixing. *preprint arXiv:2004.05068* (2020).

[143] Maree, S. C., Alderliesten, T., Thierens, D., and Bosman, P. A. N. Niching an estimation-of-distribution algorithm by hierarchical Gaussian mixture learning. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2017), ACM Press, pp. 713–720.

[144] Maree, S. C., Alderliesten, T., Thierens, D., and Bosman, P. A. N. Real-valued evolutionary multi-modal optimization driven by hill-valley clustering. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2018), ACM Press, pp. 857–864.

[145] Maree, S. C., Bosman, P. A. N., Niatsetski, Y., Koedooder, C., van Wieringen, N., Bel, A., Pieters, B. R., and Alderliesten, T. Improved class solutions for prostate brachytherapy planning via evolutionary machine learning. *Radiotherapy & Oncology 127*, Supplement 1 (2017), S96–S97.

[146] Maree, S. C., Kooreman, E. S., Luong, N. H., van Wieringen, N., Bel, A., Rodenburg, E. C. M., Hinnen, K., Westerveld, G. H., Pieters, B. R., Bosman, P. A. N., and Alderliesten, T. Better plans and easy plan selection via bi-objective optimization for HDR prostate brachytherapy. *Radiotherapy & Oncology 127*, Supplement 1 (2018), S571–S572.

[147] Maree, S. C., Luong, N. H., van Wieringen, N., Bel, A., Hinnen, K., Westerveld, G. H., Pieters, B. R., Bosman, P. A. N., and Alderliesten, T. Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy – a retrospective observer study. *Brachytherapy 18*, 3 (2019), 396–403.

[148] Martinez, A. A., Kestin, L. L., Stromberg, J. S., Gonzalez, J. A., Wallace, M., Gustafson, G. S., Edmundson, G. K., Spencer, W., and Vicini, F. A. Interim report of image-guided conformal high-dose-rate brachytherapy for patients with unfavorable prostate cancer: The William Beaumont Phase II dose-escalating trial. *International Journal of Radiation Oncology Biology Physics 47*, 2 (2000), 343–352.

[149] Mayles, P., Nahum, A., and Rosenwald, J. *Handbook of radiotherapy physics: theory and practice*. Taylor & Francis, New York, 2007.

[150] McConnell, C. R., and Brue, S. L. *Economics: Principles, problems and policies*. McGraw-Hill/Irwin, 2005.

[151] McCorduck, P. *Machines Who Think*. A. K. Peters, Ltd., Natick, MA, USA, 2004.

[152] Mehta, V. K., and Dasgupta, B. Parametric Approximation of the Pareto Set in Multi-Objective Optimization Problems. *Journal of Multi-Critea Decision Analysis 21* (2014), 335–362.

[153] Mersmann, O., Bischl, B., Trautmann, H., Preuss, M., Weihs, C., and Rudolph, G. Exploratory landscape analysis. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2011), ACM Press, pp. 829–836.

[154] Miettinen, K. *Nonlinear Multiobjective Optimization*. International Series in Operations Research & Management Science. Springer US, 1999.

[155] Milickovic, N., Lahanas, M., Papagiannopoulo, M., Zamboglou, N., and Baltas, D. Multiobjective anatomy-based dose optimization for HDR-brachytherapy with constraint free deterministic algorithms. *Physics in Medicine & Biology 47*, 13 (2002), 2263–2280.

[156] Mockus, J. *Bayesian Approach to Global Optimization: Theory and Applications*. Mathematics and its Applications. Springer Netherlands, 2012.

[157] Mohler, J. L. The 20210 NCCN clinical practice guidelines in oncology onprostate cancer. *Journal of the National Comprehensive Cancer Network 8*, 2 (2010), 145–200.

[158] Morén, B., Larsson, T., and Carlsson Tedgren. Mathematical optimization of high dose-rate brachytherapy - derivation of a linear penalty model from a dose-volume model. *Physics in Medicine & Biology 63*, 6 (2018).

[159] Morton, G. C., Sankreacha, R., Halina, P., and Loblaw, A. A comparison of anatomy-based inverse planning with simulated annealing and graphical optimization for high-dose-rate prostate brachytherapy. *Brachytherapy 7*, 1 (2008), 12–16.

[160] Moshaiov, A. The paradox of multimodal optimization: Concepts vs. species in single and multi-objective problems. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (2016), IEEE Press, pp. 1743–1748.

[161] Mostaghim, S., Branke, J., and Schmeck, H. Multi-objective particle swarm optimization on computer grids. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation - GECCO* (2007), ACM Press, pp. 869–875.

[162] Müller, B. S., Shih, H. A., Efstathiou, J. A., Bortfeld, T., and Craft, D. Multicriteria plan optimization in the hands of physicians: a pilot study in prostate cancer and brain tumors. *Radiation Oncology 12*, 1 (2017), 168.

[163] Nelder, J. A., and Mead, R. A simplex method for function minimization. *Computer Journal 7*, 4 (1965), 308–313.

[164] Nicolini, M. A two-level evolutionary approach to multi-criterion optimization of water supply systems. In *Conference on Evolutionary Multi-Criterion Optimization - EMO* (2005), vol. 3410 of LNCS, Springer, pp. 736–751.

[165] Niu, P., Niu, S., and Chang, L. The defect of the grey wolf optimization algorithm and its verification method. *Knowledge-Based Systems 171* (2019), 37–43.

[166] Pelikan, M., Goldberg, D. E., and Cantú-Paz, E. Linkage problem, distribution estimation, and Bayesian networks. *Evolutionary Computation 8*, 3 (2000), 311–340.

[167] Pelikan, M., Sastry, K., and Cantú-Paz, E., Eds. *Scalable Optimization via Probabilistic Modeling*, 1st ed. Springer-Verlag Berlin Heidelberg, 2006.

[168] Peschel, R. E., Colberg, J. W., Chen, Z., Nath, R., and Wilson, L. D. Iodine 125 versus Palladium 103 implants for prostate cancer. *The Cancer Journal 10*, 3 (2004), 170–174.

[169] Pieters, B. R., van der Grient, J. N. B., Blank, L. E. C. M., Koedooder, C., Hulshof, M. C. C. M., and de Reijke, T. M. Minimal displacement of novel self-anchoring catheters suitable for temporary prostate implants. *Radiotherapy & Oncology 80*, 1 (2006), 69–72.

[170] Pincus, M. A Monte Carlo Method for the Approximate Solution of Certain Types of Constrained Optimization Problems. *Operations Research 18*, 6 (1970), 1225–1228.

[171] Pirpinia, K., Alderliesten, T., Sonke, J. J., van Herk, M., and Bosman, P. A. N. Diversifying multi-objective gradient techniques and their role in hybrid multi-objective evolutionary algorithms for deformable medical image registration. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation - GECCO* (2015), ACM Press, pp. 1255–1262.

[172] Porter, A. T., and Forman, J. D. Prostate brachytherapy. An overview. *Cancer 71*, 3 S (1993), 953–958.

[173] Prada, P. J., Cardenal, J., Blanco, A. G., Anchuelo, J., Ferri, M., Fernández, G., Arrojo, E., Vázquez, A., Pacheco, M., and Fernández, J. High-dose-rate interstitial brachytherapy as monotherapy in one fraction for the treatment of favorable stage prostate cancer: Toxicity and long-term biochemical resultsProstate HDR brachytherapy as monotherapy in one fraction. *Radiotherapy & Oncology 119*, 3 (2016), 411–416.

[174] Preuss, M. Niching the CMA-ES via nearest-better clustering. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2010), ACM Press, pp. 1711–1718.

[175] Preuss, M. Improved topological niching for real-valued global optimization. *In Applications of Evolutionary Computation* (2012), 386–395.

[176] Preuss, M. *Multimodal Optimization by Means of Evolutionary Algorithms*, 1st ed. Springer Publishing Company, Incorporated, 2015.

[177] Preuss, M., Naujoks, B., and Rudolph, G. Pareto set and EMOA behavior for simple multimodal multiobjective functions. In *Parallel Problem Solving from Nature - PPSN* (2006), Springer, pp. 513–522.

[178] Rastrigin, L. A. The convergence of the random search method in the extremal control of a many parameter system. *Automation and Remote Control 24*, 10 (1963), 1337—1342.

[179] Reynaerd, N., van der Marck, S., Schaart, D., van der Zee, W., Tomesj, M., van Vliet-Vroegindeweij, C., Jansen, J., Coghe, M., de Wagter, C., and Heijmen, B. Monte carlo treatment planning: An introduction. Tech. rep., Netherlands Commission on Radiation Dosimetry, 2006.

[180] Rivard, M. J., Coursey, B. M., DeWerd, L. A., Hanson, W. F., Saiful Huq, M., Ibbott, G. S., Mitch, M. G., Nath, R., and Williamson, J. F. Update of AAPM Task Group No. 43 Report: a revised AAPM protocol for brachytherapy dose calculations. *Medical Physics 31*, 3 (2004), 633–674.

[181] Rodrigues, S., Bauer, P., and Bosman, P. A. N. A novel population-based multi-objective CMA-ES and the impact of different constraint handling techniques. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2014), ACM Press, pp. 991–998.

[182] Rothlauf, F. *Design of Modern Heuristics*. Springer-Verlag Berlin Heidelberg, 2011.

[183] RTOG 0924. RTOG 0924 androgen deprivation therapy and high dose radiotherapy with or without whole-pelvic radiotherapy in unfavorable intermediate or favorable high risk prostate cancer: A phase III randomized trial. Tech. rep., Radiation Therapy Oncology Group, 2016.

[184] Rudolph, G., Naujoks, B., and Preuss, M. Capabilities of emoa to detect and preserve equivalent Pareto subsets. In *Evolutionary Multi-criterion Optimization - EMO* (2007), Springer, pp. 36–50.

[185] Ruotsalainen, H., Miettinen, K., Palmgren, J.-E., and Lahtinen, T. Interactive multiobjective optimization for anatomy-based three-dimensional HDR brachytherapy. *Physics in Medicine & Biology 55*, 16 (2010), 4703–4719.

[186] Russell, S. J., and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, New Jersey, USA, 2009.

[187] Schartz, B. *The paradox of choice*. Harper Perennial, 2004.

[188] Schütze, O., Martín, A., Lara, A., Alvarado, S., Salinas, E., and Coello Coello, C. A. The directed search method for multi-objective memetic algorithms. *Computational Optimization and Applications 63*, 2 (2016), 305–332.

[189] Siauw, T., Cunha, A., Atamtürk, A., Hsu, I.-C., Pouliot, J., and Goldberg, K. IPIP: A new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices. *Medical Physics 38*, 7 (2011), 4045–4051.

[190] Snyder, K. M., Stock, R. G., Hong, S. M., Lo, Y. C., and Stone, N. N. Defining the risk of developing Grade 2 proctitis following 125I prostate brachytherapy using a rectal dose-volume histogram analysis. *International Journal of Radiation Oncology Biology Physics 50*, 2 (2001), 335–341.

[191] Spapen, S. J., Damhuis, R. A., and Kirkels, W. J. Trends in the curative treatment of localized prostate cancer after introduction of prostate-specific antigen: data from the Rotterdam Cancer Registry. *BJU International 85* (2000), 474–480.

[192] Tanabe, R., and Ishibuchi, H. A decomposition-based evolutionary algorithm for multi-modal multi-objective optimization. In *Parallel Problem Solving from Nature - PPSN* (2018), vol. 5199, Springer, pp. 249–261.

[193] Thieke, C., Küfer, K. H., Monz, M., Scherrer, A., Alonso, F., Oelfke, U., Huber, P. E., Debus, J., and Bortfeld, T. A new concept for interactive radiotherapy planning with multicriteria optimization: First clinical evaluation. *Radiotherapy & Oncology 85*, 2 (2007), 292–298.

[194] Thierens, D., and Bosman, P. A. N. Hierarchical Problem Solving with the Linkage Tree Genetic Algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2013), ACM Press, pp. 877–884.

[195] Touré, C., Hansen, N., Auger, A., and Brockhoff, D. Uncrowded hypervolume improvement: COMO-CMA-ES and the Sofomore framework. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2019), ACM Press, pp. 638–646.

[196] Unkelbach, J., Alber, M., Bangert, M., Bokrantz, R., Chan, T. C., Deasy, J. O., Fredriksson, A., Gorissen, B. L., Van Herk, M., Liu, W., Mahmoudzadeh, H., Nohadani, O., Siebers, J. V., Witte, M., and Xu, H. Robust radiotherapy planning. *Physics in Medicine & Biology 63*, 22 (2018), 22TR02.

[197] Ursem, R. K. Multinational evolutionary algorithms. In *Proceedings of the IEEE Congress on Evolutionary Computation - CEC* (1999), IEEE Press, pp. 1633–1640.

[198] van de Velde, C. J. H., van der Graaf, W. T. A., van Krieken, J. H. J. M., and Marijnen, C. A. M., Eds. *Leerboek Oncologie*. Bohn Stafleu van Loghum, 2017.

[199] van der Meer, M. C., Bel, A., Niatsetski, Y., Alderliesten, T., Pieters, B. R., and Bosman, P. A. N. Robust evolutionary bi-objective optimization for prostate cancer treatment with high-dose-rate brachytherapy. In *Parallel Problem Solving from Nature - PPSN* (2020), T. e. a. Bäck, Ed., Springer, pp. 441–453.

[200] van der Meer, M. C., Bosman, P. A. N., Niatsetski, Y., Alderliesten, T., van Wieringen, N., Pieters, B. R., and Bel, A. Bi-objective optimization of catheter positions for high-dose-rate prostate brachytherapy. *Medical Physics 47*, 12 (2020), 6077–6086.

[201] Van der Meer, M. C., Bosman, P. A. N., Pieters, B. R., Niatsetski, Y., Alderliesten, T., and Bel, A. Sensitivity of dose-volume indices to organ reconstruction settings in HDR prostate brachytherapy. *Radiotherapy & Oncology 127*, Supplement 1 (2018), S1236.

[202] van der Meer, M. C., Bosman, P. A. N., Pieters, B. R., Niatsetski, Y., van Wieringen, N., Alderliesten, T., and Bel, A. Sensitivity of dose-volume indices to computation settings in high-dose-rate prostate brachytherapy treatment plan evaluation. *Journal of Applied Clinical Medical Physics 20*, 4 (2019), 66–74.

[203] van der Meer, M. C., Pieters, B. R., Niatsetski, Y., Alderliesten, T., Bel, A., and Bosman, P. A. N. Better and faster catheter position optimization in HDR brachytherapy for prostate cancer using multi-objective real-valued GOMEA. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO* (2018), ACM Press, pp. 1387–1394.

[204] Wang, H., Deutz, A., Bäck, T., and Emmerich, M. Hypervolume indicator gradient ascent multi-objective optimization. In *Evolutionary Multi-Criterion Optimization - EMO* (2017), Springer International Publishing, pp. 654–669.

[205] Wessing, S. Two-stage methods for multimodal optimization. *PhD Thesis, Technische Universität Dortmund* (2015).

[206] Wierstra, D., Glasmachers, T., Sun, Y., Peters, J., and Schmidhuber, J. Natural evolution strategies. *Journal of Machine Learning Research 15* (2014), 949–980.

[207] Wilkinson, J. M. Geometric uncertainties in radiotherapy. *British journal of radiology 77*, 914 (2004), 86–87.

[208] Wolpert, D. H., and Macready, W. G. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation 1*, 1 (1997), 67–82.

[209] Wu, B., Pang, D., Simari, P., Taylor, R., Sanguineti, G., and McNutt, T. Using overlap volume histogram and IMRT plan data to guide and automate VMAT planning: a head-and-neck case study. *Medical Physics 40*, 2 (2013), 021714.

[210] Wu, B., Ricchetti, F., Sanguineti, G., Kazhdan, M., Simari, P., Chuang, M., Taylor, R., Jacques, R., and McNutt, T. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physics 36*, 12 (2009), 5497–5505.

[211] Yue, C., Qu, B., and Liang, J. A multi-objective particle swarm optimizer using ring topology for solving multimodal multi-objective problems. *IEEE Transactions on Evolutionary Computation 22*, 5 (2018), 805–817.

[212] Zhang, Q., and Li, H. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. *IEEE Transactions on Evolutionary Computation 11*, 6 (2007), 712–731.

[213] Zhou, A., Zhang, Q., and Jin, Y. Approximating the set of Pareto-optimal solutions in both the decision and objective spaces by an estimation of distribution algorithm. *IEEE Transactions on Evolutionary Computation 13*, 5 (2009), 1167–1189.

[214] Zhou, J., Zamdborg, L., and Sebastian, E. Review of advanced catheter technologies in radiation oncology brachytherapy procedures. *Cancer management and research 7* (2015), 199 –211.

[215] Zietman, A. L., DeSilvio, M. L., Slater, J. D., Rossi, C. J., Miller, D. W., Adams, J. A., and Shipley, W. U. Comparison of conventional-dose vs high-dose conformal radiation therapy in clinically localized adenocarcinoma of the prostate: A randomized controlled trial. *Journal of the American Medical Association 294*, 10 (2005), 1233–1239.

[216] Zitzler, E., Brockhoff, D., and Thiele, L. The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In *Conference on Evolutionary Multi-Criterion Optimization - EMO* (2007), vol. 4403 of LNCS, Springer, pp. 862–876.

[217] Zitzler, E., and Künzli, S. Indicator-based selection in multiobjective search. In *International conference on Parallel Problem Solving from Nature - PPSN* (2004), Springer, pp. 832–842.

[218] Zitzler, E., Laumanns, M., and Thiele, L. SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization. In *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems - EUROGEN 2001* (2001), International Center for Numerical Methods in Engineering (CIMNE), pp. 95–100.

[219] Zitzler, E., and Thiele, L. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation 3*, 4 (1999), 257–271.

[220] Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C. M., and Da Fonseca, V. G. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation 7*, 2 (2003), 117–132.

List of acronyms

AMaLGaM	Adapted Maximum-Likelihood Gaussian Model
BezEA	Bézier-based GOMEA
CMA-ES	Covariance Matrix Adaptation Evolution Strategies
CMSA	Covariance Matrix Self-Adaptation evolution strategy
CPU	Central Processing Unit
DFC	Density-Fitness Correlation
DTDC	Dwell Time Deviation Constraint
DTGR	Dwell Time Gradient Restriction
DVI	Dose-Volume Index
EBRT	External Beam Radiotherapy
ESTRO	European SocieTy for Radiotherapy & Oncology
GD	Generational Distance
GEC	Groupe Européen de Curiethérapie
gHSS	Greedy Hypervolume Subset Selection
GMM	Gaussian Mixture Model
GOM	Gene-pool Optimal Mixing
GOMEA	Gene-pool Optimal Mixing Evolutionary Algorithm
GPU	Graphics Processing Unit
HDR	High Dose Rate
HGML	Hierarchical Gaussian Mixture Learning
HillVallEA	Hill-Valley Evolutionary Algorithm
HIPO	Hybrid Inverse Plan Optimization
HVC	Hill-Valley Clustering
iAMaLGaM	Incremental AMaLGaM
IBMOP	Indicator-Based Multi-objective Optimization Problem
IGD	Inverted Generational Distance
IGDX	Inverted Generational Distance measured in decision space
IPSA	Inverse Planning Simulated Annealing
L-BFGS	Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm
LCI	Least Coverage Index
LSI	Least Sparing Index

MAMaLGaM	Multi-Objective AMaLGaM
MMOEA	Multimodal Multi-Objective Evolutionary Algorithm
MO-GOMEA	Multi-Objective GOMEA
MOEA	Multi-Objective Evolutionary Algorithm
MR	Mode Ratio
MRI	Magnetic Resonance Imaging
NBC	Nearest-Better Clustering
NEA2+	Nearest-better Evolutionary Algorithm
NES	Natural Evolution Strategies
NR	Niche Ratio
NSGA-II	Non-dominated Sorting Genetic Algorithm
NTCP	Normal Tissue Complication Probability
OAR	Organ At Risk
RS-CMSA	Repelling-Subpopulations CMSA
RTT	Radiation Therapy Technologist
SMS-EMOA	$\mathscr{S}$-Metric Selection Evolutionary Multi-objective Optimization Algorithm
SR	Success Ratio
TCP	Tumor Control Probability
UHV	Uncrowded Hypervolume
UHV-GOMEA	Uncrowded-Hypervolume-based GOMEA
UHVEA	Uncrowded-Hypervolume-based GOMEA
UHVI	Uncrowded Hypervolume Improvement

Samenvatting

(English summary on page 11)

Modelgebaseerde evolutionaire algoritmen voor het vinden van diverse hoogwaardige oplossingen

met een toepassing in brachytherapie voor prostaatkanker

Onze samenleving is meer en meer afhankelijk van data en algoritmen voor het oplossen van belangrijke optimalisatieproblemen met een grote impact. Deze problemen zijn vaak zeer complex, waardoor het moeilijk is om voldoende inzicht te krijgen in hun interne structuur om een probleemspecifieke optimalisatie-aanpak te kunnen ontwikkelen, of om exacte optimalisatie-algoritmen toe te kunnen passen. Om deze problemen toch te kunnen oplossen kan een *metaheuristiek* worden gebruikt. Metaheuristieken zijn zoekalgoritmen die kunnen worden toegepast zonder dat er rekening hoeft te worden gehouden met probleemspecifieke kennis. Het zijn echter, zoals de naam al aangeeft, heuristieken, wat betekent dat er doorgaans geen garanties zijn voor het vinden van de optimale oplossing. Om efficiënt te kunnen zoeken hebben metaheuristieken een bepaalde zoektendens, wat betekent dat ze proberen bepaalde probleemkenmerken uit te buiten. Wanneer het probleem in kwestie geen kenmerken vertoont die kunnen worden uitgebuit kan de zoektocht erg inefficiënt zijn of alleen oplossingen van lage kwaliteit opleveren. *Modelgebaseerde metaheuristieken* verkleinen dit risico door probleemkenmerken uit te buiten via leerbare modellen. Het gebruik van dit soort modellen maakt het expliciet welke probleemkenmerken kunnen worden uitgebuit. Hierdoor kunnen, met een bepaalde zekerheid, hoogwaardige oplossingen worden verkregen voor de klasse van problemen die goed aansluiten bij het gebruikte model.

Metaheuristieken zijn een waardevol hulpmiddel voor optimalisatie in de praktijk, waarin het vaak belangrijker is dat er snel een hoogwaardige oplossing wordt verkregen, dan dat de verkregen oplossing aantoonbaar optimaal is. In de praktijk kan het echter moeilijk zijn om alle wenselijke aspecten van een oplossing te omvatten in een kwantitatieve doelfunctie. Dit geldt met name wanneer de gewenste

oplossing een afweging is tussen tegenstrijdige aspecten, zoals de afweging tussen prijs en kwaliteit in de economie; de afweging tussen bias en variantie bij machine learning; of de afweging tussen een effectieve behandeling en bijwerkingen in de geneeskunde. Het is voor dit soort problemen namelijk op voorhand vaak onbekend hoe deze afwegingen zich zouden moeten verhouden in een gewenste oplossing. Het is daardoor niet ongebruikelijk dat de gevonden oplossing niet naar wens is. Om dit te voorkomen richten we ons in dit proefschrift op het ontwerp, de ontwikkeling en de toepassing van metaheuristieken die niet één oplossing vinden, maar een verzameling van *diverse* hoogwaardige oplossingen. Door op zoek te gaan naar diverse hoogwaardige oplossingen, en deze achteraf te vergelijken, kunnen de afwegingen van een probleem die normaal impliciet zijn, expliciet worden gemaakt. Dit geeft inzicht dat besluitvormers kan helpen om de meest wenselijke oplossing voor hun probleem te vinden. *Evolutionaire algoritmen* zijn metaheuristieken die van nature zeer geschikt zijn voor het vinden van meerdere hoogwaardige oplossingen, aangezien deze metaheuristieken tijdens het zoekproces al gebruik maken van een populatie van meerdere oplossingen. Wat dan nog rest is om tijdens optimalisatie op een bepaalde manier de diversiteit aan oplossingen binnen deze populatie te bevorderen. Dit doen we op twee verschillende manieren: via *meerdoelige optimalisatie* en via *multimodale optimalisatie*.

In dit proefschrift passen we modelgebaseerde evolutionaire algoritmen toe op een optimalisatieprobleem dat naar voren komt bij de behandeling van prostaatkanker met snelle (high-dose-rate) *brachytherapie*. Brachytherapie is een vorm van inwendige bestraling, waarbij een hoge stralingsdosis wordt gegeven om tumorcellen te doden. Daarbij kan stralingsdosis op nabijgelegen gezond weefsel echter niet worden vermeden. Tijdens de planning van de behandeling wordt bepaald hoe de stralingsdosis moet worden afgegeven, zodat de tumor zoveel mogelijk wordt bestraald, terwijl het omliggende weefsel zoveel mogelijk wordt ontzien. Deze tegenstrijdige doelen zijn de reden dat het plannen van de behandeling inherent een meerdoelig optimalisatieprobleem is, waarvoor in praktijk een aantal tegenstrijdige planningscriteria zijn geformuleerd.

Klinisch beschikbare methoden voor het plannen van brachytherapie-behandelingen vereenvoudigen de planningscriteria en gebruiken een gewogen-som benadering om deze te combineren tot een enkeldoelig optimalisatieprobleem. Deze methoden moeten echter handmatig worden ingesteld om een behandelplan te verkrijgen dat gewenste afwegingen maakt tussen de planningscriteria. Wij onderzoeken of een recent geïntroduceerd tweedoelig planningsmodel kan worden gebruikt om het planningsproces te versnellen en intuïtiever te maken. Dit model, dat direct is gebaseerd op de klinische planningscriteria, kan efficiënt worden geoptimaliseerd

met het Multi-Objective Gene-pool Optimal Mixing Evolutionary Algorithm (MO-GOMEA). Op deze manier kan een verzameling van diverse behandelplannen worden verkregen, waarbij elk plan een andere afweging heeft tussen de dosis in de tumor en de dosis in nabijgelegen gezond weefsel. Een arts kan dan het meest wenselijke plan uit deze verzameling selecteren. Om te helpen bij de selectie kan de verzameling plannen worden gevisualiseerd als een inzichtelijke afwegingscurve, waarin direct kan worden afgelezen welke behandelplannen voldoen aan alle klinische planningscriteria.

In hoofdstuk 2 laten we zien dat het tweedoelige planningsmodel eenvoudig kan worden geconfigureerd op basis van de planningscriteria die in onze kliniek worden gebruikt en dat de optimalisatie van dit model onze klinische praktijk goed weerspiegelt. Onderzoek onder ervaren artsen toont aan dat plannen die zijn verkregen via tweedoelig plannen als klinisch aanvaardbaar kunnen worden beschouwd. In 98% van de gevallen krijgen deze plannen retrospectief zelfs de voorkeur boven het klinisch gebruikte plan. De artsen waarderen bovendien de mogelijkheid om hoogwaardige plannen met diverse afwegingen te kunnen vergelijken en vinden de resultaten inzichtelijk.

In hoofdstuk 3 gebruiken we een meerdoelig evolutionair algoritme voor het automatisch instellen van de parameters van twee methoden die klinisch beschikbaar zijn voor het plannen van brachytherapie-behandelingen. Het instellen is gericht op het maximaliseren van de doelfuncties van het tweedoelige planningsmodel. We laten zien dat daarmee behandelplannen met goede doelfunctie-waarden kunnen worden verkregen. Directe optimalisatie van het tweedoelige planningsmodel met MO-GOMEA resulteert echter in behandelplannen met betere doelfunctie-waarden. Het automatisch instellen van parameters per patiënt is bovendien te tijdrovend voor de klinische praktijk. We hebben daarom via optimalisatie standaardinstellingen (*klasse-oplossingen*) bepaald die kunnen worden gebruikt voor alle patiënten als startpunt voor verdere handmatige instellingen. Deze standaardinstellingen kunnen handwerk echter niet geheel vervangen.

In hoofdstuk 4 ontwikkelen we een algoritme voor efficiënte en effectieve multimodale optimalisatie. Het doel van multimodale optimalisatie is om alle optimale oplossingen van een probleem te verkrijgen. Dit wordt meestal gedaan door meerdere *modi* te doorzoeken, dat wil zeggen, regio's van hoge kwaliteit in de zoekruimte. Hier komt de naam 'multimodale optimalisatie' dan ook vandaan. Voor de brede toepasbaarheid van deze algoritmen is het essentieel dat er geen aannames worden gedaan over het aantal modi, of over hun vorm of grootte. We introduceren *hill-valley clustering*, berg-dal clustering in het Nederlands, dat kan worden gebruikt om modi te identificeren. Hill-valley clustering maakt gebruik van

de simpele hill-valley test om te bepalen of twee oplossingen tot dezelfde berg (modus) behoren door te controleren of er een dal tussen zit. Het hill-valley evolutionaire algoritme (HillVallEA) combineert hill-valley clustering met een evolutionair algoritme. HillVallEA is vooroplopend in multimodale optimalisatie en is tweevoudig winnaar van de jaarlijkse multimodale optimalisatiewedstrijd op de Genetic and Evolutionary Computation Conference.

In hoofdstuk 5 bespreken we hoe HillVallEA kan worden uitgebreid om ook optimalisatieproblemen met meerdere doelen op te kunnen lossen. Het doel van multimodale meerdoelige optimalisatie om alle globaal optimale oplossingen te verkrijgen, ofwel, alle oplossingen in de *Pareto verzameling*. Dit voegt een complexiteitslaag toe aan meerdoelige evolutionaire algoritmen voor multimodale optimalisatie, aangezien er een balans moet worden gevonden tussen diversiteit binnen een oplossingsverzameling, met oplossingen van verschillende afwegingen, en tussen oplossingsverzamelingen in verschillende modi. We tonen aan dat de meerdoelige HillVallEA beter presteert dan andere meerdoelige optimalisatie-algoritmen in multimodale optimalisatie op een aantal bekende multimodale testproblemen. Bovendien, en misschien wel het belangrijkste, is dat we laten zien dat het in staat is om oplossingsverzamelingen in meerdere modi tegelijk te verkrijgen, te onderhouden en te verbeteren.

De meest bekende evolutionaire algoritmen voor meerdoelige optimalisatie zijn gebaseerd op Pareto-dominantie, waarbij de kwaliteit van een oplossing wordt bepaald door het aantal oplossingen in de populatie dat deze domineert. Een fundamentele beperking van deze benadering is echter dat er stagnatie optreedt wanneer er geen gedomineerde oplossingen in de populatie meer zijn. Indicator-gebaseerde meerdoelige optimalisatie kan dit verhelpen door het meerdoelige probleem te formuleren als een hoog-dimensionaal enkeldoelig probleem.

In hoofdstuk 6 bereiken we dit door de hypervolume-indicator aan te passen, zodat deze ook met gedomineerde oplossingen rekening kan houden, wat resulteert in de *uncrowded hypervolume* (UHV) indicator. Om het hoog-dimensionale enkeldoelige probleem efficiënt op te kunnen lossen, buiten we de structuur van de UHV uit met de enkeldoelige GOMEA. We laten zien dat het resulterende algoritme, dat we UHV-GOMEA noemen, kan convergeren naar een deelverzameling van de Pareto-verzameling. UHV-GOMEA behoudt daarbij een populatie van oplossingen, waardoor het beter kan presteren op multimodale problemen.

Met behulp van meerdoelige optimalisatie wordt een verzameling van diverse hoogwaardige oplossingen verkregen, waaruit een besluitvormer dan de meest wenselijke oplossing dient te selecteren. Deze selectie kan gedaan door oplossingen langs te gaan in de volgorde die overeenkomt met het doorlopen van de bijbehorende

afwegingscurve. Deze volgorde leidt echter niet noodzakelijkerwijs tot een glad en gelijkmatig verlopend traject in de beslissingsruimte. Dit dwingt de besluitvormer om de beslissingsvariabelen van alle oplossingen afzonderlijk te inspecteren, wat het selectieproces tijdrovend en onintuïtief kan maken.

In hoofdstuk 7 richten we optimalisatie expliciet op het verkrijgen van oplossingsverzamelingen die eenvoudig te doorlopen zijn, door gebruik te maken van de UHV-probleemformulering. Hiervoor formuleren we oplossingsverzamelingen als gladde Bézier-krommen in de beslissingsruimte. We lossen de Bézier-probleemformulering op met GOMEA, en noemen het resulterende algoritme BezEA. BezEA is competitief, terwijl resulterende oplossingsverzamelingen eenvoudig te doorlopen zijn doordat ze een glad traject in de beslissingsruimte vormen. We laten verder zien dat BezEA het tweedoelige brachytherapie planningsmodel efficiënt kan oplossen door gebruik te maken van het feit dat plannen op een Bézier-kromme efficiënt kunnen worden geëvalueerd. De verkregen verschillen in doelfuncie-waarden zijn klein, wat suggereert dat BezEA een goed alternatief is voor MO-GOMEA wanneer het wenselijk is om het selectieproces intuïtiever en minder tijdrovend te maken.

Concluderend is de bijdrage van dit proefschrift tweeledig. Ten eerste valideren we het tweedoelig plannen van brachytherapie-behandelingen voor prostaatkanker. Ten tweede dragen we bij aan de algemene kennis over algoritmen voor het vinden van diverse hoogwaardige oplossingen. We ontwikkelen de eenvoudige maar effectieve hill-valley clustering methode, die gebruikt kan worden voor efficiënte en effectieve multimodale optimalisatie. We verbreden daarnaast het veld van meerdoelige optimalisatie door aan te tonen hoe met hypervolume-gebaseerde meerdoelige optimalisatie efficiïente convergentie naar het optimum kan worden bereikt. Ten slotte laten we zien hoe deze aanpak kan worden gebruikt om oplossingsverzamelingen te verkrijgen die eenvoudig te navigeren zijn en daarmee de besluitvormer verder helpen bij het verkrijgen van de meest wenselijke oplossing voor het probleem in kwestie.

PhD portfolio

Name	S.C. Maree
PhD period	February 2016 – January 2020
Department	Radiation Oncology
Promotores	Prof. dr. P.A.N. Bosman
	Prof. dr. C.R.N. Rasch
Copromotores	Dr. T. Alderliesten
	Dr. A. Bel

General courses	*year*	*credits*
The AMC World of Science	2016	0.7
Oral Presentation in English	2016	0.8
Practical Biostatistics	2016	1.1

Specific courses	*year*	*credits*
SIKS course on Data Science, Vught	2016	0.5
Design and Application of Modern Heuristics, Mainz, Germany	2016	6.0
Annual OOA PhD student retreat, Renesse	2016	0.75
ESTRO 36, Vienna, Austria	2017	0.75
GECCO'17 Workshop and tutorials, Berlin, Germany	2017	1.5
CWI Machine Learning Lectures, Amsterdam	2017	0.25
Radiotherapie Kringdag – 18 november, Utrecht	2017	0.5
CCA Retreat, Noordwijkerhout	2018	0.25
ESTRO 37, Barcelona, Spain	2018	0.75
GECCO'18 Workshop and tutorials, Kyoto, Japan	2018	1.5
World of AI, Amsterdam	2018	0.25
IXA Breakfast session on valorization, Amsterdam	2019	0.1
KNAW Deep learning in the brain	2019	0.1
GECCO'19 Workshop and tutorials, Prague, Czech Republic	2019	1.5
MACODA Workshop, Leiden	2019	2.5
PPSN'20 Workshop and tutorials, Leiden,	2020	1.5

Oral Presentations	*year*	*credits*
Radiotherapie Fysica Presentatie, AMC, Amsterdam	2016	0.5
CWI Life Sciences and Health Seminar, Amsterdam	2017	0.5
GECCO 2017, Berlin, Germany	2017	0.5
NVKF, Brachytherapy project day, Amsterdam	2017	0.5
Afdelingspresentatie Radiotherapie, AMC, Amsterdam	2018	0.5
CCA NeXt (project pitch), AMC, Amsterdam	2018	0.5
CCA Retreat, Noordwijkerhout	2018	0.5
CWI Life Sciences and Health Seminar, Amsterdam	2018	0.5
GECCO 2018, Kyoto, Japan	2018	0.5
Session chair, CCA NeXt, VUmc, Amsterdam	2018	0.5
Elekta inhouse presentation, Veenendaal	2018	0.5
Radiotherapie Fysica Presentatie, AMC, Amsterdam	2019	0.5
Wetenschappelijke kringdag, UMCU, Utrecht	2019	0.5
CWI Life Sciences and Health Seminar, Amsterdam	2019	0.5
GECCO 2019, Prague, Czech Republic	2019	0.5
HUMIES competition pitch, Prague, Czech Republic	2019	0.5
MACODA workshop, Leiden	2019	0.5
PPSN 2020, MMMOO workshop, Leiden	2020	0.5

Poster Presentations	*year*	*credits*
Annual OOA PhD student retreat, Renesse	2016	0.5
ESTRO 36, Vienna, Austria	2017	0.5
NVKF, Woudschoten conferentie, Woudschoten	2017	0.5
ESTRO 37, Barcelona, Spain	2018	0.5
PPSN 2020, Leiden	2020	0.5

Teaching and supervising	*year*	*credits*
E.S. Kooreman, MSc Student (6 months)	2017	2.0
TU Delft Computer Science BSc Project, shared supervision	2017	0.5
M.A.R.C.M. Verzijl, TU Delft BSc Internship, shared supervision	2018	0.5
Teaching assistant TU Delft MSc Course Evolutionary Algorithms	2020	1.0

Parameters of esteem	*year*
Best poster award, annual OOA PhD retreat, Renesse	2016
Best Junior Brachytherapy Presentation, ESTRO 36, Vienna, Austria	2017
Junior Brachytherapy travel grant, ESTRO 37, Barcelona, Spain	2018
Niching competition, 1st place, GECCO 2018, Kyoto, Japan	2018
Niching competition, 1st place, GECCO 2019, Prague, Czech Republic	2019
HUMIES, silver award, project team, GECCO 2019, Prague, Czech Republic	2019
Best paper nomination, co-author, GECCO 2020, Cancun, Mexico	2020

List of publications

S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2020) *Ensuring Smoothly Navigable Approximation Sets by Bézier Curve Parameterizations in Evolutionary Bi-objective Optimization.* In Proceedings of the Parallel Problem Solving from Nature Conference, 215–228, Springer, Cham.

T.M. Deist, **S.C. Maree** (shared first author), T. Alderliesten, P.A.N. Bosman. (2020) *Multi-objective Optimization by Uncrowded-Hypervolume Gradient Ascent.* In Proceedings of the Parallel Problem Solving from Nature Conference, 186–200, Springer, Cham.

S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2020) *Uncrowded Hypervolume-based Multi-objective Optimization with Gene-pool Optimal Mixing.* (submitted for journal publication, preprint arXiv:2004.05068).

S.C. Maree, D. Thierens, T. Alderliesten, P.A.N. Bosman. (2020) *Two-phase Real-valued Multimodal Optimization with the Hill-Valley Evolutionary Algorithm.* In M. Preuss, M.G. Epitropakis, J.E. Fieldsend and X. Li (Eds.), *Metaheuristics for Finding Multiple Solutions*, Springer (in press).

A. Bouter, **S.C. Maree**, T. Alderliesten, P.A.N. Bosman. (2020) *Leveraging Conditional Linkage Models in Gray-box Optimization with the Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm.* In Proceedings of the Genetic and Evolutionary Computation Conference, 603–611, ACM press, New York, New York.

S.C. Maree, P.A.N. Bosman, N. van Wieringen, Y. Niatsetski, B.R. Pieters, A. Bel, T. Alderliesten. (2020) *Automatic bi-objective parameter tuning for inverse planning of high-dose-rate prostate brachytherapy.* Physics in Medicine & Biology 65 (7), 075009.

S.C. Maree, T. Alderliesten, P.A.N. Bosman. (2019) *Real-valued evolutionary multi-modal multi-objective optimization by hill-valley clustering.* In Proceedings of the Genetic and Evolutionary Computation Conference, 568–576, ACM Press, New York, New York.

S.C. Maree, N.H. Luong, E.S. Kooreman, N. van Wieringen, A. Bel, K.A. Hinnen, G.H. Westerveld, B.R. Pieters, P.A.N. Bosman, T. Alderliesten. (2019) *Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy – a retrospective observer study.* Brachytherapy 18 (3), 396–403.

S.C. Maree, T. Alderliesten, D. Thierens, P.A.N. Bosman. (2018) *Real-valued evolutionary multi-modal optimization driven by Hill-Valley Clustering.* In Proceedings of the Genetic and Evolutionary Computation Conference, 857–864, ACM Press, New York, New York.

S.C. Maree, E.S. Kooreman, N.H. Luong, N. van Wieringen, A. Bel, E.C.M. Rodenburg, K.A. Hinnen, G.H. Westerveld, B.R. Pieters, P.A.N. Bosman, T. Alderliesten. (2018) *Better plans and easy plan selection via bi-objective optimization for HDR prostate brachytherapy.* Radiotherapy and Oncology 127, Supplement 1, S571–S572.

S.C. Maree, L. Ortiz-Gracia, C.W. Oosterlee. (2018) *Fourier and Wavelet Option Pricing Methods.* In M.A.H. Dempster, J. Kanniainen, J. Keane, E. Vynckier (Eds), *High-Performance Computing in Finance: Problems, Methods, and Solutions*, 249–272, Chapman and Hall/CRC.

S.C. Maree, T. Alderliesten, D. Thierens, and P.A.N. Bosman. (2017) *Niching an estimation-of-distribution algorithm by hierarchical Gaussian mixture learning.* In Proceedings of the Genetic and Evolutionary Computation Conference, 713–720, ACM Press, New York, New York.

S.C. Maree, P.A.N. Bosman, Y. Niatsetski, C. Koedooder, N. van Wieringen, A. Bel, B.R. Pieters, T. Alderliesten. (2017) *Improved class solutions for prostate brachytherapy planning via evolutionary machine learning.* Radiotherapy and Oncology 123, Supplement 1, S96–S97.

S.C. Maree, L Ortiz-Gracia, C.W. Oosterlee. (2017) *Pricing early-exercise and discrete barrier options by Shannon wavelet expansions.* Numerische Mathematik 136 (4), 1035–1070.

About the author

Stef Maree started in 2009 with the bachelor program "Applied and Industrial Mathematics" at the Delft University of Technology, from which he received his bachelor degree in 2012. His bachelor project focused on matrix algorithms for financial applications, supervised by dr. L.A. Grzelak and prof. dr. ir. C.W. Oosterlee and was conducted at Rabobank International, Utrecht. After his graduation, he spent a year as board member of a student association in Delft. In 2013, he continued to pursue a Master's degree in Applied Mathematics at the Delft University of Technology, and obtained his degree cum laude in 2015. During his Master's, Stef completed an internship at the Numerical Algorithms Group (NAG) in Manchester, UK, on the development of fast option pricing algorithms under supervision of dr. J. du Toit. His graduation project was conducted in collaboration with dr. L. Ortiz-Gracia and prof. dr. ir. C.W. Oosterlee at Centre de Recerca Matemàtica, Barcelona, Spain, on the development of fast option pricing algorithms using wavelet expansions. After graduation, he continued to work on the topic of his graduation project as a researcher at Centrum Wiskunde & Informatica (CWI), the national research institute for mathematics and computer science in the Netherlands. In 2016, he started a PhD project at the Amsterdam University Medical Centers, location Academic Medical Center (AMC) under daily supervision of dr. T. Alderliesten (AMC) and prof. dr. P.A.N. Bosman (CWI). This project, of which the result lies in front of you, was performed in collaboration with industry partner Elekta, and the Life Sciences and Health group at CWI.

Acknowledgements

Thank you Peter, Coen, Tanja, and Arjan, for being my supervisory team. All the critical reviews and insightful suggestions have been a tremendous help to me.

Specifically, Peter and Tanja, thank you for your daily supervision and mentoring. Both of you had your own impact on me in many different ways, but I am going to address you here simultaneously, because it has always been me and you as a team. You were the ones always there, from the beginning until the very end of this journey, to help me overcome every setback. In our meetings you often suggested many new ideas, although it was not uncommon that I showed results of a completely different approach a week later. I like to take this opportunity to apologise for that. Thank you for all the effort you put into our project, and for your personal support. I am very grateful for this.

My committee members, thank you for evaluating my thesis. Although we did not work together, many of your works were an inspiration to me.

My project-partners in crime, Hoang, Anton, and Marjolein. All the research, coding, brainstorming, coffee and tea moments. I could not and would not have made it without you. Hoang, you were there to answer all my naive fresh-out-of-university questions when I started. Your knowledge and understanding of evolutionary algorithms is immensely impressive and your humour cheers anybody up. A few months later, I could share my just-gained brachytherapy knowledge with Anton and Marjolein. Anton, I could ask you questions at midnight and you would have it implemented before I even woke up. And finally, but definitely not least, Marjolein. We shared an office for the first two years or so, but at some point, we discussed so many ideas that I had to switch offices to become productive again. Your laugh is contagious and so is your motivation, it was a pleasure working together with you.

I would also like to thank the rest of our multidisciplinary team at the Academic Medical Center: Bradley, Niek, Kees, Danique, Henrike, Karel, and Lukas. Your expertise steered us researchers into developing a tool that is now in clinical use. Especially memorable was a history lesson on brachytherapy that Lukas gave me while in the operating room, which was a rather shocking experience for a mathematician. Thanks to Ernst and Matthijs, I am happy to have gotten the

opportunity to supervise you. And thanks to Inka, Emmy, Jennifer, and Dirk for allowing us to constantly look over your shoulder to get to understand the practical side of brachytherapy. I would like to thank Erik for sharing an office with me, and for his insights in how software should be designed. And of course, Rob and Ernst: sorry for mailing you at the strangest hours to fix my computer. My fellow PhD students, Laura, Caspar, Eelco, Gerben, Kleopatra, Akke, Daan, Pierre, Ziyuan, Karin and Margot. I have a slight addiction to coffee now, thanks! Special thanks go out to Janna, Sophie, and Peng for making the ESTRO conferences such amazing experiences. And thanks to the entire department of radiation oncology, you all made the Christmas drinks, karaoke evening, and all coffee moments valuable memories that I will never forget.

To our industry partner in the project, Elekta, represented by Yury, Joost and Bob: thank you for believing in us researchers who can be a little bit detached from the real world by guiding us into the right direction. I would also like to attribute a special thanks to Rob, who is sadly enough no longer with us. At the start of my project, Rob supervised me daily, well in his retirement, and he introduced me to his immense self-developed treatment planning suite QRT, which was the basis and inspiration for me to start with the development of our own code, that thanks to my many colleagues, ultimately ended up in the clinic. I will also never forget that array indices start at 0 in C++, something Rob really did not like.

My second home these last years was the Centrum Wiskunde & Informatica (CWI). Thanks to everybody in the Life Scienes & Health group, and specifically, Leen, Solon, Jasmijn, Monika, and Arkadiy, for always making me feel at home, even though I was not officially employed there. Thank you, Timo: we started with a small project that resulted in me staying for another six months (thanks Leen and Peter), and a promising new research line for both of us. And of course, thank you Marco, for introducing me to bouldering with Anton, and to the best pizza place in Amsterdam. Your house was always open for me. Or should I thank Dafni for that? And Dirk, thank you for mentoring me occasionally while I was in Utrecht, we do seem to bond quite well on the social side of research.

And thanks to all other colleagues at CWI, the Scientific Computing group, with Daan, Nikolaj, Laurent and Prashant, and thanks to Sander, Tommaso and Dmitrii. And of course, thank you Kees, you deserve a special mention as you were the one that motivated me to go into research, and the one who attempted to keep me there afterwards. And last, but definitely not least: Nada. Thank you for bringing me a cappuccino outside that time I forgot to register; your kindness is inspiring.

On a more personal note, I would like to thank my friends, where I could always

go to for discussions or distractions. Harmjan, you made me employee of the month, but you definitely deserve a lifetime achievement award. René, your dedication convinced me to cycle further and higher than I ever dreamed of. My paranimphs and good friends Pim and Wietse, one next door, the other often in a random country, or on our couch. And of course all my friends from *octo*, some of who I see regularly, but many of you way less than I wanted to. I am going to make it up to you now that this book is finished.

Dan natuurlijk mijn familie, die ik nooit genoeg zou kunnen bedanken. Paps en mams, bedankt dat jullie me altijd onvoorwaardelijk hebben gesteund, in me hebben geloofd en de mogelijkheden hebben gegeven om te komen waar ik nu ben. Bij jullie was ik altijd welkom, en jullie vonden dan altijd dat ik te snel weer weg ging. Loes en Inge, om te plagen noemde ik jullie vaak Zusje 1 en 2, maar niet altijd in die volgorde. Ik beschouw de woorden "Mijn broer is bijna 30 en ook nog niet volwassen", maar als een compliment. Oma, ik denk nog vaak aan je. Wat hebben wij een mooi afscheid gehad, met de hele familie bij elkaar, en opa die de boel met trotse blik overzag. Door die fijne herinneringen mis ik het stiekem enorm dat iedereen zo lekker om de hoek woonde, maar dat klinkt zo volwassen, en dat ben ik dus nog niet volgens Zusje 2.

En degene die twijfelde of ze in mijn dankwoord genoemd zou worden heb ik voor het laatst bewaard. Lotte, jij bent enthousiast en gemotiveerd. Dat straal je uit naar de wereld, en al helemaal naar mij. Jij hebt er altijd in geloofd dat dit boekje er zou komen, ook als ik dat zelf even niet meer zag. Jij maakt mijn slechte dagen goed, en mijn goede dagen fantastisch. Je bent mijn maatje, en een lockdown is met jou een feestje. Dus dat plekje in mijn dankwoord heb jij meer dan wie dan ook verdiend, net als dat plekje in mijn hart. Dankjewel voor alles, datameisje.

&