Incorporating Dose-Response Models in BRIGHT for Cervical Cancer Brachytherapy Treatment Planning

Author:
Thijs Penning
Student number 5266297

To obtain the degree of Master of Science in:
Computer Science, track Software Technology,
at the Delft University of Technology,
Faculty Electrical Engineering, Mathematics and Computer Science.

Thesis committee:

Prof. dr. Peter A.N. Bosman	CWI, TU Delft, supervisor
Dr. Anton Bouter	CWI, supervisor
Dr. Tanja Alderliesten	LUMC, supervisor
Dr. David M.J. Tax	TU Delft

TU Delft Delft University of Technology

CWI Centrum Wiskunde & Informatica

Preface

This thesis represents the end of my academic journey at the Delft University of Technology. I followed both my Bachelor in Computer Science and Engineering and Master in Computer Science here. Over the years, I have gained extensive knowledge and developed many valuable skills in the computer science field. Besides the general curriculum, the flexible structure allowed me to focus my studies on my interests in software development and artificial intelligence research.

With this thesis, I was able to further explore these interests through the application of evolutionary algorithms to an interesting and impactful problem. The journey was challenging and at times, it seemed progress had stalled. However, through hard work and persistence, this thesis took shape resulting in the current version I am proud to present.

Finally, I want to acknowledge the people who have supported me. I thank Prof. dr. Peter Bosman, Dr. Anton Bouter, and Dr. Tanja Alderliesten for their supervision and guidance during my thesis. I am also grateful for my family, friends, and colleagues who assisted and encouraged me all these years.

Thijs Penning
Delft, June 2025

Abstract

Crafting and refining high-dose-rate brachytherapy treatment plans for cervical cancer is a time-consuming process. In recent years, BRIGHT was developed, an AI-based automated treatment planning method that provides not just one, but a set of optimized, patient-specific treatment plans, each with a different trade-off between objectives of interest. BRIGHT's plans are optimized using protocols that define guidelines on the delivered doses. In this thesis, we explore an alternative approach using dose-response models. These models provide insights into the estimated outcomes and risks of a treatment plan. Additionally, they offer great adaptability by including external patient characteristics. Currently, the use of these models remains limited to a feedback role. However, we instead include these models in BRIGHT to optimize them directly. Using one Tumor Control Probability (TCP) model and five Normal-Tissue Complication Probability (NTCP) models, we designed and tested several dose-response objective formulations and optimization techniques. We found that with the current models, these new objectives are insufficient as a replacement for BRIGHT's protocol-based coverage and sparing objectives. The produced plans have greatly improved dose-response outcomes but fall short in protocol compliance. Extending the existing objectives rather than replacing them proved more favorable. Average model improvements around 0.004 for NTCP are observed among the best coverage-sparing plans satisfying the protocol. Additionally, by sacrificing some sparing in the protocol-satisfaction range, improvements around 0.005 are possible for TCP and NTCP. Moreover, these dose-response-focused plans show distinct differences in their dose distribution favoring the dose-response targets compared to regular BRIGHT. Ultimately, the improvements we obtained are only marginal, and the clinical implications of this are unclear. The covariates of the models used in this thesis mostly overlapped with BRIGHT's objectives and did not fully extract their potential. Nevertheless, this thesis proves that the concept is viable and builds a foundation for this technique for when more and better dose-response models become available.

Contents

List of Figures

List of Abbreviations

Abbreviation	Definition
AI	Artificial Intelligence
CPU	Central Processing Unit
CTV_{HR}	High-Risk Clinical Target Volume
CTV_{IR}	Intermediate-Risk Clinical Target Volume
DCP	Dose-Calculation Point
DVI	Dose-Volume Index
EA	Evolutionary Algorithm
EBRT	External Beam Radiation Therapy
EQD2	Equivalent Dose in 2 Gy Fractions
FOS	Family of Subsets
GBO	Gray-Box Optimization
GOM	Gene-pool Optimal Mixing
GOMEA	Gene-pool Optimal Mixing Evolutionary Algorithm
GPU	Graphics Processing Unit
GTV_{res}	Residual Gross Tumor Volume
HDR	High-Dose-Rate
ICRU	International Commission on Radiation Units and Measurements
LT-FOS	Linkage-Tree Family of Subsets
MO	Multi-Objective
MO-RV-GOMEA	Multi-Objective Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm
MRI	Magnetic Resonance Imaging
NT	Normal-Tissue
NTCP	Normal-Tissue Complication Probability
OAR	Organ at Risk
RI	Response Index
RV	Real-Valued
TCH	Weighted Tchebycheff Distance
TCP	Tumor Control Probability

Chapter 1

Introduction

Cancer is a disease affecting millions each year worldwide. In this thesis, we specifically focus on cervical cancer, a type that affects many women. A very common and effective treatment for this type of cancer is brachytherapy in combination with External Beam Radiation Therapy (EBRT) [24]. In recent years, an automated treatment planning tool called BRIGHT has been developed that helps create and optimize these brachytherapy treatment plans [17, 33]. Our work focuses on extending this tool to include dose-response models in this plan-optimization process to improve tumor control and normal-tissue complications outcomes.

1.1 Problem Statement

Many factors, from one's lifestyle to their genetics, influence the risk of developing cancer. It is a diverse and widespread disease that is hard to prevent and difficult to cure. In 2022 alone, there were almost 20 million new cases and 9.7 million deaths spread over 36 types of cancer in 185 countries. Cervical cancer ranks amongst the top five cancers affecting women here, counting around 661 thousand cases and 348 thousand deaths [11].

Several of these cancer types, including cervical cancer, are commonly treated with EBRT and brachytherapy [23, 24]. First, the EBRT uses high-energy beams to irradiate the cancer from outside the body. Next, the brachytherapy applies an additional dose of radiation from inside the body directly to the tumor with reduced spread to the surrounding organs. Brachytherapy has a number of variants; this thesis focuses on treatment planning for the High-Dose-Rate (HDR) variant.

In HDR brachytherapy, a radioactive source is temporarily inserted using an applicator with several tubes to irradiate the tumor close by from all directions. By controlling how long the source is halted at various locations, the distributed dose can be controlled. A treatment plan then defines all these halting times. In practice, these plans are often handcrafted and refined until an acceptable distribution is found. This refinement process and the acceptability of a plan are based on protocols and dose-response relationships extracted from studies on past cases [18, 44] and on dose distributions.

Manually creating these treatment plans is an iterative and time-consuming process.

This limits the options that can be explored, making it harder to fine-tune the patient's needs. To speed this up and explore more alternatives, Artificial Intelligence (AI) can be used. Brachytherapy treatment planning can be defined as a multi-objective problem trying to optimize multiple doses delivered to the tumor and the Organs at Risk (OARs). Such problems can efficiently be solved using Evolutionary Algorithms (EAs), which the tool BRIGHT leverages to tackle this problem [17, 33]. Given a protocol of planning aims for the dosages, BRIGHT produces a set of plans considering various tumor coverage and OAR-sparing trade-offs. It can produce more plans faster and often with better coverage and sparing as well. In clinical use, these plans initially still required some fine-tuning by the physicians [4]. However, further developments of BRIGHT have since addressed these problems and reduced the amount of fine-tuning needed. Moreover, as an assisting tool, it still helped speed up the planning process and provided useful alternative trade-offs [4].

BRIGHT optimizes plans based on a protocol consisting of dosage planning aims. This approach matches part of the clinical procedure and fits plans to general guidelines. Another possible plan-evaluation method is to use dose-response models. Using patient- and plan-variables, these models provide feedback on plans through predictions of their expected outcomes, such as tumor control and normal-tissue complications. They offer a direct relation between a plan and its post-treatment effects, and their flexibility allows them to better adapt to specific patient characteristics.

Currently, the use of these dose-response models remains limited to a feedback role such as in [19]. However, given that this feedback can be used when determining the clinical acceptability of plans, these models might be worth optimizing. This thesis researches the validity of this. Integrating these models into BRIGHT's optimization process to produce plans with an increased focus on the dose-response outcomes.

1.2 Research Questions

The goal of this research is to integrate dose-response models that predict cervical cancer HDR brachytherapy treatment outcomes into the planning optimization of BRIGHT. At its core, this thesis aims to answer the following main research question:

MRQ: *How can dose-response models be integrated into BRIGHT to optimize high-dose-rate brachytherapy treatment plans for cervical cancer that balance tumor control and normal-tissue complications?*

The first stage of this research focuses on how the dose-response models can be integrated into BRIGHT and the results this brings. On their own, a dose-response model provides useful insights about a specific post-treatment effect. To optimize a set of these models, they need to be carefully combined into objectives that preserve their essence while striking a good optimization balance. Next, optimizing these dose-response objectives in place of BRIGHT's current target-coverage- and OAR-sparing-focused objectives should provide similar quality plans based on the clinical protocols. We address these challenges through the following two research questions:

RQ1: *How can dose-response models be represented as objectives suitable for optimization?*

RQ2: *How do dose-response objectives perform when replacing BRIGHT's coverage and sparing objectives in optimization?*

Based on results from the first stage, the second stage of this research focuses on extending BRIGHT with the dose-response objectives rather than substituting its existing objectives. This transforms the optimization problem from multi-objective to many-objective. Given the increased complexity of the many-objective space, a specialized extension using Tchebycheff scalarizations is also explored. Additionally, the area of interest relatively shrinks with this increased search space. Further efforts are made to direct the optimization towards the area of interest. To achieve all of this, we defined three more research questions:

RQ3: *Can extending BRIGHT to many-objective optimization with dose-response objectives improve Tumor Control Probability (TCP) and Normal Tissue Complication Probability (NTCP) outcomes of the generated plans?*

RQ4: *How can the use of Tchebycheff scalarizations in the optimization algorithm improve outcomes in the extended many-objective space?*

RQ5: *What methods can be employed to steer optimization outcomes in the many-objective space toward clinically desirable regions?*

1.3 Thesis Outline

This thesis is structured as follows. Chapter 2 explains the relevant background information and techniques. Next, the methodology and experiments about the formulation and performance of dose-response optimization in BRIGHT are described in Chapter 3 and Chapter 4, respectively. Continuing, Chapter 5 details the techniques used to solve our problem in the complex many-objective space. Furthermore, the experiments and results for the various many-objective solutions are shown in Chapter 6. Then, in Chapter 7, we analyze our findings and discuss potential directions for future research. Finally, Chapter 8 concludes this work.

Chapter 2

Background

This chapter explains the relevant background and gives an overview of the related work that is key to this thesis. To start, we describe the clinical practices within cervical HDR brachytherapy and outline the aspects targeted to optimize treatment plans. Next, we will delve into the available options and the specific variant of dose-response models used, as this is the main focus of this work. Then, we introduce several EA techniques used by BRIGHT and the contributions introduced in this thesis. To conclude, the specifics of BRIGHT are explained in detail.

2.1 High-Dose-Rate Brachytherapy

Brachytherapy is a form of internal radiation treatment used to treat cancer patients. Specifically, high-dose-rate brachytherapy is a variant of this treatment with a short operating time. Brachytherapy and its HDR variant are used to treat various types of cancer including prostate [16, 38], head-and-neck [39], eye [36, 40], breast [1, 3], and cervical [31, 42, 59], with this work focusing on the latter. For cervical cancer treatments, an applicator consisting of thin, flexible tubes is inserted in the region of the tumor; additional hollow needles can be inserted if necessary. A miniaturized radioactive source, commonly Iridium-192 or Cobalt-60 [53, 54], is passed through the applicator and the needles for several minutes to irradiate the tumor. A treatment plan defines how to move the source to maximize tumor coverage while minimizing exposure to the surrounding organs. The radioactive source does not move in a continuous motion. Rather, it moves from one stopping point to the next. At these stopping points, called dwell positions, it is held for a corresponding dwell time defined by the treatment plan before moving on.

2.1.1 Treatment Plan Quality

To evaluate a treatment plan and gain insights into the dose received by each organ, the Dose-Volume Indices (DVIs) are used. In practice, the quality of a plan is often assessed based on a protocol and the dose distribution. These protocols consist of several dose and volume indices for the various tumor targets and surrounding organs. These indices are defined as the following:

- D_v^o is a dose index reporting the dose received by the most irradiated volume v of organ o.

- V_d^o is a volume index reporting the volume of organ o that received at least a dose d.

The dose and volume values used by these indices can be represented in multiple formats. Volumes are a percentage % of the organ's volume or a specific value, commonly reported in cubic centimeters cm^3 or cc. Doses also have a relative representation as a percentage of the planned radiation dose (commonly 7 Gy for 4 fractions [28, 41, 43]) and a direct option reported in Gy. This dose representation in Gy is also used in two ways, either with the reported dose value or as an Equivalent Dose in 2 Gy Fractions (EQD2). The formula to calculate the EQD2 given a reported dose value is provided in Equation 2.1.

$$\text{EQD2} = D \cdot \frac{d + \alpha/\beta}{2 + \alpha/\beta} \tag{2.1}$$

Here D is the total dose over all the fractions, d is the dose per fraction, and α/β represents the intrinsic radiosensitivity of the target tissue. Note that if the patient underwent external beam radiation therapy first, then that delivered dose is added to the EQD2.

In a protocol, these DVIs are formulated as planning aims with a target dose or volume to reach. For instance, an organ or target O might have the following DVIs, $D_{90}^O > 90\ Gy$ and $D_{90}^O < 95\ Gy$. These DVIs indicate that the 90% of the volume of O should receive more than 90 Gy and less than 95 Gy (both in EQD2). Similarly, $V_{100}^O > 80\%$ indicates that 80% of the volume of O should receive at least 100% of the planning radiation dose.

The exact DVI values cannot be calculated as that would require infinite precision. Instead, they are approximated using so-called Dose-Calculation Points (DCPs) [46]. Thousands of these points are sampled uniformly for each target volume and OAR to model the volumes. Then, for each point, the received dose is calculated as the sum of all received doses from each dwell position for the corresponding dwell times. Given the dose received in each DCP, the dose and volume indices can be calculated as shown in Equation 2.2 and Equation 2.3, respectively. Let $\vec{P}^o$ denote the vector of all DCP doses p in organ o, with $\vec{P}^{o,s}$ the descending dose-sorted version of $\vec{P}^o$. Furthermore, v_p^o is the volume represented by each DCP in organ o.

$$D_v^o = \vec{P}^{o,s}_{\lfloor \frac{v}{v_p^o} \rfloor} \tag{2.2}$$

$$V_d^o = \frac{|\{p | p \in \vec{P}^o \wedge\ p > d\}|}{|\vec{P}^o|} \tag{2.3}$$

2.1.2 The EMBRACE II study

EMBRACE II [44] is an observational multi-center study about the effects of EBRT and Magnetic Resonance Imaging (MRI) guided brachytherapy treatments of cervical cancer. Throughout several checkups up to five years post-treatment, they gather insightful data about outcomes, such as tumor control, survival rates, side effects, and quality of life. This

data is useful to build predictive models that can estimate treatment outcomes before delivery. We will delve deeper into the concept of these dose-response models and the ones applied in this work in Section 2.2 and Section 3.1, respectively.

As with any research observing outcome trends, all treatments included in the study adhered to a standardized setup. This setup prescribed $4 \times 7Gy$ fractions of HDR brachytherapy, using EQD2 values with $\alpha/\beta = 10$ for targets and $\alpha/\beta = 3$ for OARs, and the EQD2 EBRT inclusion as 45 Gy/25 fractions [44]. Using Equation 2.1, the EQD2 EBRT addition equates to $45 \cdot \frac{1.8+10}{2+10} = 44.25Gy \approx 44.3Gy$ for targets and $45 \cdot \frac{1.8+3}{2+3} = 43.2Gy$ for OARs. Given this, we can refine our EQD2 equations to Equation 2.4 and Equation 2.5, with d as the dose per fraction.

The EMBRACE II setup also provides a protocol of planning aims and constraints for the DVIs. These are expressed in EQD2 doses and percentage volumes for the targets and cubic centimeters volumes for the OARs, consistent with the other setup parameters. It covers various targets, including the High-Risk (CTV_{HR}) and Intermediate-Risk (CTV_{IR}) Clinical Target Volumes, the Residual Gross Tumor Volume (GTV_{res}), and point A where the uterine artery and ureter cross. The surrounding organs, such as the bladder, rectum, sigmoid, and bowel. And an International Commission on Radiation Units and Measurements (ICRU) point, the ICRU Recto-vaginal. The protocol is summarized in Table 2.1.

$$\mathrm{EQD2}_{\mathrm{Target}} = \mathrm{EQD2}_{10} = (4 \cdot d) \cdot \frac{d+10}{2+10} + 44.3 \tag{2.4}$$

$$\mathrm{EQD2}_{\mathrm{OAR}} = \mathrm{EQD2}_{3} = (4 \cdot d) \cdot \frac{d+3}{2+3} + 43.2 \tag{2.5}$$

Table 2.1: The planning aims and constraints of the EMBRACE II protocol. Table adapted from [44].

Target	D90 CTV_{HR} $EQD2_{10}$	D98 CTV_{HR} $EQD2_{10}$	D98 GTV_{res} $EQD2_{10}$	D98 CTV_{IR} $EQD2_{10}$	Point A $EQD2_{10}$
Planning Aims	$>$ 90 Gy $<$ 95 Gy	$>$ 75 Gy	$>$ 95 Gy	$>$ 60 Gy	$>$ 65 Gy
Constraints	$>$ 85 Gy	-	$>$ 90 Gy	-	-
OAR	Bladder D_{2cc} $EQD2_3$	Bowel D_{2cc} $EQD2_3$	Rectum D_{2cc} $EQD2_3$	Sigmoid D_{2cc} $EQD2_3$	ICRU Recto-vaginal $EQD2_3$
Planning Aims	$<$ 80 Gy	$<$ 70 Gy	$<$ 65 Gy	$<$ 70 Gy	$<$ 65 Gy
Constraints	$<$ 90 Gy	$<$ 75 Gy	$<$ 75 Gy	$<$ 75 Gy	$<$ 75 Gy

2.2 Dose-Response Models

Dose-response models describe the relationship between the quantity of dose applied and the effect it causes. This can be applied to various instances, from predicting the impact of pollutants on ecosystems [45] to estimating infection risk from pathogen exposure [55].

In this work, we focus on its use with radiation doses and other relevant patient-health factors to predict tumor control and OAR toxicity. Estimates for these post-treatment effects provide valuable insights when assessing plans. Furthermore, the inclusion of additional health factors helps clarify patient-specific needs.

In our radiotherapy use case, the dose-response models are split into two categories. These categories are tumor control probabilities, which predict the chance of killing the tumor, and normal-tissue complication probabilities, which predict the chance of specific health complications [5]. Together, these categories balance a treatment plan. Increasing the dose will improve the TCP values but decrease the NTCP values and vice versa when decreasing the dose.

2.2.1 Cox Proportional Hazards Models

The Cox proportional hazards model, also known as a Cox regression model, is a dose-response model for time-to-event data [13]. Given a dose and possibly other covariates, it models the risk of an event over time. A key assumption of this model is that the hazard ratios between individuals are proportional over time. The formula to calculate this risk or hazard given a time is formulated in Equation 2.6.

$$h(t|\mathbf{X}) = h_0(t) \cdot e^{\beta_1 X_1 + \beta_2 X_2 + ... + \beta_p X_p} \tag{2.6}$$

In the equation above, $h(t|\mathbf{X})$ is the hazard at time t given covariates $\mathbf{X} = (X_1, X_2, ..., X_p)$, $h_0(t)$ is the baseline hazard at time t, where all covariates are zero, and $\beta_1, \beta_2, ..., \beta_p$ are the covariate coefficients. This setup makes Cox regression both flexible in its handling of time and interpretable in how the covariates increase the hazard given their coefficients.

With their flexible and interpretable setup, Cox regression has been used to model the influence and risk of various factors regarding radiotherapy and cervical cancer treatment. These uses range from survival analysis [12] to investigating external factors such as someone's age or their age of marriage [26] and research into recurrence and its effects [30]. Furthermore, using multiple models, they can also provide more general feedback on treatment plans before applying them as EviGUIDE does [19].

2.3 Evolutionary Algorithms

Evolutionary algorithms are metaheuristics algorithms within the class of evolutionary computation that employ biologically-inspired techniques to solve difficult problems. The optimization procedure of an EA is inspired by Darwinian evolution to converge towards the best solutions [21]. Based on their biological counterparts, these solutions, also called individuals, have a genotype, which is defined as a string of variables or genes. In the algorithm, a group of these individuals, referred to as a population, is optimized simultaneously. In this population, traits or genes are refined through variation of the better-performing solutions as evaluated by their fitness, similar to how this would occur in nature.

A typical EA is initialized with a random population, and through selection, crossover, and mutation, its solutions should gradually improve over many generations [2, 25]. Each

generation starts with a population of individuals, which, based on their fitness and random chance, are used to pass on their genes to the next generation of the population. The selection, crossover, and mutation operations guide the process of variation to, over many generations, explore the search space and converge toward high-quality solutions. These operations are explained further in Section 2.3.1.

Optimization algorithms can be categorized into three types depending on their ability to exploit domain knowledge. On one extreme, there is white-box optimization, where everything about a problem is well understood and can be exploited to solve it more efficiently. On the other extreme, there is black-box optimization, where no information about the problem is available except an evaluation function. In between exists Gray-Box Optimization (GBO), where limited information about the problem structure is present. This information could unlock partial evaluations. For partial evaluations, only the part of the evaluation that changed is recalculated rather than the whole thing, which can result in substantial speedups. Fortunately, our treatment planning is such a problem allowing for more efficient evaluations. How these partial evaluations are possible for our problem is briefly explained in Section 2.4.

2.3.1 Selection, Crossover, and Mutation

Following the Darwinistic ”Survival of the fittest” principle, the selection operation filters out weaker solutions based on a fitness function. Fitness is a problem-specific attribute defining what constitutes a better or fitter individual. For instance, in the well-known OneMax problem, a binary string is optimized to include the maximum number of ones. The fitness for a six-variable solution ”001100” would be only 2, while ”110010” would get 3 for its number of ones. In the selection procedure, these higher-fitness individuals are more likely to get selected and pass their genes on to the next generation. In our OneMax example, this would result in ”110010” with fitness 3 being more likely to be passed on than ”001100” with fitness 2. With a properly defined fitness function, the selection pressure pushes the solutions in the desired direction [6, 20]. Note that real-world problems are often more complex than the OneMax example shown, and a clearly defined fitness function might not even exist.

Applied either before or after the selection operation, crossover combines the population or selected individuals to create new solutions. The complexity and implementation of the crossover operation depends on the EA, but the goal is always to combine one or more ”parent” solutions into one or more ”offspring” solutions [57]. To continue our simple running example, we will apply the basic one-point crossover method to it. This method takes two parent solutions and creates two offspring solutions by splitting both parents’ genes at one point. If both our individuals from before (”001100” and ”110010”) were selected and matched for crossover, then given ”random” crossover point 2, we would get the following offspring: ”000010” and ”111100”. Both offspring contain the first two variables from one parent and the other four from the other parent creating two brand-new solutions. As shown, crossover allows us to combine genes in an attempt to create even fitter offspring. One-point crossover is a simple method that explains this concept nicely, but these simple methods fail on more complicated and deceptive tasks [14, 15]. In modern EAs, many new

advanced methods exist, each with its niche.

Purely selecting fit individuals and performing crossover with them fails to explore outside of its initial information. These steps can only work with the genes present within the population. If some allele values are not present from the start or lost along the way, they are lost forever. A possible way to combat this would be to increase the population size, but especially for complex problems, this is unsustainable. Mutation stimulates exploration by introducing a small chance that an allele will randomly flip during crossover [29, 52]. When creating an offspring solution, each allele will have a tiny chance to mutate to a different value. This chance is kept low to not disturb the convergence too much and only apply about a handful of mutations per generation. Returning to our example once more, if given only those two individuals, the optimal solution is unreachable as the sixth allele was 0 for both. With the mutation operation in place, this allele will eventually mutate XXXXX0 $\longrightarrow$ XXXXX1, allowing the algorithm to find better solutions later on. It should be acknowledged that this effect works both ways. However, good genes are often more prevalent in the population, thereby mitigating the risks if this mutation chance is kept low enough.

2.3.2 GOMEA

Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) [56] is a state-of-the-art evolutionary algorithm. Compared to the regular EA structure, this algorithm varies most in the crossover operation by introducing Gene-pool Optimal Mixing (GOM). GOM uses a linkage model to represent the dependencies of the optimization problem. Exploiting these dependencies, GOM separately tests various minor modifications for dependent sets of genes. These modifications are taken from a pool of parents and only keep the beneficial changes. The design of this process allows for considerable speed-ups, particularly in a GBO setting where partial evaluations can be used for all these minor modifications.

The dependencies to exploit during GOM are defined by a Family of Subsets (FOS), most commonly a Linkage-Tree FOS (LT-FOS). As the name indicates, a FOS is a set of subsets originating from the same set. A mathematical definition is provided in Equation 2.7.

$$\mathcal{F} = \{\mathbf{F}^0, \mathbf{F}^1, ..., \mathbf{F}^{|\mathcal{F}|-1}\} \quad \text{where } \mathbf{F}^i \subseteq \{0,1,...,l-1\} \wedge i \in \{0,1,...,|\mathcal{F}|-1\} \tag{2.7}$$

The LT-FOS is a special FOS variant where the subsets can be placed in a hierarchical tree structure as displayed in Figure 2.1. An LT-FOS tree has all the variables in independent subsets at the bottom. Then, bottom-up, the subsets are hierarchically merged till the top-level subset contains all the variables completing the tree. The LT-FOS requirement is defined in Equation 2.8. This indicates that each FOS element of size larger than one is composed of exactly the elements of two smaller FOS elements.

$$\forall \mathbf{F}^i \; (|\mathbf{F}^i| > 1) \Rightarrow \exists \mathbf{F}^j, \mathbf{F}^k \Rightarrow \mathbf{F}^j \cap \mathbf{F}^k = \emptyset, \; |\mathbf{F}^j| < |\mathbf{F}^i|, \; |\mathbf{F}^k| < |\mathbf{F}^i|, \; \textit{and} \; \mathbf{F}^j \cup \mathbf{F}^k = \mathbf{F}^i \tag{2.8}$$

The LT-FOS creation depends on the problem. With clear gene-dependencies, it could be done beforehand. Otherwise, it will be created and refined while running GOMEA. In the

latter case, the LT-FOS is initialized randomly. Next, patterns in gene-relations are tracked in what is called mutual information. Based on this mutual information, the strongest gene-relations are merged iteratively until the LT-FOS tree structure is complete.

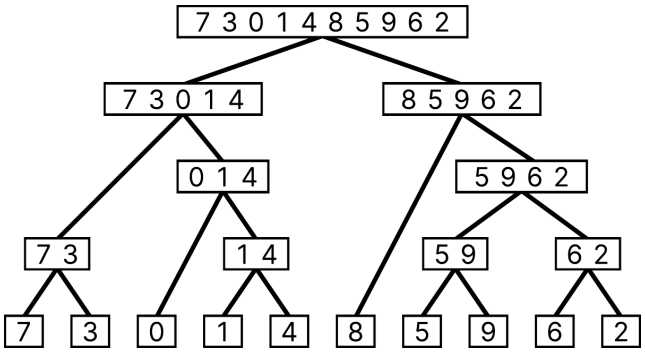

Figure 2.1: An illustration of the LT-FOS structure and how variables are dependent and independent depending on the subset.

Using these gene-dependency subsets from the LT-FOS, the GOM step transforms parent solutions into new offspring for the next generation. First, a parent solution is copied into a provisional offspring. Next, iteratively, the genes indicated by an LT-FOS subset are injected from a random donor parent in the pool. If this change results in a better or equal fitness value, then it is kept. Otherwise, it is discarded. This step is performed for each LT-FOS subset in random order. The donor parent for the gene-mixing operations is selected randomly each time. The entire process, starting from copying a parent solution into a potential offspring, is repeated for the whole population such that each original solution initializes one offspring. This GOM approach to crossover allows GOMEA to efficiently explore while preserving performance and exploiting dependencies in the problem structure. Moreover, in a GBO setting, partial evaluations substantially reduce the evaluation cost, further amplifying GOMEA's efficiency.

2.3.3 Real-Valued Optimization

In Real-Valued (RV) optimization, the genes are represented with values in a continuous range. Conventional crossover, consisting of swapping or mixing genes between solutions, cannot explore continuous ranges. Instead, real-valued GOMEA [8] uses variation through multivariate Gaussian distribution sampling to update the offspring. For each FOS element, instead of injecting genes from a donor parent, the injected genes are sampled from a multivariate Gaussian distribution defined over the FOS-indicated genes. Lastly, some extra techniques to control variance and evolution rates similar to [7] are used, but a detailed explanation falls outside the scope of this background.

2.3.4 Multi-Objective Optimization

Multi-Objective (MO) optimization aims to simultaneously optimize for multiple objective functions, finding solutions with varying objective trade-offs. A multi-objective problem

with m-objectives is described in Equation 2.9. Here, $\mathbf{x}$ is a solution from the set of solutions $\mathbf{X}$, and $f_i(\mathbf{x})$ is the i-th objective value for solution $\mathbf{x}$.

$$\min_{\mathbf{x}\in\mathbf{X}} \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), ..., f_m(\mathbf{x})) \tag{2.9}$$

Given the MO trade-offs, it is impossible to pinpoint an optimal solution as they each offer different benefits and drawbacks, thus giving no clear answer. Instead, MO optimization takes a different approach using Pareto dominance. Defined in Equation 2.10, a solution $\mathbf{x}^i$ Pareto dominates a solution $\mathbf{x}^j$, if and only if it is equal or better for each objective value and strictly better for at least one. Note that the equation considers objective minimization.

$$\mathbf{x}^i \succ \mathbf{x}^j \ : \ (\forall_k \in \{1,...,m\} \ : \ f_k(\mathbf{x}^i) \leq f_k(\mathbf{x}^j)) \ \wedge \ (\exists_k \in \{1,...,m\} \ : \ f_k(\mathbf{x}^i) < f_k(\mathbf{x}^j)) \tag{2.10}$$

Following this Pareto dominance definition, a solution is Pareto optimal if it is not dominated by any other solution. Formally, given a set of solutions $\mathbf{X}$, $\mathbf{x}^i \in \mathbf{X}$ is Pareto optimal if and only if, $\neg\exists \mathbf{x}^j \in \mathbf{X} \ : \ \mathbf{x}^j \succ \mathbf{x}^i$. The set of all these Pareto optimal solutions forms the Pareto form the Pareto set, which in objective space describes a Pareto front. Figure 2.2 visualizes an approximation of such a front. This front can be described as the boundary of the best solutions found for all trade-offs. With these techniques, MO optimization can optimize and present solutions in the multi-objective space without bias.

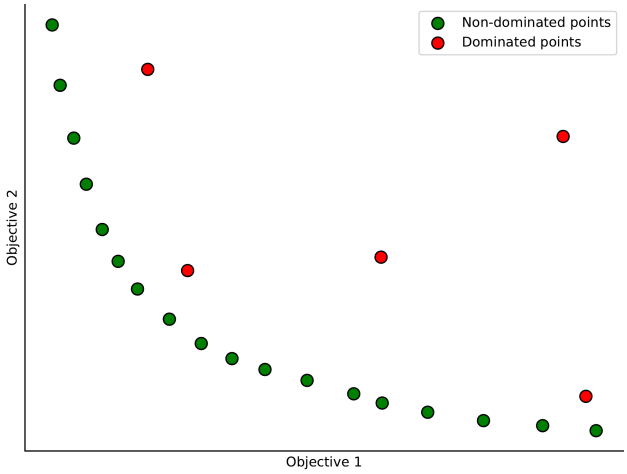

Figure 2.2: A visualization of an approximation front in a two-dimensional space, illustrating solutions for a minimization problem.

These MO optimization techniques have been applied to GOMEA producing multi-objective GOMEA [32] and Multi-Objective Real-Valued GOMEA (MO-RV-GOMEA) [9]. In this thesis, we focus on the real-valued variant. This variant includes the RV techniques explained in Section 2.3.3 with a slight MO adjustment. The multivariate Gaussian distribution sampling does not use the whole population. Instead, a selected subset of the best individuals in the population is partitioned into overlapping clusters grouping similar individuals. The resulting clusters are centered around different trade-offs in the multi-objective

space. Then, with separate LT-FOS and corresponding distributions for each cluster, they are used for variation with the whole population.

Another important technique used by EAs, especially for MO optimization, is elitism. An elitist archive stores the best solutions found throughout the generations. MO-RV-GOMEA exploits this archive to accelerate improvements and balance the population over the entire front. To replace the worst solutions that got stuck in local optima, random elitists are reintroduced into the population to spread their genes during variation. This improves the gene-pool and prevents parts of the multi-objective front from vanishing.

2.3.5 Tchebycheff Scalarizations

Pareto dominance-based multi-objective evolutionary algorithms, such as MO-RV-GOMEA, struggle with selection pressure toward the Pareto front in higher dimensional spaces. When optimizing problems involving four or more objectives, defined as many-objective problems, the performance of these algorithms diminishes. To solve this, Luong et al. [34] integrated Tchebycheff scalarizations into MO-RV-GOMEA to improve the performance on these many-objective problems.

The Weighted Tchebycheff Distance (TCH) metric in a multi-dimensional setting is defined as the largest weighted dimensional distance between two points. In an m-dimensional space with weight vector $\mathbf{w} = (w_1, ..., w_m) \ : \ w_i \geq 0 \ \wedge \ \sum_{i=1}^{m} = 1$, the TCH between a point $\mathbf{z}$, and a reference point $\mathbf{z}^*$, is defined in Equation 2.11.

$$TCH(\mathbf{z}, \mathbf{z}^*, \mathbf{w}) = \max_{1 \leq i \leq m} \{w_i | z_i - z_i^* |\} \tag{2.11}$$

Combining the multi-objective definition of Equation 2.9 and the TCH of 2.11, Luong et al. [34] defines the Tchebycheff scalarizations formulation for many-objective problems as Equation 2.12. Here, the reference point $\mathbf{z}^*$ is set as the utopian point ($z_i \leq \min\{f_i(\mathbf{x}) | \mathbf{x} \in \mathbf{X}\}$, and a weight vector $\mathbf{w}$ is assigned to each individual $\mathbf{x}$ from the population P.

$$\min_{\mathbf{x} \in \mathbf{P}} \{TCH(\mathbf{f}(\mathbf{x}), \mathbf{z}^* \mathbf{w})\} \ = \ \min_{\mathbf{x} \in \mathbf{P}} \{ \max_{1 \leq i \leq m} \{w_i | f_i(\mathbf{x}) - z_i^*\}\} \tag{2.12}$$

The TCH integration of [34] adds a few steps to the MO-RV-GOMEA process. First, at the start of the algorithm, the TCH weight vectors, defined as $\mathbf{w} = (w_1, ..., w_m) \ : \ w_i \geq 0 \ \wedge \ \sum_{i=1}^{m} = 1$, are generated. A dense pool of vectors is generated to ensure they are well-distributed. From this dense pool, a subset matching the population size is selected using greedy scattered subset selection [47]. Next, this same set of weight vectors is (re)distributed over the population each generation. This weight vector association uses TCH with the current population's ideal objective vector $\mathbf{z}^P$ as the reference point, defined in Equation 2.13.

$$\mathbf{z}^P = (z_1^P, ..., z_m^P) \ : \ z_i^P = \min\{f_i(\mathbf{x}) | \mathbf{x} \in P\} \tag{2.13}$$

Source code was provided by the authors of [34]. In this code, the authors normalize $\mathbf{z}^P$ and all points based on the objective ranges. Normalized, $\mathbf{z}^P$ should end up as the zero-vector; the other point's values should be between $[0, 1]$ for each dimension. Next, with each weight vector, the TCH between each (normalized) unassociated point and (normalized) $\mathbf{z}^P$ is calculated, after which it is assigned to the TCH-nearest individual. Finally, an individual's

associated weight vector is used to check for improvements during the GOM operation. Offspring modifications are only accepted if the improvement has a lower TCH; that is, it improves in the direction of the associated weight vector. Effectively, this Tchebycheff scalarizations integration introduces weight vectors that push the population into the shape and direction of the Pareto front.

2.4 BRIGHT

BRIGHT is a tool for automated brachytherapy treatment planning, aiming to improve planning time and quality. In standard treatment planning, creating the plan is a time-consuming process. Additionally, only a single plan is iteratively updated and checked, thus severely limiting the insight of the clinical expert into the treatment planning process. Alternatively, BRIGHT solves this by using MO-RV-GOMEA to optimize a whole population of plans. In only a few minutes, it returns a Pareto approximation front of plans, allowing the physician to choose the preferred trade-offs for the patient. Altogether, it generates more plans for better patient-plan matching, takes less time than the clinical approach, and the optimized plans tend to be preferred over the clinical result [35]. In clinical use, these improvements are slightly reduced, though still relevant, as many plans still required some fine-tuning [4].

A treatment plan is defined by its dwell times. This information determines how the source is moved through the patient and how the dose is delivered. To optimize these plans using an EA, each individual is defined exactly the same way, with the genes dictating how the dose is distributed. This definition will have clear gene-dependencies, which is good for the GOM crossover. Additionally, it enables partial evaluations during the GOM step. The DVIs used to calculate the objectives are calculated using large matrix operations with the dwell times and the dose calculation points. When testing a GOM modification, only a few of these dwell times are updated. Usually, all values would be expensively recalculated. Instead, partial evaluations are possible by calculating only the difference, considerably reducing computation time. This speedup allows BRIGHT to optimize further within the same time frame or increase precision by setting a higher number of DCPs.

2.4.1 Objective Formulation

BRIGHT combines the DVIs into grouped objectives while keeping the balance between tumor coverage and OAR sparing. First, the DVIs, defined in Equations 2.2 and 2.3, are represented by the normalized distance to their planning aims. The normalized distance calculation, differing for coverage and sparing targets, are shown in Equations 2.14 and 2.15 and Equations 2.16 and 2.17, respectively. Here, $D^{o,max}$ is the highest desirable dose for organ o and $V^{o,tot}$ the total volume of organ o.

$$\delta_c^{norm}(D_v^o) = (D_v^o - D_v^{o,aim}) \;/\; (D^{o,max} - D_v^{o,aim}) \tag{2.14}$$

$$\delta_c^{norm}(V_d^o) = (V_d^o - V_d^{o,aim}) \;/\; (V^{o,tot} - V_d^{o,aim}) \tag{2.15}$$

$$\delta_s^{norm}(D_v^o) = (D_v^{o,aim} - D_v^o) \;/\; D_v^{o,aim} \tag{2.16}$$

$$\delta_s^{norm}(V_d^o) = (V_d^{o,aim} - V_d^o) \;/\; V_d^{o,aim} \tag{2.17}$$

An objective of grouped DVIs, whose composition is explained in Section 2.4.2, is represented with a single objective value based on the largest DVI violations. Initially, this only considered the single worst-case. Later on, this was replaced with a weighted worst-case to also improve non-worst-case targets where freely possible [10]. In this weighted worst-case approach, the least violated DVI is assigned a weight of 1. Then, the weight is exponentially increased by a factor of 10 for each remaining least violated DVI. After the weights have been assigned they are normalized between 0 and 1. Given a grouped objective $\mathbf{O}$, its DVI set $\Delta_{\mathbf{O}}$, and w_i as the weight for the i-th DVI, the objective value is calculated as shown in Equation 2.18.

$$\mathbf{O} = \sum_{i \in |\Delta_{\mathbf{O}}|} w_i \delta^{norm}(\Delta_{\mathbf{O},i}) \tag{2.18}$$

2.4.2 Coverage, Sparing, and Added Aims

Currently, BRIGHT considers three objectives when optimizing cervical cancer treatment plans: coverage, sparing, and added aims. The first two, coverage and sparing, are the main objectives used to quantify plan quality in accordance with the EMBRACE II protocol, whereas the latter, added aims, is extra optimization guidance, customizable to align with the preferences of any potential medical clinic.

The coverage objective contains only dose-maximizing planning aims. These planning aims focus on the tumor targets containing several different risk areas of the tumor. An overview of all DVIs and their respective planning aims is provided in Table 2.2. The sparing objective contains only dose-minimizing planning aims. These planning aims primarily focus on the various OARs in the direct vicinity of the tumor. However, this objective also contains limits on the dose applied to the tumor. Table 2.3 provides an overview of the DVIs and respective planning aims. The added aims objective contains a mix of maximizing and minimizing planning aims over the tumor targets and Normal-Tissue (NT) regions. This objective can be used to introduce extra aims on some of the DVIs of the coverage and sparing objective, as well as introducing some new DVIs. Once more, an overview of the DVIs and their respective planning aims as used in this study is displayed in Table 2.4.

To quantify plan performance, BRIGHT has defined the ”golden corner” illustrated in Figure 2.3. The golden corner is an area in the objective space where both the coverage and the sparing objective values are positive. As the objectives are based on DVI distances to the planning aims, a positive value indicates that all planning aims are met. Therefore, plans in this corner adhere to all the coverage and sparing aims, making them generally preferred options.

Table 2.2: The DVIs used in the coverage objective along with their planning aims and constraints.

DVI	D98 GTV_{Res} $EQD2_{10}$	D90 CTV_{HR} $EQD2_{10}$	D98 CTV_{HR} $EQD2_{10}$	D98 CTV_{IR} $EQD2_{10}$
Planning Aims	$>$ 95 Gy	$>$ 90 Gy	$>$ 75 Gy	$>$ 60 Gy
Constraints	$>$ 90 Gy	$>$ 85 Gy	-	-

Table 2.3: The DVIs used in the sparing objective along with their planning aims and constraints.

DVI	D90 CTV_{HR} $EQD2_{10}$	Bladder D_{2cc} $EQD2_3$	Bowel D_{2cc} $EQD2_3$	Rectum D_{2cc} $EQD2_3$
Planning Aims	< 95 Gy	< 80 Gy	< 70 Gy	< 65 Gy
Constraints	-	< 90 Gy	< 75 Gy	< 75 Gy
DVI	Sigmoid D_{2cc} $EQD2_3$	ICRU Recto-vaginal $EQD2_3$		
Planning Aims	< 70 Gy	< 65 Gy		
Constraints	< 75 Gy	< 75 Gy		

Table 2.4: The DVIs used in the added objective along with their planning aims and constraints.

DVI	V200 Core CTV_{HR}	V50 CTV_{IR}	V100 Mid CTV_{IR}	V100 CTV_{HR}
Planning Aims	$> 99.5\%$	$> 99.9\%$	$< 25.0\%$	$> 99.9\%$
Constraints	$> 90.0\%$	$> 90.0\%$	$< 35.0\%$	$> 90.0\%$
DVI	V100 Mid NT	V100 Top NT	D90 Bot NT	
Planning Aims	$< 0.1\%$	$< 0.2\%$	$< 25.0\%$	
Constraints	$< 1.5\%$	$< 7.0\%$	$< 40.0\%$	

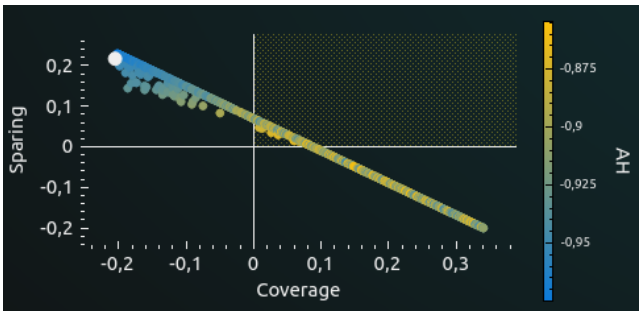

Figure 2.3: A front of generated plans displayed in the coverage-sparing space with the ”golden corner” marked, as shown in the BRIGHT application.

Chapter 3

Dose-Response Optimization in BRIGHT

To include dose-response relationships in BRIGHT, the models must be combined into optimizable objectives. This section aims to answer the first research question about how this can be realized. We first describe the models we use. Then, given the models, we introduce several normalization and aggregation methods to combine them to create the objectives.

3.1 Applied Dose-Response Models

We use the models from EviGUIDE [19], another tool assisting the brachytherapy cervical cancer treatment pipeline. EviGUIDE employs one TCP model for the high-risk clinical target volume and five NTCP models for the bladder, ICRU bladder point, rectum, ICRU recto-vaginal point, and bowel. All models included are Cox proportional hazards models. The covariates and coefficients are based on several works analyzing the EMBRACE data [27, 37, 48, 49, 50, 51, 60]. Due to limitations of available clinical data, we use an adapted version of the TCP model using only a subset of the covariates used by the version in [19]. A list of these models with their covariates and coefficients is displayed in Table 3.1.

We divide these models into two objectives categorized by their TCP and NTCP values, referred to as Response Index (RI) objectives RI_T and RI_N, respectively. As there is only one TCP model, it will make up the whole objective. On the other hand, the five NTCP models require a balanced aggregation method to combine them. Figure 3.1 displays the NTCP models plotted over their respective EQD2 dose. The figure shows that, in our operating range from the EBRT dose onward, these models operate in varying ranges. Therefore, some form of normalization is required to apply certain aggregation methods. These normalization and aggregation methods are discussed in the upcoming sections.

Table 3.1: An overview of the Cox proportional hazards covariates and coefficients used for the TCP and NTCP models.

Endpoint	**Covariate**	**Category**	**Baseline cumulative hazard**[1]	**Coefficient**
Local Tumor Control	Histopathological Type	SQ	0.2754	Ref
		AdSq		1.0976
	CTV_{HR} Volume	$<$ 30 cm3		Ref
		30-45 cm3		0.1903
		$>$ 45 cm3		1.0749
	Overall treatment time in days	continuous		0.0351
	CTV_{HR} D90	continuous		-0.0389
Pooled; Bleeding, Cystitis, Fistula G $\geq$ 2	Bladder D_{2cc}	continuous	0.0108	0.0333
Urinary Incontinence G $\geq$ 2	ICRU Bladder	continuous	0.0602	0.0132
Pooled; Proctitis, A/R bleeding G $\geq$ 2	Rectum D_{2cc}	continuous	0.0008	0.0752
Vaginal Stenosis G $\geq$ 2	ICRU Recto-vaginal	continuous	0.0662	0.0248
Flatulence G $\geq$ 2	Bowel D_{2cc}	continuous	0.0152	0.0335

[1] Risk over 5 years

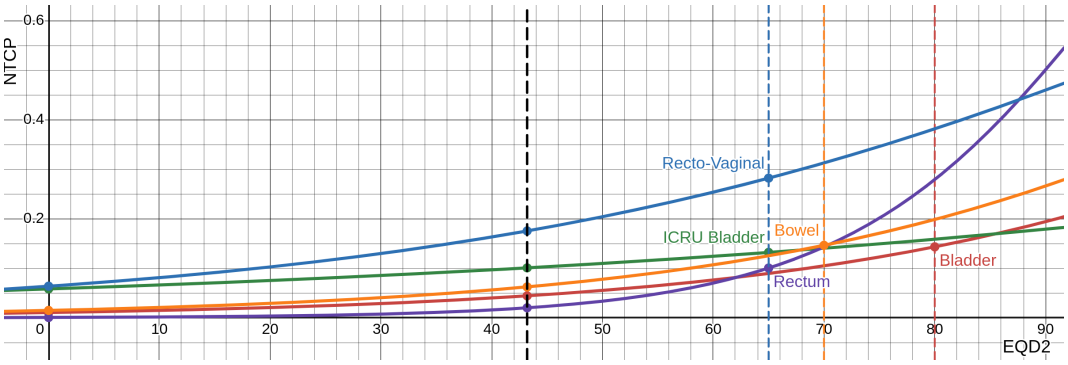

Figure 3.1: The NTCP models plotted over their respective EQD2 doses.

3.2 Target Normalization

The initial strategy for evaluating the objectives was a worst-case aggregation approach similar to the current objectives. The issue with direct worse-case comparisons was the ICRU Recto-vaginal model. As shown in Figure 3.1, it always had a worse value than the other models for any dose in the expected operating range. The output ranges had to be normalized, which meant fixing these models to reference points. With a lower and upper reference point, we normalize a part of the range to $[0,1]$ with the rest following the normalized formula.

The models use EQD2 doses, and all cases in this work treated the patients with EBRT before HDR brachytherapy. This provided two natural lower bounds to consider. It could be set either where the EQD2 in total was zero or where it was equal to the EBRT dose, thus zero for the HDR brachytherapy treatment. The upper reference had only one logical target, the planning aims of the DVI covariates. Four of the DVIs were part of the EMBRACE II protocol and had a planning aim defined. From Table 2.1, Bladder $D_{2cc} < 80$ Gy, Rectum $D_{2cc} < 65$ Gy, ICRU Recto-vaginal < 65 Gy, and Bowel $D_{2cc} < 65$ Gy. For the remaining model, we set ICRU Bladder < 65 Gy as the planning aim based on [44, 49].

These choices lead to two target-based normalization variants: long-range and short-range. Long-range normalization uses the whole range from the lower reference point at zero EQD2. Short-range normalization focuses on only the added dose with its lower reference point at the EBRT dose EQD2. Figure 3.2a and Figure 3.2b show the normalized models plotted over their respective doses for long-range and short-range, respectively.

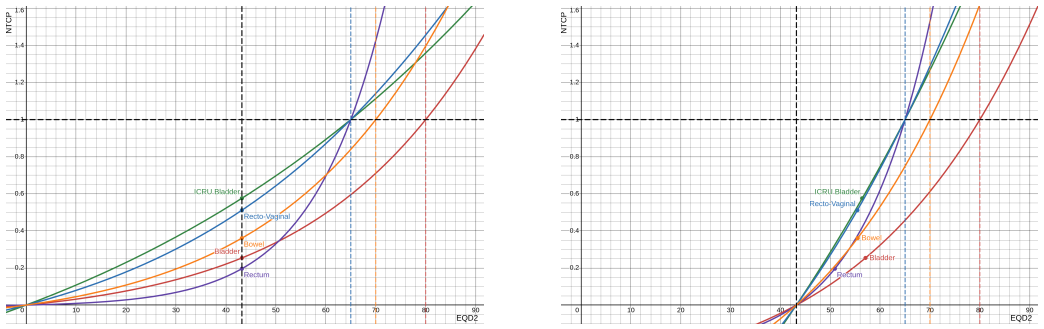

(a) Long-range normalized NTCP models.

(b) Short-range normalized NTCP models.

Figure 3.2: The target-normalized NTCP models plotted over their respective EQD2 doses.

Compared to before, these normalized models overlap in the expected operating range. With these overlaps, each model can be the worst case if it gets too close to its target. This is crucial for the worst-case aggregation to find balanced treatment plans. Equations 3.1, 3.2, and 3.3 formulate the long-range normalization, short-range normalization, and the target-based RI objective, respectively. Here, f^l and f^s are the long-range and short-range normalization functions, respectively. Additionally, m_i is the i-th model which has the current dose x_i, the 0 EQD2 dose $x_{i,0}$, the EBRT dose $x_{i,e}$, and the target dose $x_{i,t}$. Finally, DR is the set of dose-response models used to calculate the objective, and f can be either f^l or f^s depending on which normalization is applied.

$$f_l(m_i, x_i) = \frac{m_i(x_i) - m_i(x_{i,0})}{m_i(x_{i,t}) - m_i(x_{i,0})} \tag{3.1}$$

$$f_s(m_i, x_i) = \frac{m_i(x_i) - m_i(x_{i,e})}{m_i(x_{i,t}) - m_i(x_{i,e})} \tag{3.2}$$

$$RI = \max_{i \in |DR|} \{f(m_i, x_i)\} \tag{3.3}$$

Ultimately, both methods allow for effective worst-case aggregation. However, this target-based approach loses key aspects of the dose-response relationship. Inherently, these models cannot be weighted equally without sacrificing information. Normalizing between these zero- and target reference points only captures the dose-response model shape. The values attached to that range and the clinical relevance behind those values are lost. For instance, two models with identical shapes, normalized between NTCP values $[1\%, 5\%]$ and $[10\%, 50\%]$, are treated the same. Increases in complication probability are only relevant relative to the model range and shape. An NTCP increase of $2\% \rightarrow 3\%$ has the same effect as a $20\% \rightarrow 30\%$ increase in our example. Additionally, the severity of the health complications related to these probabilities is also ignored. Ultimately, this target-based normalization has its flaws concerning these dose-response relationships. Nevertheless, both variants are evaluated, though other options are also explored.

3.3 Multiplicative Aggregation

Multiplicative aggregation aims to combine the models while preserving their complication-predicting relevance. Additionally, all models contribute to the objective value of every evaluation. In this method, the objective value represents the probability of the desired effects by multiplying the individual model outputs. The probabilities multiplied are taken directly in the case of TCP. For NTCP, they are flipped to represent the desired non-complication outcome. The TCP and NTCP objective calculations are formulated in Equation 3.4 and Equation 3.5, respectively. When optimizing, these objectives will be minimized as $1 - RI$. Note that these objective values only represent the probability of the desired effects. They cannot be taken as the true probability since the multiplications do not consider model dependencies.

$$RI_T = \prod_{p \in TCP} p \tag{3.4}$$

$$RI_N = \prod_{p \in NTCP} (1 - p) \tag{3.5}$$

This natural approach to combining a set of probabilities strikes a delicate balance in model influence. Each percent risk increase affects the objective value, and these effects are more similar thanks to the flipped NTCP probabilities. The NTCP model complications tend to be the more unlikely outcome. Thus, the flipped non-complication results are more probable. When multiplied, percent changes are more influential in the lower ranges. A $5\% \rightarrow 6\%$ increase applies the same change factor as $20\% \rightarrow 24\%$ in another model. Using the higher flipped ranges, this change becomes $95\% \rightarrow 94\%$ and is matched by $80\% \rightarrow$

79.16%, a much smaller difference. Notably, the flip reduces the range disparity, but it also switches it to be survivability-based.

This aggregation method preserves much of the dose-response relationship, though complication severity is still not covered. However, all used models predict moderate-to-severe (Grade ≥ 2) events [58], thereby mitigating this issue. Additionally, if different severity models are combined, this method could easily be expanded by adding powers to the multiplied probabilities to add weight to specific models. Another minor drawback is that the objective value range is inconsistent. It depends on the number of models as probability multiplications only reduce a value. Small values and inconsistent ranges could prove awkward when extracting information from them or combining this method with other techniques, such as epsilon dominance [22].

3.4 Additive Aggregation

Additive aggregation is an adaptation of the multiplicative approach. Instead of combining the models through multiplication, they are added together and averaged. This approach entirely eliminates the range disparities as each percent change is uniformly represented in the mean operation. Due to this, the NTCP flip is not used as it adds no value. Without flipping, both objectives use the same formula shown in Equation 3.6. Here, DR is the set of dose-response models used to calculate the objective, which is TCP and NTCP for RI_T and RI_N, respectively. When optimizing, the NTCP objective will directly minimize RI_N, while the TCP objective minimizes $1 - RI_T$.

$$RI = \frac{1}{|DR|} \sum_{p \in DR} p \tag{3.6}$$

Similarly to the multiplicative setup, this approach does not cover complication severity but could easily introduce it. Models could be weighted using multiple entries in the mean operation to increase their influence. Furthermore, the mean operation generally stabilizes the objective value ranges, making them independent of the number of models. Consistent ranges could simplify integration with other techniques and visual information extraction. Finally, there is a duality in removing the range disparities. It improves model balance when dealing with diverse output ranges. However, compared to the multiplicative approach, the influence of models no longer slightly scales as their organs get more irradiated. This influence reduction could increase the chance of over-radiating organs.

Chapter 4

Experiments for Dose-Response Objectives

To answer the second research question, BRIGHT is optimized with the dose-response objectives for each of the objective formulations described in Chapter 3. The quality of the resulting fronts are compared to the regular BRIGHT version. This evaluation analyses how closely the generated fronts match for the DVI-based objectives and what the corresponding dose-response outcome change is. Ultimately, these experiments show the validity of the dose-response model objectives as an alternative to BRIGHT's current approach and the value that this outcome-based approach adds.

4.1 Experimental Setup

To ensure transparency and reproducibility, this section outlines all relevant aspects of the experiment setup. These aspects detail how the experiments were conducted, describing the dataset, hardware specifications, and experimental design. Additionally, they explain the various settings optimized and evaluated during the experimental process.

4.1.1 Dataset and Hardware Specifications

The experiments were performed using private real-world data from 16 anonymized patients treated at the Leiden University Medical Center. The data for each patient consists of a set of annotated MRI scans and several recorded data points. The annotated MRI scans indicate the dose applicator, tumor, and OAR positions. BRIGHT uses these positions and scans to calculate the applied dosage of a potential treatment plan and visualize them in a graphical interface. The recorded data points include information about the tumor and patient treatment, among other details. The dose-response model covariates use the cancer type and the duration of the treatment procedure extracted from this data.

Given the computationally expensive nature of AI and EAs, the hardware used has a powerful impact on the performance in a given time frame. Specifically, the Graphics Processing Unit (GPU) architecture is the most important factor as it handles the majority of the computations. These experiments used the GPU model *NVIDIA GeForce RTX 2080 Ti*

with a GPU memory of *11 GB*, driver version *560.35.03*, and CUDA version *12.6*. The supporting operations were performed using model *Intel(R) Xeon(R) Silver 4110 CPU @ 2.10GHz* Central Processing Units (CPUs). The server that ran the experiments had 4 GPUs and 2 CPUs of the aforementioned models.

4.1.2 Experimental Design

Running each patient with all available settings is a time-intensive process. To address this, the experiments were divided into two parts, tuning and evaluation. The tuning phase used a subset of the patients to gather information about the behavior of various parameter options. Then, based on the tuning results, the parameter values were fixed for the rest of the patients during the evaluation runs. Finally, in the evaluation phase, conclusions on the performance of the tested run types were drawn based on all patients; this included the fixed-parameters evaluation runs and the matching setting tuning runs.

Four patients were utilized for the tuning subset; the other twelve only appeared in the evaluation runs. The tuning patients were semi-randomly selected to cover scenarios of varying difficulties. These difficulties were determined using regular BRIGHT and preliminary dose-response optimization results. First, patient 1 represents a simple case where BRIGHT and dose-response optimization manage to achieve prominent results. Second, patient 2 is moderately difficult, with BRIGHT still reaching the golden corner, but the dose-response optimization is far off. Third, patient 3 is also moderately difficult, with both approaches barely falling short of the golden corner. Fourth, patient 4 is a challenging case where the best plans found are far off the desired results. Together, this set of patients provided insights into how each parameter option performs under varying circumstances to better determine the best values to extend to the entire dataset.

The performance of the settings tested in both phases was evaluated by analyzing the results over ten runs to ensure robust statistical analysis. Moreover, to reduce the variance between run types, we used ten predetermined seeds for sampling the DCPs, one for each run. The plans of all ten runs were compared in DVI- and dose-response-based spaces to determine the relative effects of each parameter and run type. Despite the results being three-dimensional, the comparisons use the two-dimensional fronts to reduce cluttering and focus attention on the potential of the tested methods. The combined results from all evaluated patients provided a detailed analysis of the effectiveness of each setting.

4.1.3 Run Types and Parameters

These experiments compare the performance of three-dimensional optimized plans using BRIGHT for various objective definitions; these are the run types. The regular version of BRIGHT using the coverage, sparing, and added objectives detailed in Section 2.4.2 serves as the baseline; further referenced as ”base”. Four candidate run types are compared to this baseline. First and second, run types ”long” and ”short” are target normalization methods described in Section 3.2. Third, ”multp” represents the multiplicative aggregation approach detailed in Section 3.3. Fourth, run type ”add” is the additive aggregation method described in Section 3.4. Each candidate BRIGHT variation is run with TCP and NTCP

objectives defined using its normalization or aggregation method and the added objective. Including the added objective for the candidate run-types assures a fair comparison as this objective focuses on additional goals beyond the other objectives. In total, five run-types are evaluated; in the rest of the experiments, we will reference these based on their tested methods: ”base”, ”long”, ”short”, ”multp”, and ”add”.

Optimizing treatment plans with BRIGHT uses several parameters for the plan-quality precision and the MO-RV-GOMEA algorithm. Here, we focus on five impactful parameters in particular. First, the number of DCPs used to calculate the DVIs for evaluation accuracy. Second, the time limit for how long we optimize. Third, the population size for how many solutions we optimize simultaneously. Fourth, the number of GOM clusters to control how many trade-offs the algorithm can focus on. Fifth, the elitist archive size to limit how many of the best-found solutions are stored throughout the optimization process. In these experiments, we fixed four of these parameters: 100.000 DCPs, 288 population size, 12 clusters, and 1.000 archive size. The number of DCPs was empirically established, while the population size, the number of clusters, and the archive size are well-defined for BRIGHT with three objectives. For the time limit, multiple options are tested in the tuning phase. Given the available hardware, time limits of 5, 10, and 15 minutes were chosen as improvements are expected to plateau within this range. Additionally, a time limit option of 30 minutes was added to verify if this was indeed the case. To summarize, 100.000 DCPs, 288 population size, 12 clusters, and 1.000 archive size are fixed, and time limits ”5m”, ”10m”, ”15m”, and ”30m” are tested during tuning.

4.2 Tuning Results

We determine the time limit parameter for the evaluation phase using the dose-response run types results in the TCP-NTCP space. These run types are the main focus of these experiments. Furthermore, their optimized spaces best display the improvements over time. Most cases in these settings showed little to no improvements when increasing the time limit. Specifically, the results of ”easy case” patient 1 were effectively identical for all options. Additionally, the aggregation approaches, ”multp” and ”add”, were less affected by the increased optimization time than their target-based counterparts. Figure 4.1 visualizes the time-limit performances over a select number of cases that best display the differences; additional plots are available in Appendix A.

The fronts plotted in Figure 4.1 show a general trend of the 15- and 30-minute time limit runs edging out the shorter variants. Though the 10-minute fronts remain relatively close, they cannot consistently extract the desired performance. If time is of the essence, it could be a viable option. However, these experiments evaluate the potential of the tested methods which better fit the longer options. Between the longer time limits, the difference is minimal. Results stagnate here; increasing the optimization time further is unlikely to improve results. The 30-minute fronts slightly outperform the 15-minute fronts for some of the target-based normalization cases in Figures 4.1a and 4.1b. Nevertheless, these improvements are negligible and not worth doubling the optimization time. Given these results, we set the time limit to 15 minutes for the evaluation phase. These findings align with intuition:

challenging cases needing a finer dose-distribution balance demand more optimization time, but beyond 15 minutes, the marginal gains are too small to warrant the extra runtime.

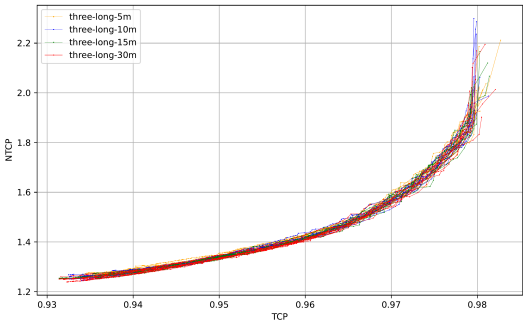

(a) Patient 2: Time limit comparison for normalization method ”long”

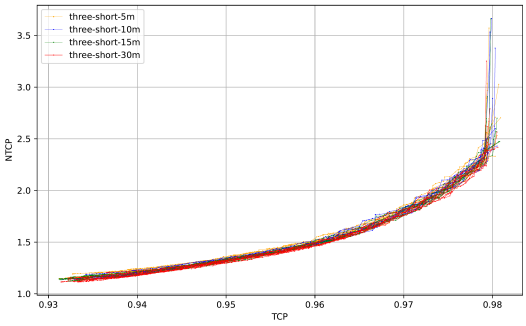

(b) Patient 2: Time limit comparison for normalization method ”short”

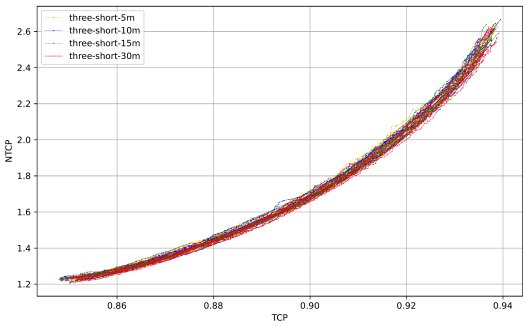

(c) Patient 3: Time limit comparison for normalization method ”short”

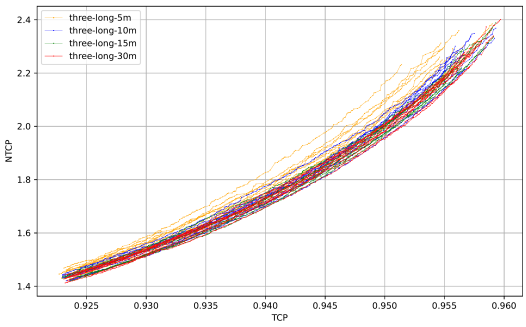

(d) Patient 4: Time limit comparison for normalization method ”long”

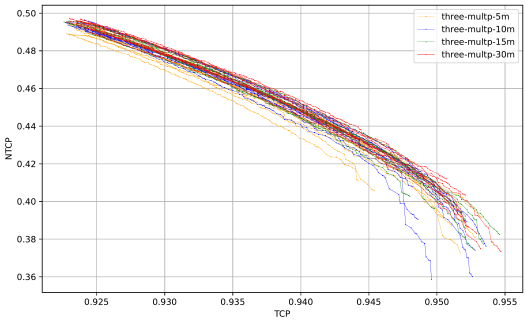

(e) Patient 4: Time limit comparison for aggregation method ”multp”

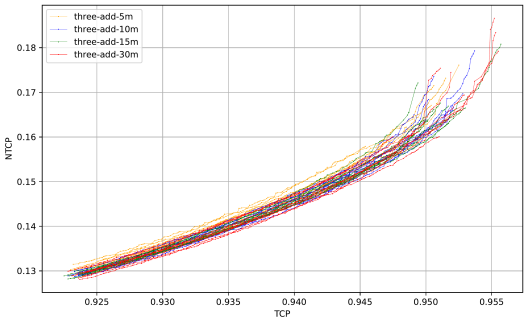

(f) Patient 4: Time limit comparison for aggregation method ”add”

Figure 4.1: The evaluated time-limits’ fronts for the used objective formulations plotted in their own TCP-NTCP space. The displayed plots show the largest time limit differences found over all tested patients and methods.

4.3 Evaluation Results

In the evaluation phase, we analyze the effectiveness of substituting the coverage and sparing objectives for the dose-response version. Moreover, we compare the various dose-response objective formulations to determine which approach produces the best results. There are vast differences in how close the dose-response optimization fronts match the ”base” results in the coverage-sparing space. The distance between the fronts depends on the patient that was optimized for and the formulation method used. A few cases that demonstrate these differences are displayed in Figure 4.2.

The few cases visualized in Figure 4.2 show the various patterns of front distances in the evaluation set. In some cases, the target-based normalization methods outperform ”multp” and ”add”, Figures 4.2a and 4.2b. Other times, this effect is the other way around, and ”multp” and ”add” provide better results, Figures 4.2c and 4.2d. And for some patients, they produce very similar fronts, Figures 4.2e and 4.2f. Additionally, the distances of the best matching fronts seem to vary independent of which method performs best. Furthermore, the aggregation methods ”multp” and ”add” generally produce comparable results, and the normalization approaches ”long” and ”short” match each other as well, albeit not as closely.

Based on these results, BRIGHT optimized using dose-response objectives cannot consistently match the protocol-based coverage and sparing results of ”base”. However, this is expected as the ”base” run type that directly optimizes these objectives should prevail over the other run types that can only improve them through secondary effects. The large difference shows a clear misalignment between the protocol-based objectives and the dose-response objectives. Additionally, these objectives are broader than the set of DVIs covered by the dose-response models; the objective-leading DVI for a patient might not be included. Instead, the results should be analyzed in the TCP-NTCP space. The dose-response objectives are designed to improve the post-treatment effects of the BRIGHT-generated plans. Plotting the fronts in this space better visualizes these improvements that are integral to determining the effectiveness of these objectives.

Figure 4.3 displays such plots. Showing the TCP-NTCP space results for all evaluated run types for each patient. Given the differing dose-response objective formulations of the evaluated run types, the plotted values are recalculated utilizing the various methods tested. To limit figure size, the results for each patient are only shown with one recalculation. With an equal representation of each method, these plots clearly show the differences in dose-response outcome improvements.

As expected, all candidate run-types improve their dose-response outcomes compared to the baseline, which does not consider the dose-response objectives. Additionally, Figure 4.3 shows that the methods generally align, as in most cases, they also outperform ”base” when displayed in a recalculated space. Specifically, given the same TCP values, we see that for ”add” the average complication probability represented improved as much as 0.01, Figures 4.3b, 4.3f, 4.3h, and 4.3p. The combined no-complication probability of ”multp” increases nearly 0.02, Figures 4.3a, 4.3c, 4.3g, and 4.3k. The target-based normalizations do not directly characterize probabilities. Instead, they represent the worst-case relative to the risk increase from a base to the planning aim dose. Utilizing the zero-dose base for ”long”, the NTCP objective reduced up to almost 0.1 worst-case relative risk increase, Figures 4.3d,

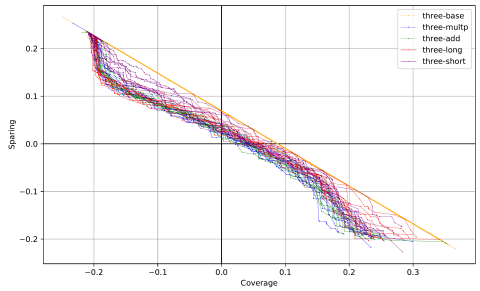

(a) Patient 11: Run type comparison with time limit ”15m”

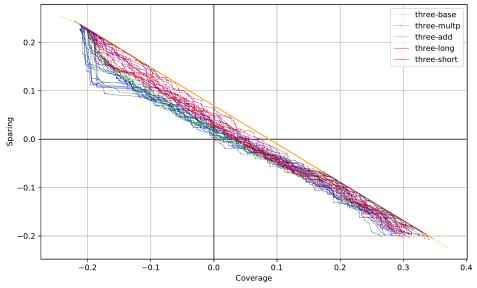

(b) Patient 16: Run type comparison with time limit ”15m”

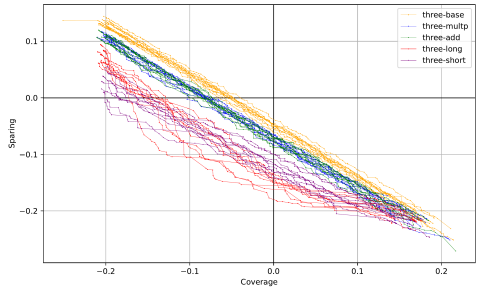

(c) Patient 8: Run type comparison with time limit ”15m”

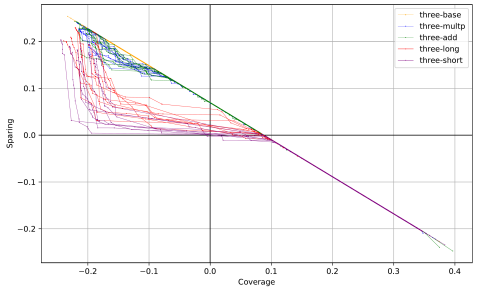

(d) Patient 9: Run type comparison with time limit ”15m”

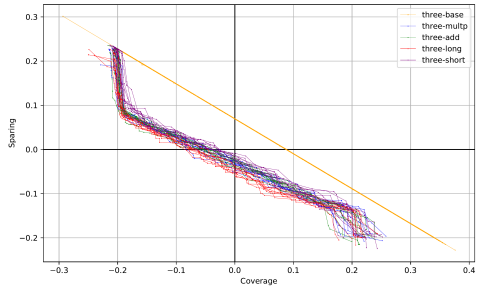

(e) Patient 13: Run type comparison with time limit ”15m”

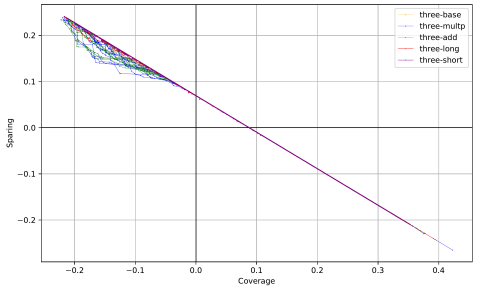

(f) Patient 15: Run type comparison with time limit ”15m”

Figure 4.2: A display of a select few evaluation phase results that demonstrate how close the dose-response optimization fronts match the ”base” results in the coverage-sparing space.

4.3e, 4.3j, 4.3l. With the steeper relative risk increases from the EBRT-base of ”short”, this is even reduced multiple tenths to the extent of 0.5, Figures 4.3i, 4.3m, 4.3n, 4.3o.

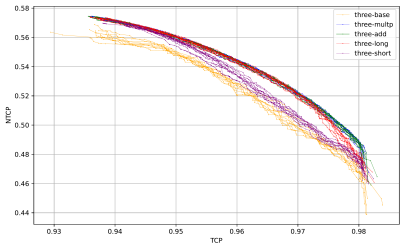

(a) Patient 1: Run type comparison with time limit ”15m” displayed using ”multp” aggregation

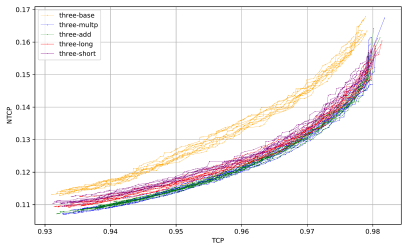

(b) Patient 2: Run type comparison with time limit ”15m” displayed using ”add” aggregation

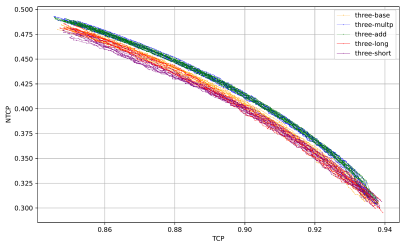

(c) Patient 3: Run type comparison with time limit ”15m” displayed using ”multp” aggregation

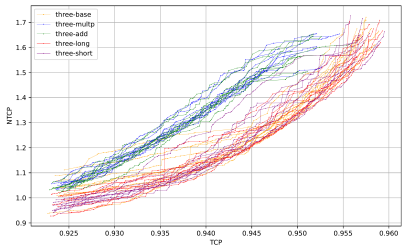

(d) Patient 4: Run type comparison with time limit ”15m” displayed using ”long” normalization

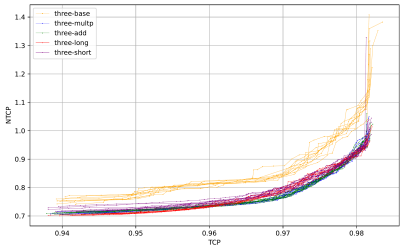

(e) Patient 5: Run type comparison with time limit ”15m” displayed using ”long” normalization

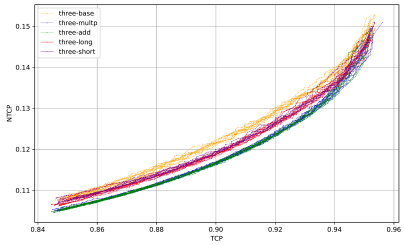

(f) Patient 6: Run type comparison with time limit ”15m” displayed using ”add” aggregation

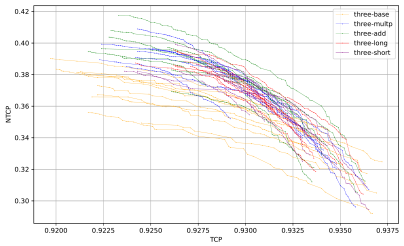

(g) Patient 7: Run type comparison with time limit ”15m” displayed using ”multp” aggregation

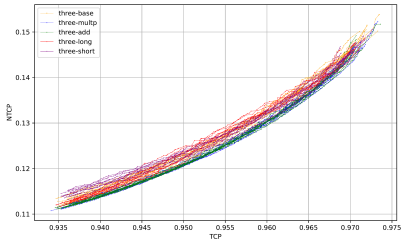

(h) Patient 8: Run type comparison with time limit ”15m” displayed using ”add” aggregation

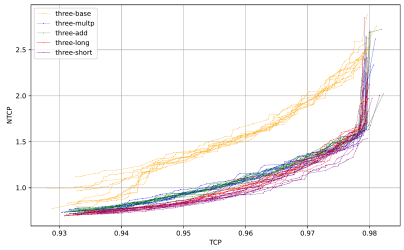

(i) Patient 9: Run type comparison with time limit ”15m” displayed using ”short” normalization

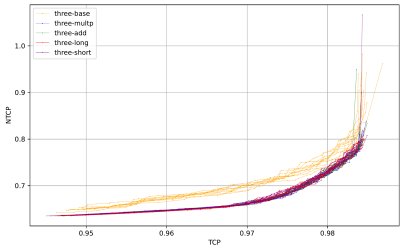

(j) Patient 10: Run type comparison with time limit ”15m” displayed using ”long” normalization

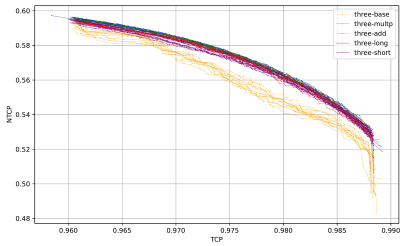

(k) Patient 11: Run type comparison with time limit ”15m” displayed using ”multp” aggregation

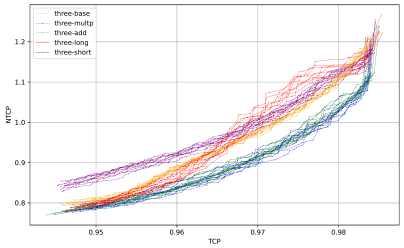

(l) Patient 12: Run type comparison with time limit ”15m” displayed using ”long” normalization

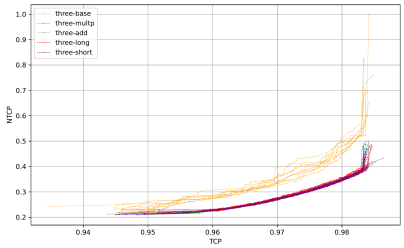

(m) Patient 13: Run type comparison with time limit ”15m” displayed using ”short” normalization

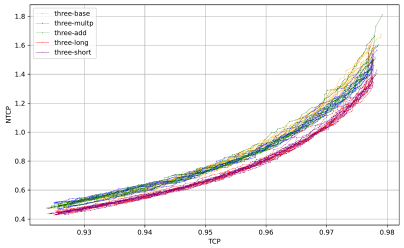

(n) Patient 14: Run type comparison with time limit ”15m” displayed using ”short” normalization

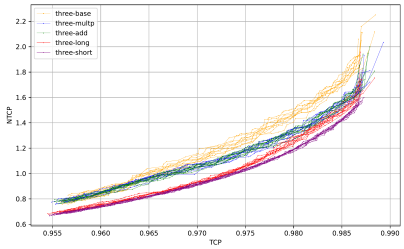

(o) Patient 15: Run type comparison with time limit ”15m” displayed using ”short” normalization

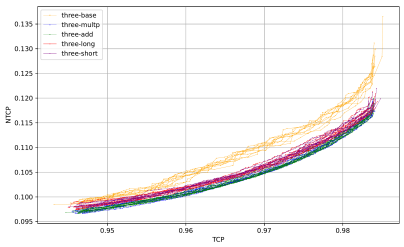

(p) Patient 16: Run type comparison with time limit ”15m” displayed using ”add” aggregation

Figure 4.3: The results fronts of the run types for each patient. Displayed in the dose-response space using recalculated values for the tested objective formulations to show their relative differences.

Comparing the dose-response run types amongst themselves, there are some peculiar differences between the target-based approaches and ”multp” and ”add”. Conceptually, they differ. The target-based variants follow BRIGHT’s worst-case-based approach, while the others combine the models prioritizing the dose-response relationships. As a result, ”long” and ”short” focus on the individual models during optimization, whereas ”multp” and ”add” focus on them together. In practice, both approaches appear beneficial in select cases. Additionally, improvements are often shared as they are generally aligned. However, given the natural dose-response-probability combinations of ”multp” and ”add”, one would expect these values to improve when optimizing the dose-response outcomes with any method. Nevertheless, for some patients, the ”long” and ”short” optimized results have either, partially, Figure 4.3h, or entirely, Figures 4.3c and 4.3l, worse average-complication and no-complication probabilities than ”base”.

Altogether, the target-based normalizations sacrifice essential aspects of the dose-response relationship and can potentially worsen the naturally interpretable probabilities of the combined results. Considering this, we will answer our second research question based solely on the ”multp” and ”add” results. Based on Figure 4.3, optimizing with these two aggregation methods greatly improves the dose-response outcomes. However, as shown in Figure 4.2, they fail to consistently match BRIGHT in protocol compliance. Moreover, the plans with the biggest dose-response outcome improvements are often not the best coverage-sparing space plans. Figure 4.4 shows two extreme cases of this effect, where the fronts lie remarkably lower than previously in Figure 4.2. To conclude, although substituting BRIGHT’s coverage and sparing objectives for dose-response objectives substantially decreases the post-treatment risks, they lack protocol compliance and, therefore, are not a suitable replacement.

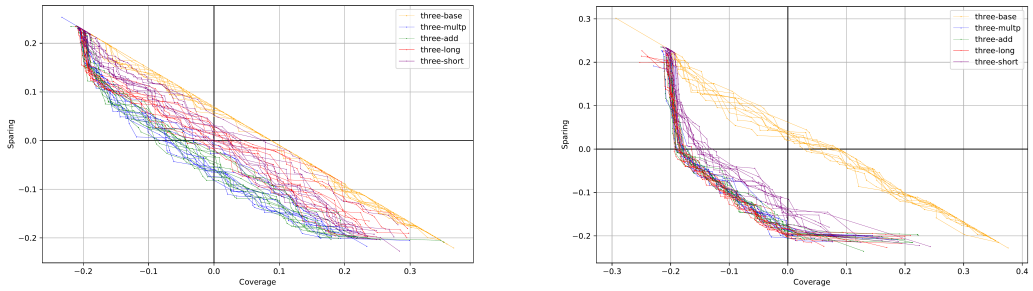

(a) Patient 11: Run type comparison with time limit ”15m”

(b) Patient 13: Run type comparison with time limit ”15m”

Figure 4.4: The coverage-sparing fronts using only filtered results that have the best TCP-NTCP values. The two cases displayed are extreme cases in which the front deteriorates considerably due to the filtering.

Based on these results demonstrating the potential and flaws of our approach, we extend our research question to explore an alternative direction. Rather than attempting to handle the complex requirements of treatment planning entirely through the dose-response objectives, we use these as extra objectives, extending this to a many-objective problem. With

this formulation, BRIGHT's current objectives should optimize the protocol's planning aims while the additional objectives steer those found plans to improve their dose-response outcomes. The extra research questions created to address these challenges are provided in Section 1.2. These questions bring new challenges that require solving. To focus our efforts effectively, we only continue with one of the aggregation methods. The "multp" and "add" methods produced comparable results and conceptually captured similar dose-response relationship qualities; both are viable options without a clear favorite. Ultimately, we continue with "add" due to its extensibility; the design scales better and allows smoother integration with other techniques.

Chapter 5

Many-Objective BRIGHT

Many-objective optimization is used to address the research questions that arose from the findings of the multi-objective experiments in Section 4.3. The current MO-RV-GOMEA used by BRIGHT can be extended to five objectives. However, as such multi-objective algorithms struggle with more than four objectives, many-objective adaptations were also explored. Given the Tchebycheff scalarizations adaptation from the literature [34], we explain the utopian point definition for usage with BRIGHT. Furthermore, we introduce two alterations to this approach designed to improve objective balance and golden-corner focus.

5.1 BRIGHT Utopian Point

The proper utopian point definition is vital for the Tchebycheff scalarizations to function optimally. After the weight vectors are assigned based on the best current objective factor, the utopian point distance determines whether solutions have improved. If the optimal objective values are known, this definition can be provided at the start of the algorithm; otherwise, it needs to be refined with the best values found during the optimization. For our BRIGHT objectives, these optimal values are known. However, as the solutions are pushed into a constrained section of the search space, this refining approach is opted for instead.

Considering the whole search space, the optimal objective values for BRIGHT are generally easy to find. Note that the following values do not necessarily match with the results plots as the objectives are redefined to minimize them during optimization. By minimizing or maximizing the entire dose applied, the coverage, sparing, TCP, and NTCP objectives can be fully optimized. The corresponding optimal values are -1 for coverage and sparing, 0 for TCP, and the NTCP value is calculated from its models at the EBRT dose; given our models, this is approximately 0.081. Conversely, the added aims value is more complicated as it comprises minimization and maximization targets. Therefore, it has no clear optimal value, though for simplicity, it could be defined as -1.

The problem with using these values for the utopian point is that BRIGHT constrains plans to the space where coverage and sparing are lower or equal to 0.2. To satisfy these constraints, the plans must strike a minimum dose-distribution balance. As a result of this minimum balance, the optimal values are unreachable. Given these restrictions, preferably,

the best reachable values would be used for the utopian point instead. However, these values are patient-specific and cannot be determined in advance. To solve this, we refine a constrained utopian point during optimization.

The constrained utopian point is initialized as the worst case for all objectives. Then, in each generation, all the individuals are checked for utopian point improvements. Importantly, these utopian point improvements are only accepted if the containing individual satisfies the constraints. This is possible as improvements are checked differently based on whether the constraints are satisfied. If unsatisfied, the distance to the constrained space is used as a quality measure; when this space is reached, the utopian point takes over. Altogether, this approach provides a meaningful, constrained utopian point in the relevant operating space that allows for more effective Tchebycheff-scalarizations guidance than the "true" utopian point of the whole space.

5.2 Normalized Tchebycheff Evaluations

In the code of [34], the solutions are normalized during the weight vector association. In the non-normalized space, the shape of the population might not fit that of the weight vector distribution. These mismatching shapes shift the relative TCH distances, which leads to relatively incorrect associations. Given the iterative association procedure, this effect amplifies with associations worsening throughout the loop as the intended weight vectors become unavailable. The range normalization applied brings the population to a similar shape as the weight vector distribution, fixing the relative TCH distances and reducing this problem. A similar issue occurs for Tchebycheff evaluations where the objective ranges influence their contributions. In this section, we explain these range disparity influences and how range normalization can be applied to fix them.

5.2.1 Objective Range Disparity

When there are large discrepancies between the objective ranges, some objectives might dominate or get ignored when evaluating TCH distances. This metric, defined as the largest weighted distance, benefits the objectives that operate in a larger range and punishes those with a smaller spread. Additionally, the uniform selection of weight vectors, designed to balance the objectives, only perpetuates the existing imbalance and cannot actively enforce balance.

With an example of the sparing and TCP objectives under the conditions of normalized assignment, uniform weight vector selection, and an over-time refined utopian point, we show this imbalance. The range for sparing and TCP is commonly around $0.3 - 0.5$ and $0.03 - 0.05$, respectively. Taking ranges 0.5 and 0.05, an order of magnitude difference, the sparing objective dominates the TCH distance, even for a weight vector that should prioritize TCP, such as $[0.1, 0.9]$. Given the normalized assignment and assuming a well-spread population, the solution associated with this weight vector should have sparing and TCP values around 10% and 90% of their respective ranges, which given minimization is $(1-0.1)\cdot 0.5 = 0.45$ and $(1-0.9)\cdot 0.05 = 0.005$, respectively. Considering these values, the TCH is $0.1 \cdot |0.45 - 0| = 0.045$ for sparing and $0.9 \cdot |0.005 - 0| = 0.0045$ for TCP. This much

larger TCH for sparing hinders TCP improvements from being accepted as they need to be much bigger. The range disparity effectively disables TCP improvements in this example despite the chosen weight factor that should have favored this objective.

5.2.2 Range Normalized Evaluations

Applying the range normalization from the weight vector association to the Tchebycheff evaluations removes all range disparity. Each objective range is normalized to $[0,1]$, where 0 represents the best value and 1 is the worst. Recalculating the TCH for the example in Section 5.2.1, the normalized versions of the objective values are around $(1-0.1)\cdot 1 = 0.9$ for sparing and $(1-0.9)\cdot 1 = 0.1$ for TCP. Using these normalized values, the new TCH is $0.1\cdot|0.9-0| = 0.09$ and $0.9\cdot|0.1-0| = 0.09$, for sparing and TCP, respectively. Using the range normalized evaluations, the individual dimensions' TCH distances converge to the same value, allowing improvements for all objectives. Additionally, from this point, the TCH influence of relatively similar improvements is properly scaled based on the associated weight vector.

The use of range-normalized objective values also requires a redefinition of the utopian point. One approach is to normalize the previously used refining utopian point during evaluation. However, the fixed normalized range allows for a more elegant yet balanced definition. Setting the utopian point at -0.1 for all dimensions provides a consistent improvement direction with an equal focus on all objectives. Additionally, as it is an extension of the reference point at values 0, it improves the effectiveness of the weight vectors. Adding a constant to each distance has a relatively larger effect on the smaller distances that coincide with the focus of the weight vector, thus increasing its effectiveness. Combined, this utopian point definition and the range normalization ensure that objectives are balanced and effectively explored despite any range disparities.

5.3 Guided Weight Vector Spread

One of BRIGHT's advantages is that it provides a front of plans, considering various objective trade-offs. However, the main area of interest is the golden corner, where all the coverage and sparing planning aims are met. When extending BRIGHT to the five-dimensional space, this area only represents a fractional part of the search space. The constraints in place slightly reduce this effect by guiding the solutions toward a smaller space surrounding the golden corner. Nevertheless, the selection pressure remains overly diffused in the many-objective space.

To increase the golden-corner focus, we modify the weight vector spread to increase the vector density around this area. By increasing the amount of vectors with an even objective trade-off, more improvements in this direction should be accepted. To achieve this, we introduce quadratic guidance scaling for the weight vector spread. During the greedy scattered subset selection to select the weight vectors, a scalar increases the inter-vector distances with the effect strongest around the center and fading outwards. These scalars quadratically scale based on a vector's distance to the center weight vector ($\frac{1}{n}$ for all n objectives) as formulated in Equation 5.1. Here, gc_v is the guidance scalar for vector v, v_c

is the center weight vector, and $d(v, v_c)$ is the Euclidean distance between vectors v and v_c.

$$gc_v = \frac{1}{\max(0.1,\ d(v, v_c))} \tag{5.1}$$

The inclusion of the 0.1 minimum scalar fraction in the denominator is designed to limit the scaling effect. The limited scaling prevents extreme values excessively increasing center distances and limits the guidance to a set radius inside which the selection is uniform again. Additionally, the quadratic scaling always increases the inter-vector distances as it is always greater than 1. The largest possible distances to the center weight vector are at the corners weight-vector space; this maximum distance, shown in Equation 5.2, is always smaller than 1 given $n \geq 1$ dimensions.

$$\sqrt{(1-\frac{1}{n})^2 \ +\ ((n-1)\ \cdot\ (\frac{1}{n})^2)} \tag{5.2}$$

With this quadratic scaling, the weight vector spread is reshaped from a uniform spread to a center-focused distribution. Figure 5.1 illustrates this change, depicting the uniform selection on the left and the guided approach on the right. The figure shows that the guided spread is much denser populated around the center weight vector, which should align with the golden-corner's direction. Furthermore, vector density reduces moving outwards from the center yet keeps some focus here for exploration.

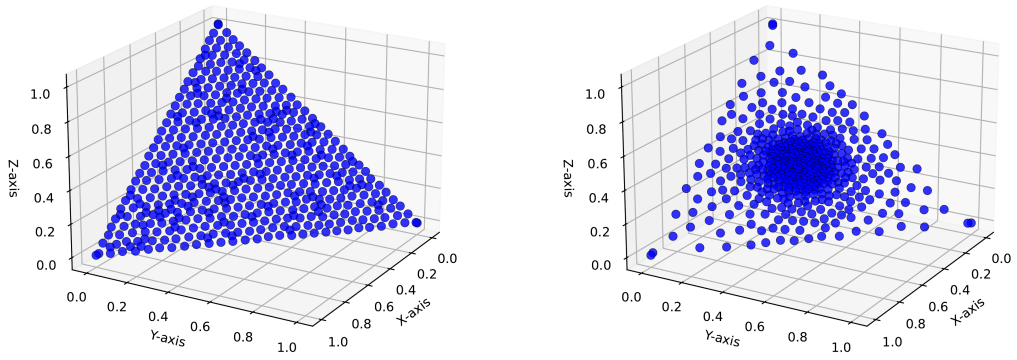

(a) The uniform weight vector distribution.

(b) The guided weight vector distribution.

Figure 5.1: The weight vector spreads for the uniform and guided selection approaches visualized for a three-dimensional example.

Chapter 6

Experiments for Many-Objective BRIGHT

These experiments aim to answer the third, fourth, and fifth research questions. Various setups using the techniques detailed in Chapter 5 are evaluated on their ability to improve the dose-response outcomes for plans protocol with good protocol compliance. Based on these improvements, we determine the impact of including dose-response objectives in treatment planning optimization.

6.1 Experimental Setup

The setup for these experiments is similar to that of the first experiments described in Section 4.1. First, the dataset and hardware specifications remain unchanged. Second, the experimental design is only slightly modified. To analyze the evaluation phase results, we compare the achieved TCP and NTCP values displayed over the coverage-sparing space. Third, new run types and parameters are evaluated and tuned during the experimental process. To avoid redundancy, this section covers only the run types and parameters used. Details on the dataset and hardware specifications can be found in Section 4.1.1, and the experimental design in Section 4.1.2.

One addition to the experimental setup is the inclusion of a mini baseline experiment. This mini-experiment applies the techniques developed for many-objective BRIGHT to the three-dimensional baseline. The different many-objective techniques are designed for the extra complications of the higher dimensional space. However, some additions, such as the extra golden-corner focus, might also benefit the multi-objective definition. This mini-experiment evaluates this. Additionally, it also serves as a reference for the effectiveness of the designed techniques on simpler problems. This extra experiment will not have a tuning phase. Instead, they continue with the fixed and tuned parameters from the experiments in Chapter 4 and are run with the run types explained below.

6.1.1 Run Types and Parameters

The extension of BRIGHT with the dose-response objectives should improve the TCP and NTCP values of good coverage and sparing plans. As before, our baseline run type is the regular 3-dimensional BRIGHT; we define this as ”base”. Given the complexity of the many-objective space, we apply several versions of the MO-RV-GOMEA algorithm to explore it. This gives us four candidate run types. First, run type ”normal” is the regular MO-RV-GOMEA employed by BRIGHT. Second, ”default” is the standard Tchebycheff scalarizations variant from the literature [34] using our utopian point definition described in Section 5.1. Third, we define ”norm” as our altered Tchebycheff approach using normalized evaluations as described in Section 5.2. Fourth, run type ”guide” uses the extension of ”norm” with the guided weight vector spread detailed in Section 5.3. In total, we evaluate five run-types: ”base”, ”normal”, ”default”, ”norm”, and ”guide”.

Similar to the previous experiments, we need to address the same five impactful parameters and choose which values to use. The number of DCPs remains unchanged at 100.000 as this amount still provides sufficient accuracy without excessive overhead. The population size, the number of clusters, and the archive size should scale with the increased dimensionality of our problem. Before, these values were well-defined for BRIGHT. For the five-objectives variant, this is only the case for the population size, which becomes 480. Instead, we need to tune the number of clusters and the archive size. Scaling the values defined for the lower-dimensional variants of BRIGHT, we decided to test 20 and 30 clusters and an elitist archive size of 5.000 and 10.000. The time limits also need to be re-tuned as the many-objective space and evaluated run types operate differently than the three-objective optimization. Given the available hardware, while keeping it relatable to the clinical pipeline, the same values of 5, 10, 15, and 30 minutes are tested. To summarize, the DCPs and population size are fixed at 100.000 and 480, respectively. The other parameters are all tuned with clusters ”20c” and ”30c”, elitist archive sizes ”5ks” and ”10ks”, and time limits ”5m”, ”10m”, ”15m”, and ”30m”.

6.2 Multi-Objective Application Results

As explained in the experimental setup above, this mini-experiment demonstrates the effectiveness of the many-objective techniques in the simpler multi-objective setting. Figure 6.1 shows the results of this extra experiment. These plots show little to no differences between the run types. For many patients, such as shown in Figures 6.1b, 6.1e, and 6.1f, the fronts of all runs and run types overlap, forming almost a single line. Regular BRIGHT can already find this front; the evaluated techniques do not improve upon this. In the other cases where there is some variance between the runs, like Figures 6.1a, 6.1c, and 6.1d, they also achieve similar results. The runs for each setting are generally spread across the same range with near-equal consistency.

Overall, the application of many-objective techniques in this setting generally matches the performance of regular BRIGHT but cannot improve on it. The three-dimensional definition of BRIGHT is simple enough that the use of these adaptions designed for complicated cases is not required. For most patients, the convergence of all runs onto a single

line suggests that regular BRIGHT already extracts the maximum it can given the current objective definitions. In the cases where the results vary more, the adaptations still fail to improve. If anything, one could argue that they are even marginally worse. The patients likely require complex inner-objective balances. However, with minimal objectives, these complexities might not extend to the objective relations that the adaptations target. Ultimately, these techniques do not add any value in the multi-objective setting. It can match the performance of regular BRIGHT. However, it is not worth the added complexity.

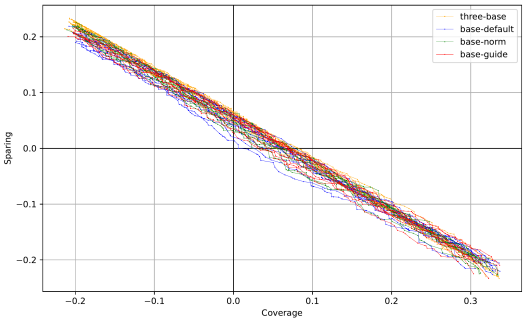

(a) Run type comparison for patient 2

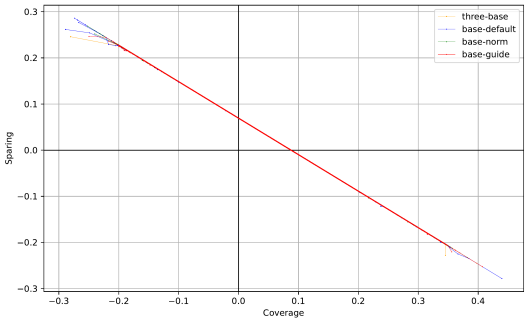

(b) Run type comparison for patient 5

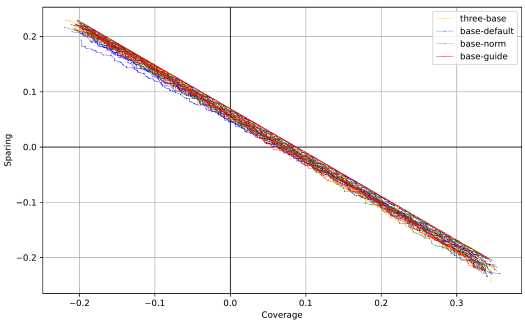

(c) Run type comparison for patient 6

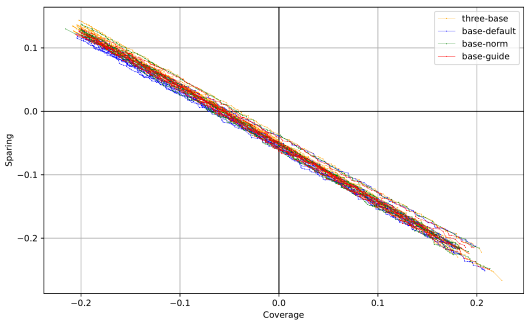

(d) Run type comparison for patient 8

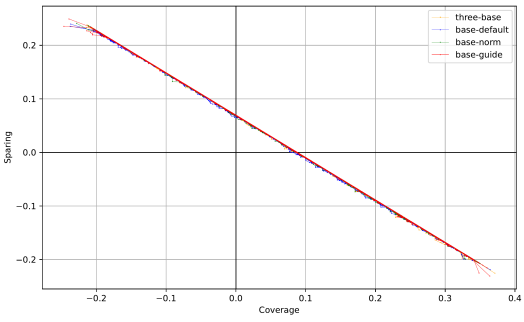

(e) Run type comparison for patient 12

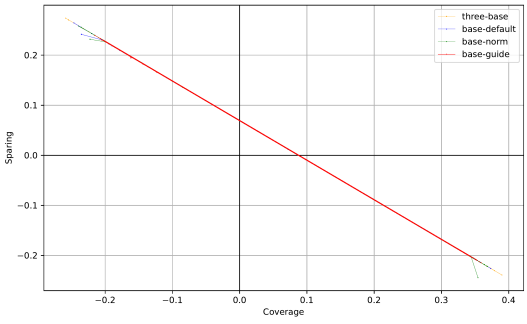

(f) Run type comparison for patient 15

Figure 6.1: The results of applying the designed many-objective techniques to the three-dimensional definition of BRIGHT. The results were run with a time limit of 15 minutes and are displayed in the coverage-sparing space.

6.3 Tuning Results

Given the three tunable parameters, each run type was run for all possible combinations. With these results, the differences for each set of parameter options are analyzed and set in order from least to most dependent on the other parameters. This dependency ordering dictates that the elitist archive size be set first, followed by the number of clusters, and finally, the time limit. Considering the goal of the many-objective runs to achieve coverage-sparing results similar to ”base” while improving TCP-NTCP outcomes, both spaces are used to determine the best parameter values.

When inspecting the results, we found our ”harder” case, patient 4, to be more complicated than previously anticipated. Due to this, BRIGHT’s optimization strategy of splitting the time limit into three parts negatively impacts the results. To achieve fast initial improvements but end with precision, a quarter, half, and all of the DCPs are used for the three parts, resampling them between each split. The complexity of this case leads to many plans falling outside the constraints after these resamplings. This resets most progress, notably reducing the effective time used for optimization. As a result, the achieved fronts are slightly inconsistent; we consider this when tuning the parameters.

6.3.1 Elitist Archive Size

The evaluated elitist archive sizes of 5.000 and 10.000 showed little to no distinctions for most of the 128 patients, settings, and many-objective run-type combinations. Figure 6.2 visualizes six cases representing the differences found over all the results. More cases are shown in Appendix A. A few settings show minor improvement for ”5ks” over ”10ks”, Figures 6.2c and 6.2d. Larger differences are also observed for patient 4, Figure 6.2f, though this is largely due to the inconsistency with this patient. However, the vast majority of cases show no distinctions, as demonstrated by Figures 6.2a, 6.2b, and 6.2e.

Based on these observations, we decided to continue with an elitist archive size of 5.000. Both options achieved comparable results. However, the minor differences were more commonly favoring ”5ks”, and a smaller archive size reduces computation requirements. We suspect both archive sizes can sufficiently cover our five-dimensional space, consequently causing the observed similar results. The slight distinction could potentially be explained by the discretization algorithm used to shrink the elitist archive when it exceeds the maximum size. The larger size of ”10ks” reduces the occurrence of this event. However, this comes at the cost of increased computational scaling, which grows quadratically with the archive size. With the elitist archive size set as ”5ks”, we continue to tune the other parameters with this value fixed.

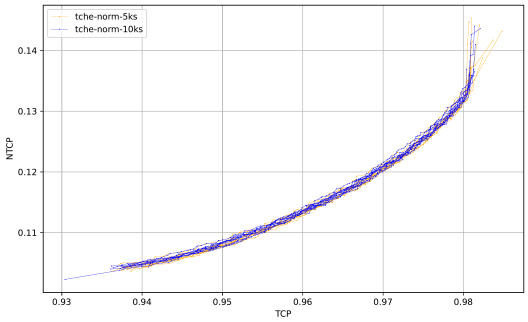

(a) Patient 1: Elitist archive size comparison for run type ”norm”, clusters ”20c”, and time limit ”10m” in TCP-NTCP space

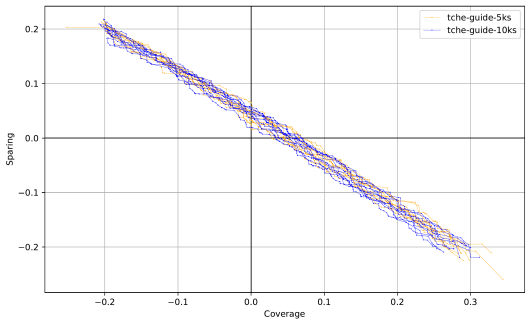

(b) Patient 2: Elitist archive size comparison for run type ”guide”, clusters ”30c”, and time limit ”15m” in coverage-sparing space

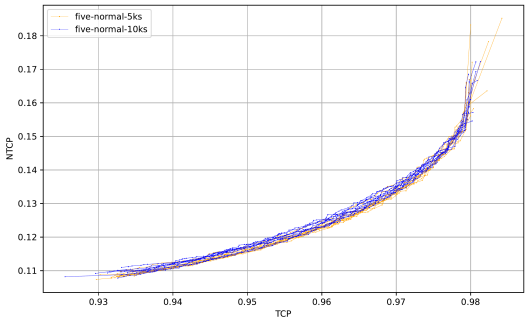

(c) Patient 2: Elitist archive size comparison for run type ”normal”, clusters ”30c”, and time limit ”15m” in TCP-NTCP space

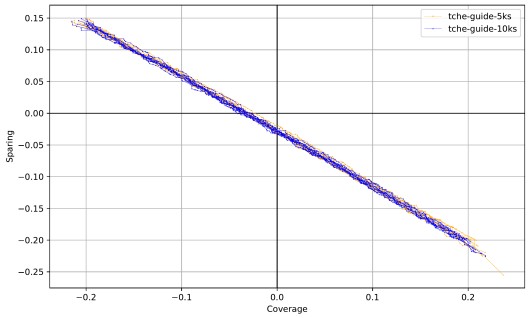

(d) Patient 3: Elitist archive size comparison for run type ”guide”, clusters ”20c”, and time limit ”30m” in coverage-sparing space

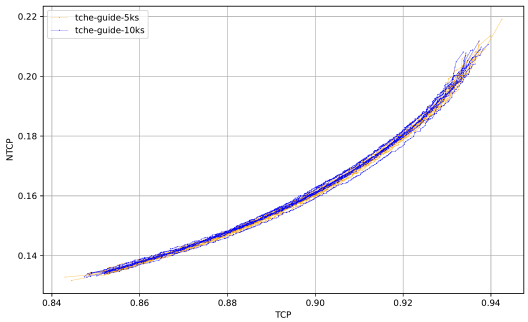

(e) Patient 3: Elitist archive size comparison for run type ”guide”, clusters ”30c”, and time limit ”30m” in TCP-NTCP space

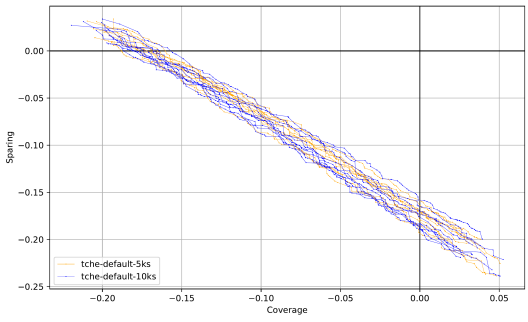

(f) Patient 4: Elitist archive size comparison for run type ”default”, clusters ”20c”, and time limit ”5m” in coverage-sparing space

Figure 6.2: The evaluated elitist archive sizes’ fronts plotted in the coverage-sparing and TCP-NTCP space. These plots visualize the difference in optimization potential between the archive-size options.

6.3.2 Number of Clusters

The number of clusters also did not particularly impact the resulting fronts. Comparisons for some of these fronts are displayed in Figure 6.3, with additional cases in Appendix A. In most cases, the fronts overlap, Figures 6.3b, 6.3c, and 6.3d. Slightly larger differences are visible in Figures 6.3a and 6.3e, favoring "20c" and "30c", respectively. However, these improvements do not consistently carry over to similar settings and spaces, as demonstrated by the varying results for patient 2. Lastly, patient 4 again shows the biggest differences for both options as shown in Figure 6.3f.

Although the results are generally consistent, small trends with patients favoring either option can be observed. Patient 2 remains indifferent towards the number of clusters, but for patients 1 and 3, the front overlaps are frequently slightly shifted with common beneficiaries "20c" and "30c", respectively. Additionally, patient 4 also faintly favors "20c", although the results vary too much to make any definitive claims.

We speculate that the similar results and patient-specific preferences depend on the objective correlation. The five objectives comprise DVIs targeting tumor coverage or OAR sparing. These two types of DVIs are dependent, with one improving as the other decreases and vice versa. Additionally, the use of these DVIs across multiple objectives further increases objective correlation. This correlation reduces the five-dimensional search-space spread, which affects the number of clusters needed. Due to this, 20 clusters are likely sufficient to cover our heavily-correlated space. Notably, part of this correlation depends on the main-contributing DVIs of the sparing objective and whether these match the DVIs used for NTCP. Differences in the contributions of these DVIs could influence the optimal number of clusters for a patient and cause the favoritism observed.

Ultimately, we decided to continue with "20c". Albeit marginally, it was the better value for most tuning runs. Additionally, the population over clusters ratio $\frac{480}{20} = 24$ matches that of the multi-objective approach $\frac{288}{12} = 24$. We fix this value of 20 clusters to determine the value for the last tuned parameter and the evaluation runs.

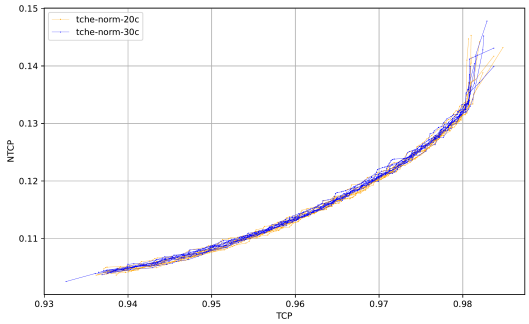

(a) Patient 1: Number of clusters comparison for run type ”norm”, archive size ”5ks”, and time limit ”10m” in TCP-NTCP space

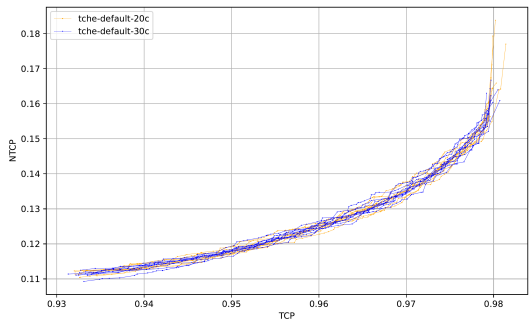

(b) Patient 2: Number of clusters comparison for run type ”default”, archive size ”5ks”, and time limit ”15m” in TCP-NTCP space

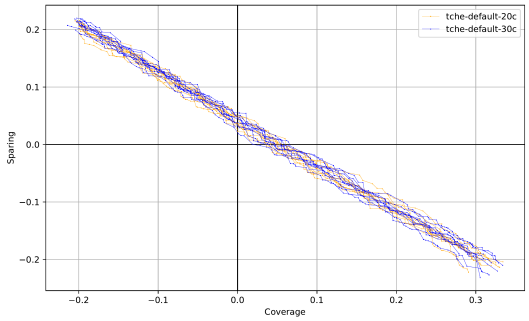

(c) Patient 2: Number of clusters comparison for run type ”default”, archive size ”5ks”, and time limit ”30m” in coverage-sparing space

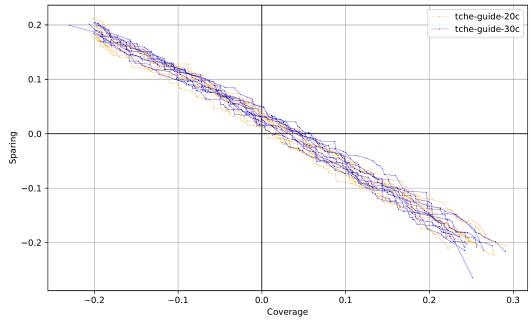

(d) Patient 2: Number of clusters comparison for run type ”guide”, archive size ”5ks”, and time limit ”5m” in coverage-sparing space

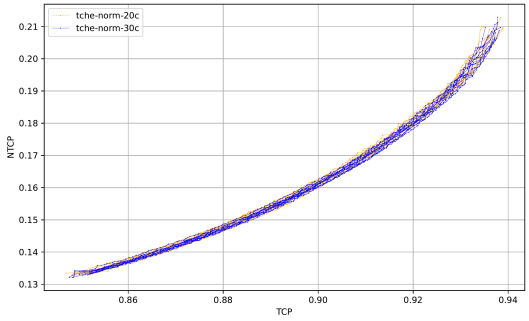

(e) Patient 3: Number of clusters comparison for run type ”norm”, archive size ”5ks”, and time limit ”30m” in TCP-NTCP space

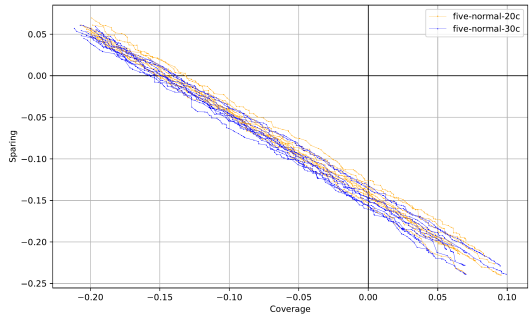

(f) Patient 4: Number of clusters comparison for run type ”normal”, archive size ”5ks”, and time limit ”15m” in coverage-sparing space

Figure 6.3: The evaluated number of clusters’ fronts plotted in the coverage-sparing and TCP-NTCP space. These plots visualize the difference in optimization potential between the clusters’ options.

6.3.3 Time Limit

The optimization time limit is an important parameter trading off time usage for solution quality. The relative front improvements in trade for our evaluated time limits are visualized in Figure 6.4 using a few select cases. More cases are shown in Appendix A. For a simpler case, like Figure 6.4a, shorter times such as ”10m” are sufficient to find the best plans possible given the used algorithm. However, most cases are more complicated and show clearer improvements for the increased time limits. Some of these improvements continue for all tested limits, such as in Figures 6.4b, 6.4c, and 6.4f. Others plateau after an adequate time limit is reached, commonly after ”15m” as demonstrated by Figures 6.4d and 6.4e. Ideally, our time limit aligns with the moment the improvements plateau or are small enough to be considered negligible.

Based on Figure 6.4, the longer options of ”15m” and ”30m” fit this description best. These options show clear improvements on their predecessors. Moreover, ”30m” even exceeds ”15m” in some cases, albeit less consistently. However, comparing these longer options, ”15m” would be the preferred choice for the evaluation runs considering the performance differences and clinical relevance of these limits. Over the entire tuning set, the improvements of using ”30m” are, although apparent, marginal given the increased time requirements. Additionally, the ”30m” time limit was designed to show whether improvements plateaued, not as a realistic option since it does not fit in the clinical pipeline. One could even argue that ”15m” is a stretch regarding clinical relevance. However, given the essence of a higher limit in the scientific context of researching the viability of this method, we justify it.

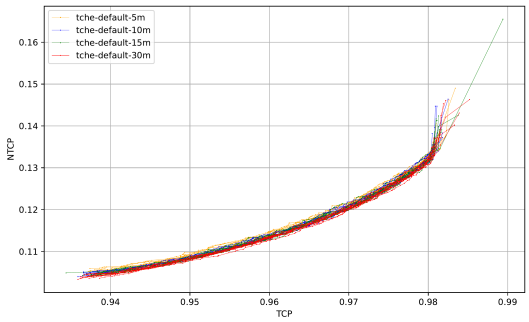

(a) Patient 1: Time limit comparison for run type ”default”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

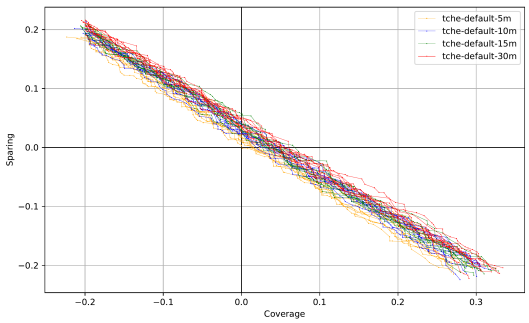

(b) Patient 2: Time limit comparison for run type ”default”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

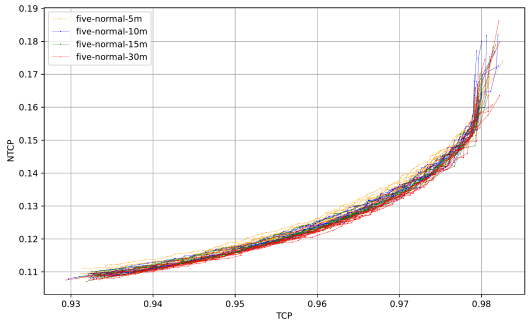

(c) Patient 2: Time limit comparison for run type ”normal”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

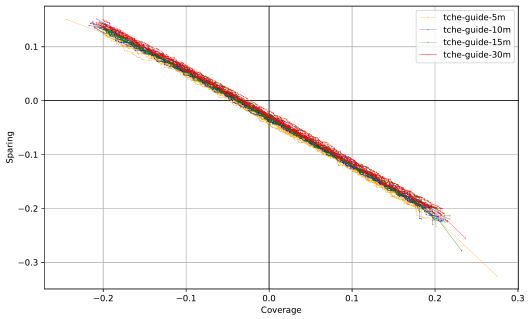

(d) Patient 3: Time limit comparison for run type ”guide”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

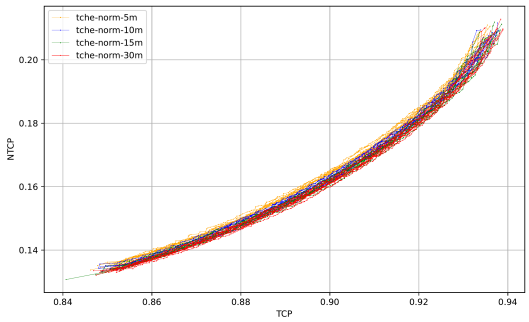

(e) Patient 3: Time limit comparison for run type ”norm”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

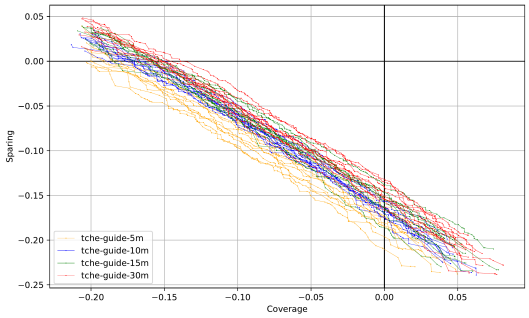

(f) Patient 4: Time limit comparison for run type ”guide”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

Figure 6.4: The evaluated time limits’ fronts plotted in the coverage-sparing and TCP-NTCP space. These plots visualize the difference in optimization potential between the time-limit options.

6.4 Evaluation Results

Using the evaluation set, we analyze the impact of optimizing dose-response models in the many-objective formulation. Furthermore, we compare the performance of the various approaches designed to assist this complex optimization. To start, Figure 6.5 shows how combining all objectives allows the candidate run types to achieve results similar to ”base” in the coverage-sparing space while also finding plans with better dose-response outcomes. Additionally, the candidates appear to generate comparable fronts. Only a rare few cases show marginal differences between them, as displayed in Figure 6.5b.

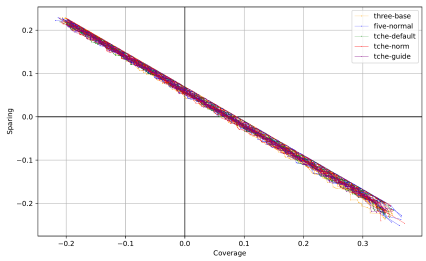

(a) Patient 6: Run type comparison in coverage-sparing space

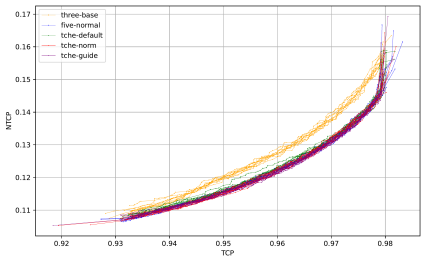

(b) Patient 9: Run type comparison in TCP-NTCP space

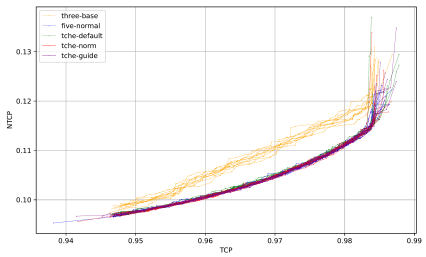

(c) Patient 10: Run type comparison in TCP-NTCP space

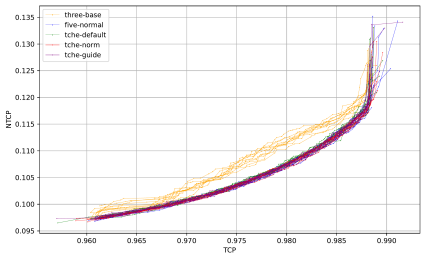

(d) Patient 11: Run type comparison in TCP-NTCP space

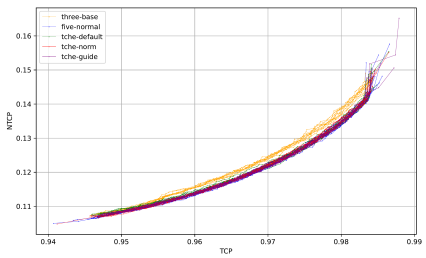

(e) Patient 12: Run type comparison in TCP-NTCP space

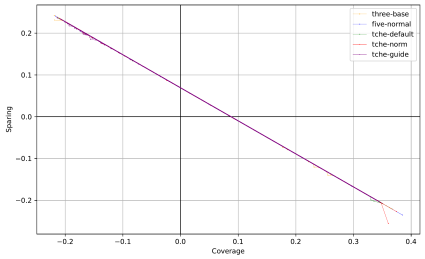

(f) Patient 14: Run type comparison in coverage-sparing space

Figure 6.5: Evaluation-run plots that display the difference between fronts of the candidate run types and baseline in the coverage-sparing and TCP-NTCP space. All results used settings clusters ”20c”, archive size ”5ks” and time limit ”15m”.

Given that the many-objective runs show admirable results when observing the objective types independently, we are interested in how their combined results compare to ”base”. To analyze these trade-offs, we plot the results in the coverage-sparing space with colors that represent the TCP and NTCP values or differences. Figure 6.6 visualizes these color-represented dose-response objective values and -differences used to evaluate the results for a single patient. Using the same TCP and NTCP color scales, Figures 6.6a and 6.6b visualize the colors for our baseline ”three-base” and Figures 6.6c and 6.6d shows the same for a candidate approach ”tche-guide”. Additionally, Figures 6.6e and 6.6f focus on the dose-response improvements of our candidate over the baseline in the golden corner.

These color-visualized insights for patient 9 demonstrate that the best dose-response values gradually improve as the plans move toward the optimal found coverage and sparing results. However, given the gradient-based color assignments, the exact values are hard to extract from the color plots. Instead, the heatmaps highlight the differences found in the golden corner. These differences show the direct improvements possible at the front and those obtainable through sacrificing sparing while remaining at the same coverage. An initial assessment of this single patient already shows various improvements possible at the front and through sacrifices.

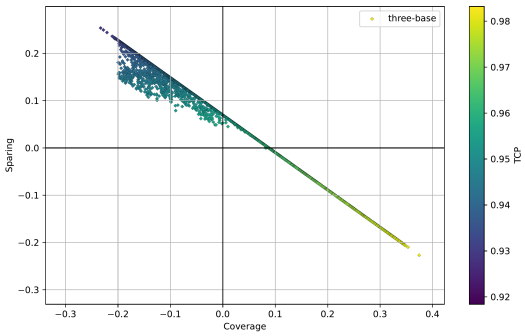

(a) TCP overview for run type "base"

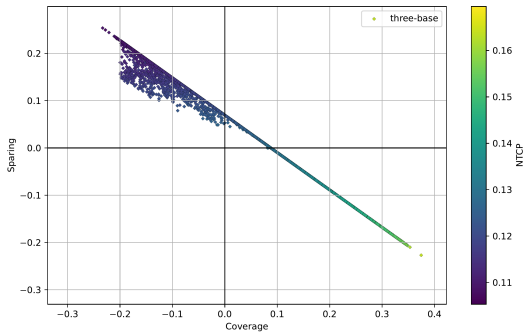

(b) NTCP overview for run type "base"

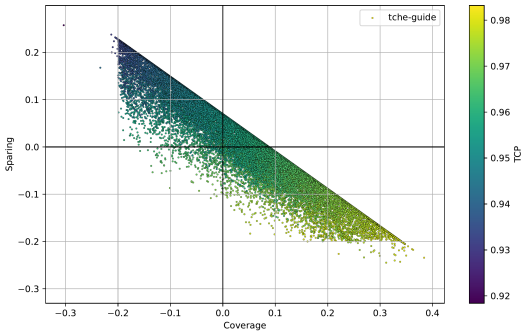

(c) TCP overview for run type "guide"

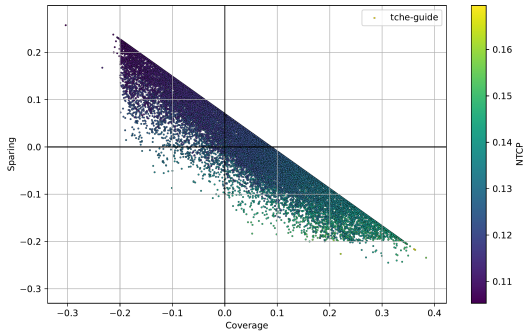

(d) NTCP overview for run type "guide"

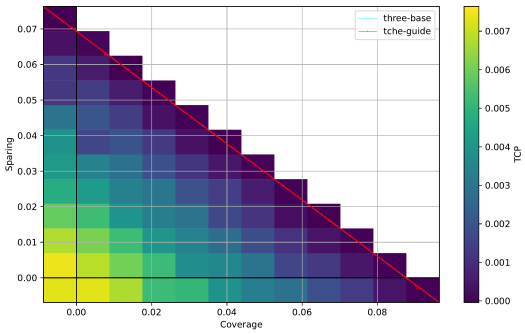

(e) Golden-corner TCP differences between run types "base" and "guide"

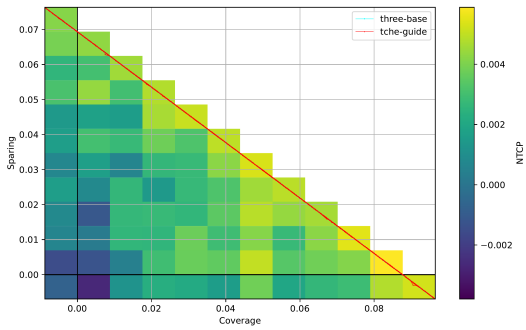

(f) Golden-corner NTCP differences between run types "base" and "guide"

Figure 6.6: The color-represented TCP and NTCP values and differences for run types "base" and "guide". The displayed results visualize these dose-response qualities for the plans for patient 9 in the coverage-sparing space.

Notably, each of the candidate run types shows different levels of improvement compared to the baseline. Figures 6.7 and 6.8 show the performances of the tested run-types for two patients with different optimal found dose-response regions. For all patients, the TCP improvement pattern observed was the same as shown in the figures. Improvements only really show when sacrificing sparing. Additionally, the specific run type appears only

marginally impactful on the obtainable TCP improvements. For NTCP, there are two main patterns. The biggest improvements are either found at the front, such as in Figure 6.7, or require balanced sacrifice, as in Figure 6.8. Commonly, improvements up to 0.006 of average NTCP model risk are possible on or near the front. Furthermore, larger performance variations between the run types are observed for this more complicated objective representing multiple models.

Table 6.1 compares the improvements obtained by the run types. Following the heatmap pattern, we analyze the NTCP improvements on the front and TCP and NTCP improvements achievable by sacrificing some sparing. For each category, we measure the best and average cell improvements. Similar to the plots, these values suggest that the run types are most influential for the NTCP results. We analyze the results using the Wilcoxon signed-rank test with a p-value smaller than or equal to 0.05 indicating significance. The NTCP results for ”guide” are statistically significantly better than the others, while there is no clear significant best option for TCP. Overall, run type ”guide” showed the most consistent and biggest improvements. Next, run types ”normal” and ”norm” achieved similar results, with the best option usually depending on patient-specific characteristics. Lastly, ”default” improved the least overall. Despite finding the largest TCP improvements, it is already matched by the others for the average TCP. Furthermore, it performed statistically significantly worse compared to all other run types for the NTCP-objective metrics. Based on the results, we continue with run type ”guide” as our final candidate. Statistical tests proved this method to either significantly improve or match the best-found results for NTCP and TCP, respectively.

Table 6.1: The dose-response improvements obtained by the tested candidate run types compared to the baseline in and around the golden corner. Bold values are statistically significantly better than all others with a p-value smaller than or equal to 0.05 based on the Wilcoxon signed-rank test. Values marked with * indicate a p-value smaller than 0.001.

Metric	**Normal**	**Default**	**Norm**	**Guide**
Best TCP improvement	0.00597 ± 0.00185	0.00639 ± 0.00212	0.00578 ± 0.00182	0.00591 ± 0.00138
Average TCP improvement	0.00166 ± 0.00067	0.00166 ± 0.00064	0.00165 ± 0.00064	0.00168 ± 0.00062
Best NTCP front improvement	0.00392 ± 0.00174	0.00337 ± 0.00171	0.00389 ± 0.00176	$\mathbf{0.00417} \pm 0.00166$
Average NTCP front improvement	0.00311 ± 0.00176	0.00264 ± 0.00191	0.00315 ± 0.00178	$\mathbf{0.00352}^* \pm 0.00169$
Best NTCP non-front improvement	0.00472 ± 0.00125	0.00434 ± 0.00122	0.00479 ± 0.00125	$\mathbf{0.00532}^* \pm 0.00142$
Average NTCP non-front improvement	0.00243 ± 0.00109	0.00185 ± 0.00102	0.00251 ± 0.00107	$\mathbf{0.00303}^* \pm 0.00122$

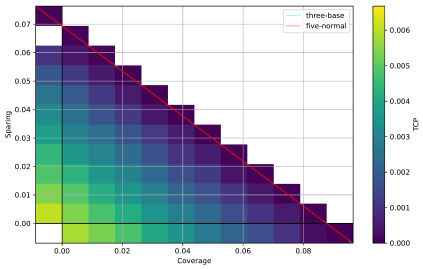

(a) Golden-corner TCP differences between run types ”base” and ”normal”

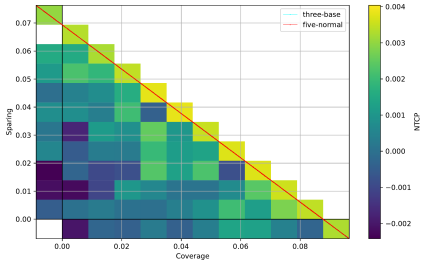

(b) Golden-corner NTCP differences between run types ”base” and ”normal”

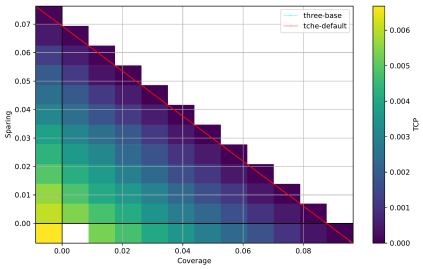

(c) Golden-corner TCP differences between run types ”base” and ”default”

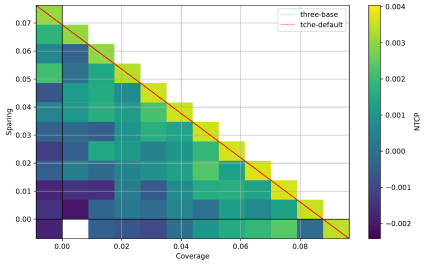

(d) Golden-corner NTCP differences between run types ”base” and ”default”

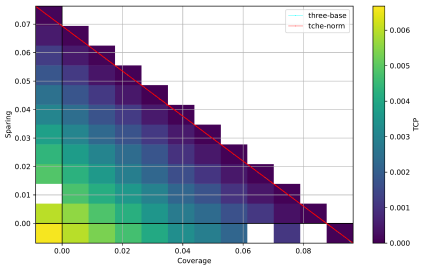

(e) Golden-corner TCP differences between run types ”base” and ”norm”

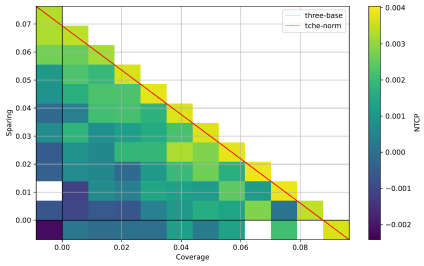

(f) Golden-corner NTCP differences between run types ”base” and ”norm”

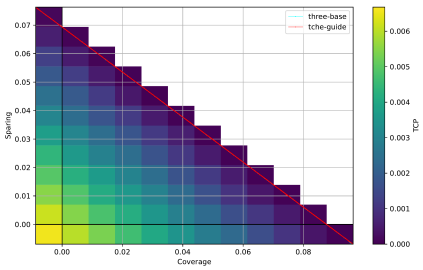

(g) Golden-corner TCP differences between run types ”base” and ”guide”

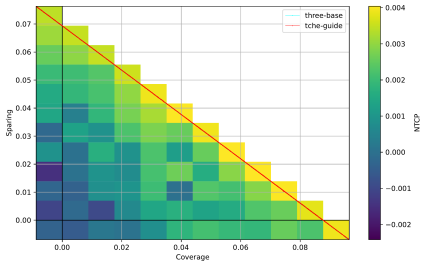

(h) Golden-corner NTCP differences between run types ”base” and ”guide”

Figure 6.7: The color-represented TCP and NTCP differences for candidate run types ”normal”, ”default”, ”norm”, and ”guide” compared to baseline ”base”. These plots visualize the dose-response qualities for the plans of patient 16 in the coverage-sparing space.

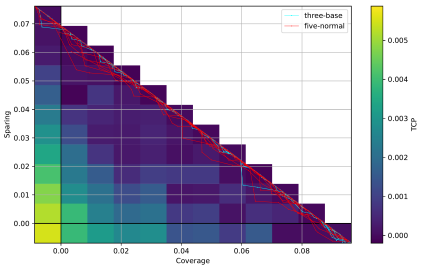

(a) Golden-corner TCP differences between run types ”base” and ”normal”

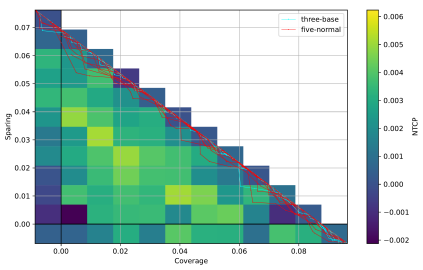

(b) Golden-corner NTCP differences between run types ”base” and ”normal”

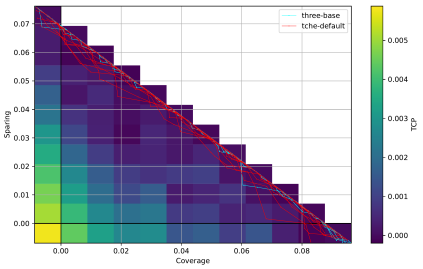

(c) Golden-corner TCP differences between run types ”base” and ”default”

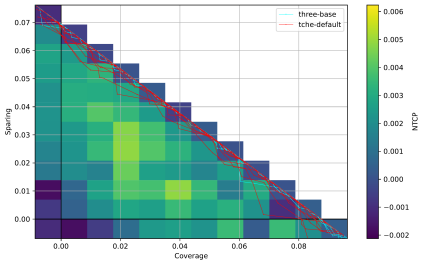

(d) Golden-corner NTCP differences between run types ”base” and ”default”

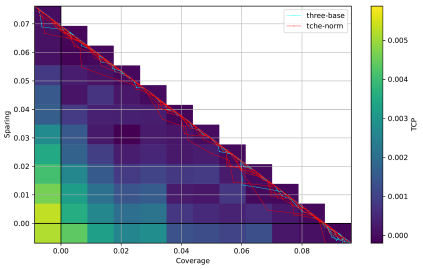

(e) Golden-corner TCP differences between run types ”base” and ”norm”

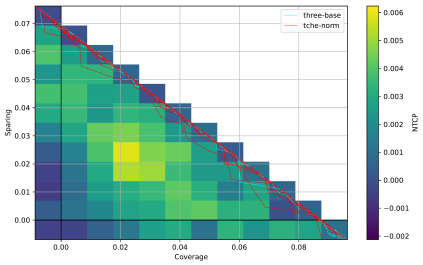

(f) Golden-corner NTCP differences between run types ”base” and ”norm”

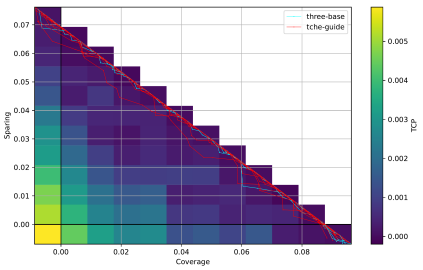

(g) Golden-corner TCP differences between run types ”base” and ”guide”

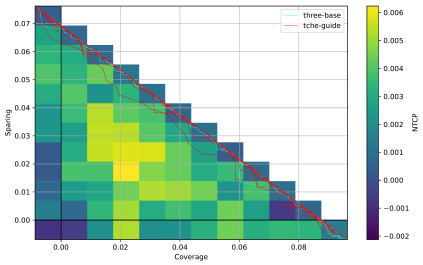

(h) Golden-corner NTCP differences between run types ”base” and ”guide”

Figure 6.8: The color-represented TCP and NTCP differences for candidate run types ”normal”, ”default”, ”norm”, and ”guide” compared to baseline ”base”. These plots visualize the dose-response qualities for the plans of patient 12 in the coverage-sparing space.

With our final candidate "guide", we have observed improvements around 0.005 TCP with sparing sacrifices, 0.004 NTCP at the front, and 0.005 NTCP with sparing sacrifices, as shown in Table 6.1. Additional visualizations of these improvements achieved by "guide" for the not yet shown patients are available in Appendix A. However, besides these numerical improvements, it is also imperative to analyze how the inclusion of these dose-response objectives influenced the plans themselves. BRIGHT has developed an app in which the optimized plans can be examined and compared. Figures 6.9 and 6.10 show some distinct differences observed between the dose distributions of baseline "base" and candidate "guide" plans given near-similar coverage and sparing values. The compared plans are taken from the golden corner at the maximum coverage and maximum sparing cells following the discretization used for the heatmaps.

Figure 6.9 compares random front-plans for "base" where Figure 6.10 preselects them for the best NTCP values. The "guide" plans are preselected for NTCP for both cases to properly show what the dose-response objectives achieve. As the cases in Figure 6.9 demonstrate, the dose distributions for the random front-plans can vary more compared to "guide", whereas the preselected plans of Figure 6.10 have more similar shapes. However, both comparisons still show the same effects, albeit to different extents. Additional comparisons that further show these effects on more cases are available in Appendix A.

These dose distributions demonstrate a shift in focus around the regions considered by the dose-response models. The most important regions to visually analyze the shift for are the CTV_{HR} in red, the bladder in blue, the rectum also in red, the sigmoid in yellow, and the bowel in beige. Generally, the many-objective plans have a higher dose on the CTV_{HR} covered by the TCP objective. Furthermore, given the higher spiking dose in the center of the tumor, the dose at its edges close to the NTCP-covered organs is reduced. Notably, the edge close to the sigmoid, which the NTCP objective does not cover, is commonly irradiated more as a trade-off. The extent of these effects naturally varies per patient and plan. However, it can be observed that the inclusion of the dose-response objective produces distinctly different plans with targeted dose concentrations and organ preservation for the TCP and NTCP subjects, respectively.

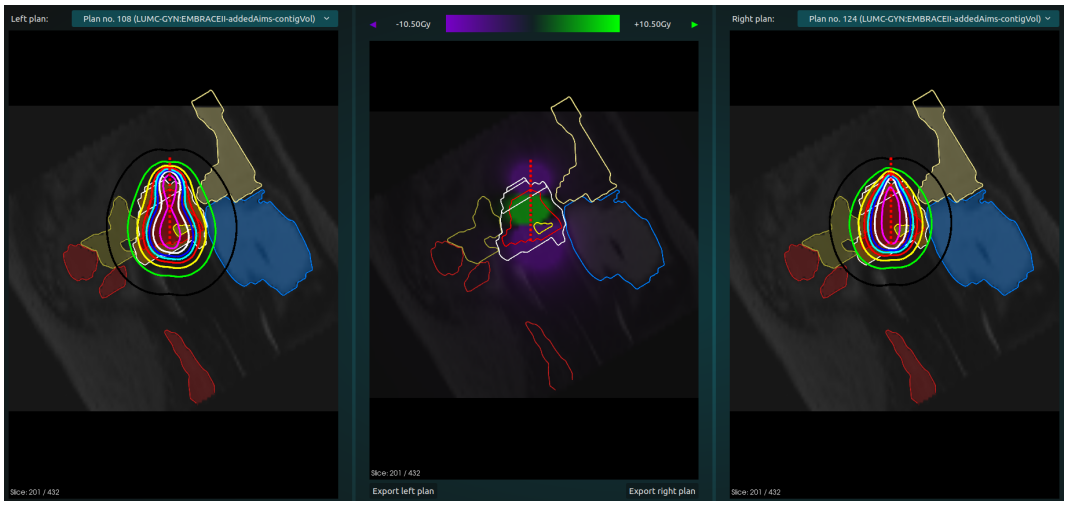

(a) Dose distribution comparison on sparing-oriented plans for patient 10.

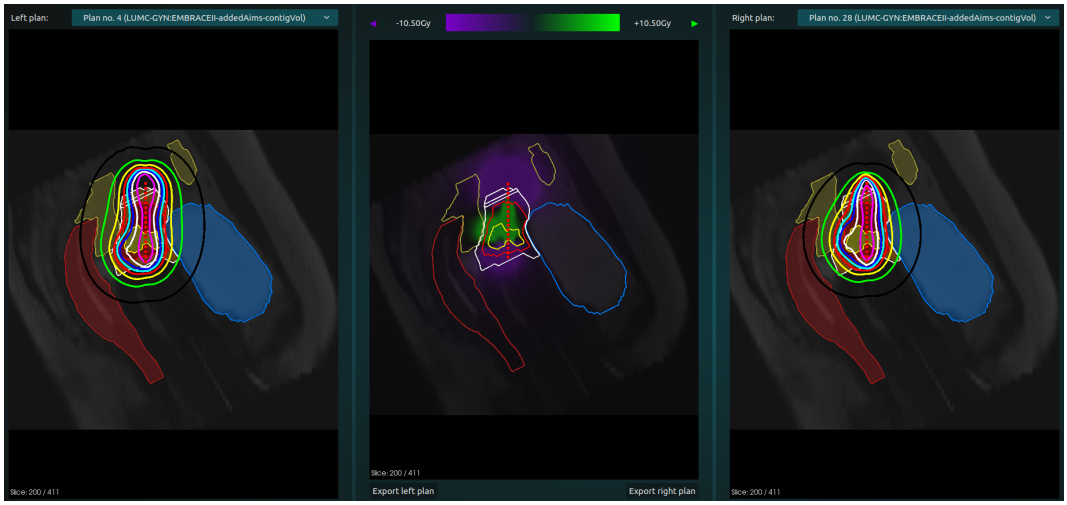

(b) Dose distribution comparison on coverage-oriented plans for patient 16.

Figure 6.9: The dose distribution comparisons of BRIGHT-base plans on the left and guide-candidate plans on the right. The plans show their isodose lines, and the middle highlights their differences. The purple areas indicate a higher BRIGHT-base dose, and the green areas have a higher candidate-guide dose.

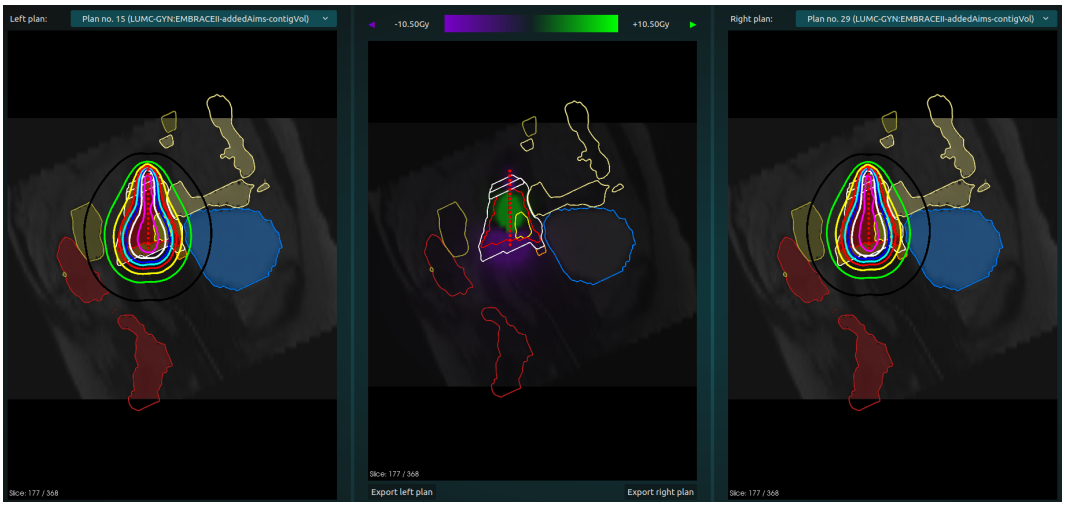

(a) Dose distribution comparison on coverage-oriented plans for patient 1.

(b) Dose distribution comparison on sparing-oriented plans for patient 9.

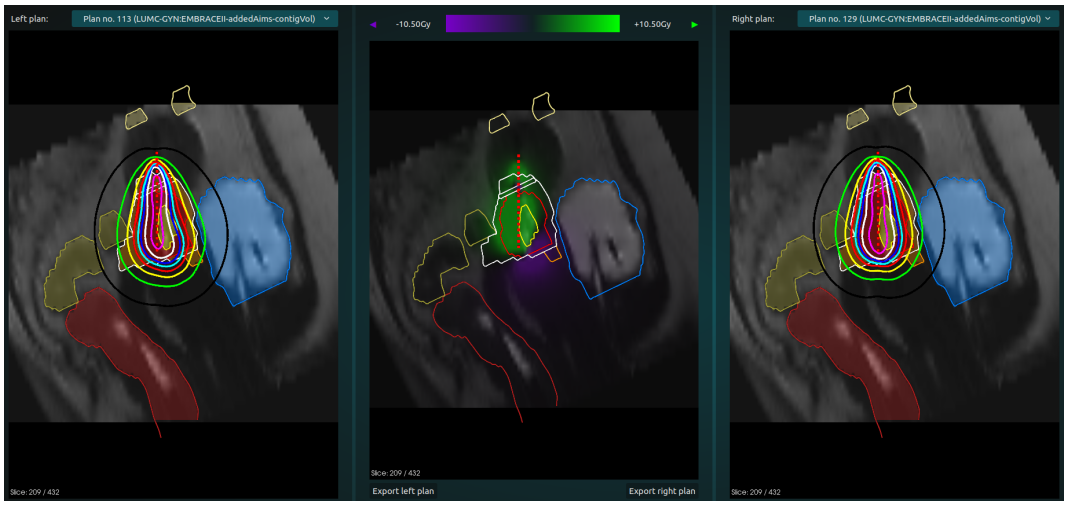

(c) Dose distribution comparison on sparing-oriented plans for patient 13.

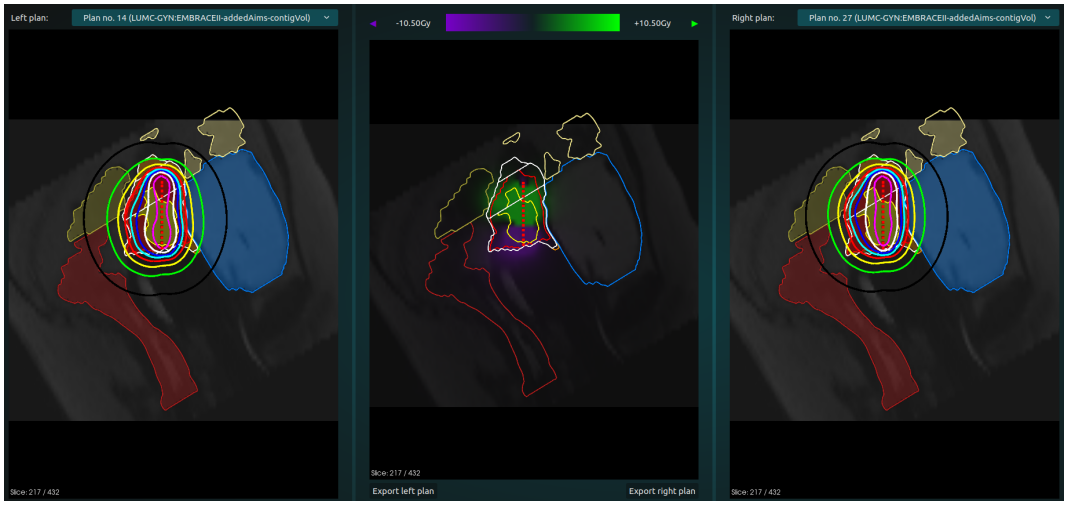

(d) Dose distribution comparison on coverage-oriented plans for patient 14.

Figure 6.10: The dose distribution comparisons of NTCP-preselected BRIGHT-base plans on the left and guide-candidate plans on the right. The plans show their isodose lines, and the middle highlights their differences. The purple areas indicate a higher BRIGHT-base dose, and the green areas have a higher candidate-guide dose.

Chapter 7

Discussion

Throughout this thesis, various methods have been developed and tested to include dose-response models into BRIGHT's optimization process and balance it with the existing objectives during optimization. The experiments with these methods showed several limitations regarding the integration of such models. However, the results also demonstrate distinct differences and improvements that pave the way for future work. In this section, we examine these limitations and achievements to detail the areas that future research could focus on.

7.1 Multi-Objective Usage

The results of the first experiments described in Chapter 4 showed that the tested dose-response objective definitions were insufficient replacements for BRIGHT's existing coverage and sparing. Based on the results, our objectives produced plans with greatly improved dose-response outcomes. However, they were unable to adhere to the protocol in many cases and generally fell short considering the DVI-based metrics. The current dose-response objectives generally align with coverage and sparing but not enough for this substitution to be viable.

This outcome was not entirely unsuspected. The protocol considers several DVIs for various regions, many of which the dose-response models do not cover. Additionally, the added objective mostly shapes the plans and cannot satisfy the remaining parts either. Nevertheless, the idea of this approach can work. Dose-response models are still heavily researched and keep improving. More or better models covering all aspects of the protocol are required. With these, the objectives might be able to optimize the predicted dose-response outcomes and the protocol compliance simultaneously. Only when this is achieved can these objectives be considered a suitable replacement. Up till then, their usage should remain limited to the many-objective setting.

Nevertheless, despite the multi-objective usage not being viable considering the protocol compliance, it is useful to evaluate the dose-response objective formulations. Optimizing them separately from the coverage and sparing objectives provides useful insights into their correlation and performance in the other space. Using this, we decided to continue with the additive aggregation method in the many-objective setting. Importantly, compar-

isons with this setup are not straightforward. Furthermore, there currently is no optimal way to combine the predictions of several models for use in clinical practice. Through empirical evaluation with the described technique, we landed at our choice. However, concerning both clinical use and optimization use, more research into an optimal combination technique is required here.

7.2 Many-Objective Usage

The inclusion of the dose-response objectives in a many-objective setting besides the existing objectives proved more successful. The results reported in Chapter 6 show improved TCP and NTCP values while remaining on or close to the coverage-sparing front. Furthermore, distinct shifts in the dose distributions of the new plans have been observed. Plans with improved dose-response outcomes often show increased dose intensities on CTV_{HR} and reduced intensities at the tumor edges close to the bladder, rectum, and bowel. Despite improvements being small, the apparent change in distribution proves the potential of our concept and execution.

Currently, the obtained improvements are around 0.004-0.005, an average half percent improvement for each model. The biggest limitation here is that the current version of BRIGHT already optimizes most DVIs used by the dose-response models, thus already achieving good results. Alternatively, this limitation can be viewed as a lack of models or model quality. Through the covariates and their slight non-linearity, these models can provide unique insights and consider patient-specific needs that the DVI-based objectives cannot cover. However, the current models used in this thesis fail to extract this potential.

The NTCP models do not consider any non-DVI covariate. Additionally, the single TCP model does have multiple, but they are ineffective as there are no other TCP models. The non-DVI covariate values are unaffected by the dose distributions and only shift the risk functions. When combining multiple functions, these shifts change the relative balance. However, for a single model, they act as a fixed shift for the objective value; this does not affect the Pareto dominance or normalized Tchebycheff evaluations. Nevertheless, the slight non-linearity and the reduced set of organs compared to the sparing objective still managed to improve plans. However, it is up to future research to explore and exploit the power of non-DVI covariates. To bring more models and covariates and study what improvements are possible when this unique and patient-specific information directly influences the optimization process.

A final direction that future work could focus on regards the Tchebycheff adaptations we designed to improve the many-objective optimization process. The results from Chapter 6 showed that directly using the default implementation was statistically significantly worse than or equal to using the normal MO-RV-GOMEA. This was most likely due to the range discrepancies as explained in Chapter 5. Our fixed version for this using range-normalized evaluations was better and managed to match the normal setting with the best approach depending on patient-specifics. Lastly, we introduced the guided weight vector spread on top of the range-normalized evaluations. It outperformed all and was statistically equal or significantly improved upon the others for TCP and NTCP, respectively.

Notably, these techniques were developed for their use in BRIGHT with its highly correlated objectives and its golden-corner focus. Nevertheless, they could be relevant in both general and niche applications. The range-normalized evaluations should transfer to almost all use cases, as range discrepancies are common and generally undesired. Even when no range discrepancies are present, its normalized utopian point definition can still improve optimization guidance. The guided weight vector spread is more niche but also configurable. Its general use can shift the focus to one or possibly multiple areas of the search space with variable intensities based on scalars and cutoffs. Future research can explore these directions on more general many-objective problems.

Chapter 8

Conclusion

This thesis demonstrates how dose-response models can be directly included in the optimization of cervical cancer HDR brachytherapy treatment plans using BRIGHT. Several dose-response objective formulations and optimization techniques have been developed and evaluated to investigate their impact as either a replacement or extension of the current objectives.

Empirical results show that for the current set of models, these dose-response objectives are not a viable replacement for the existing coverage and sparing objectives. Although plans with better dose-response outcomes are found, these plans fail to satisfy the protocol for regions not covered by the models. A more complete set of models is required for this to be feasible. The results also show that the natural objective formulations through multiplicative and additive aggregations perform best with the current models. Given this, the additive approach is preferred due to its advantages concerning scalability and extensibility.

Optimizing these dose-response objectives alongside BRIGHT's existing objectives performs better. This approach produces plans with better dose-response outcomes while remaining on or near the coverage-sparing front. However, this many-objective definition does complicate the optimization process, thereby worsening results. To combat this, we applied the many-objective MO-RV-GOMEA variant with Tchebycheff scalarizations [34] and developed two new extensions for it: normalized evaluations and the guided weight vector spread. Combining all of these techniques significantly improves the results, as indicated by the Wilcoxon signed-rank test.

Final results in and around BRIGHT's golden corner show average dose-response improvements of around 0.004 for NTCP models at the front and around 0.005 for TCP and NTCP models by sacrificing some sparing. Furthermore, the plans with improved dose-response outcomes show distinct differences in their dose distributions. Ultimately, the results seem only marginal as the current models overlap too much with BRIGHT and lack non-DVI covariates. However, these results prove that the concept is valid and the execution works. It sets a strong foundation to build from when more and better models become available.

Acknowledgments

We want to thank Dr. Stefan Ecker, Prof. dr. Christian Kirisits, and Dr. Nicole Eder-Nesvacil from the Medical University of Vienna and Dr. Bradley Pieters and Danique Barten from the Amsterdam University Medical Center for their engagement with the work. Particularly, we are grateful to Stefan Ecker for his advice on adapting the dose-response models to suit the available patient data. Additionally, we thank Lavinia Verhagen, Anouk Corbeau, and Dr. Stephanie de Boer from the Leiden University Medical Center for their efforts in collecting and providing the patient data used in this study.

Bibliography

[1] Douglas W Arthur, Derrick Koo, Robert D Zwicker, Shidong Tong, Harry D Bear, Brian J Kaplan, Brian D Kavanagh, Laurel A Warwicke, Diane Holdford, Cyrus Amir, et al. Partial breast brachytherapy after lumpectomy: Low-dose-rate and high-dose-rate experience. *International Journal of Radiation Oncology* Biology* Physics*, 56 (3):681–689, 2003.

[2] Thomas Bäck and Hans-Paul Schwefel. An overview of evolutionary algorithms for parameter optimization. *Evolutionary computation*, 1(1):1–23, 1993.

[3] Kathy L Baglan, Alvaro A Martinez, Robert C Frazier, Vijay R Kini, Larry L Kestin, Peter Y Chen, Greg Edmundson, Elizabeth Mele, David Jaffray, and Frank A Vicini. The use of high-dose-rate brachytherapy alone after lumpectomy in patients with early-stage breast cancer treated with breast-conserving therapy. *International Journal of Radiation Oncology* Biology* Physics*, 50(4):1003–1011, 2001.

[4] Danique LJ Barten, Bradley R Pieters, Anton Bouter, Marjolein C van der Meer, Stef C Maree, Karel A Hinnen, Henrike Westerveld, Peter AN Bosman, Tanja Alderliesten, Niek van Wieringen, et al. Towards artificial intelligence-based automated treatment planning in clinical practice: A prospective study of the first clinical experiences in high-dose-rate prostate brachytherapy. *Brachytherapy*, 22(2):279–289, 2023.

[5] Michael Baumann and Cordula Petersen. TCP and NTCP: A basic introduction. *RAYS-ROME-*, 30(2):99, 2005.

[6] Tobias Blickle and Lothar Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4(4):361–394, 1996.

[7] Peter AN Bosman. The anticipated mean shift and cluster registration in mixture-based EDAs for multi-objective optimization. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 351–358, 2010.

[8] Anton Bouter, Tanja Alderliesten, Cees Witteveen, and Peter AN Bosman. Exploiting linkage information in real-valued optimization with the real-valued gene-pool opti-

mal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 705–712, 2017.

[9] Anton Bouter, Ngoc Hoang Luong, Cees Witteveen, Tanja Alderliesten, and Peter AN Bosman. The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 537–544, 2017.

[10] Anton Bouter, Tanja Alderliesten, Bradley R Pieters, Arjan Bel, Yury Niatsetski, and Peter AN Bosman. GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy. *Medical physics*, 46(9):3776–3787, 2019.

[11] Freddie Bray, Mathieu Laversanne, Hyuna Sung, Jacques Ferlay, Rebecca L Siegel, Isabelle Soerjomataram, and Ahmedin Jemal. Global cancer statistics 2022: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA: A cancer journal for clinicians*, 74(3):229–263, 2024.

[12] Anthony TC Chan, SF Leung, Roger KC Ngan, Peter ML Teo, WH Lau, WH Kwan, Edwin P Hui, HY Yiu, Winnie Yeo, FY Cheung, et al. Overall survival after concurrent cisplatin–radiotherapy compared with radiotherapy alone in locoregionally advanced nasopharyngeal carcinoma. *Journal of the National Cancer Institute*, 97(7):536–539, 2005.

[13] David R Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.

[14] Kalyanmoy Deb and David E Goldberg. Analyzing deception in trap functions. In *Foundations of genetic algorithms*, volume 2, pages 93–108. Elsevier, 1993.

[15] Kalyanmoy Deb and David E Goldberg. Sufficient conditions for deceptive and easy binary functions. *Annals of mathematics and Artificial Intelligence*, 10:385–408, 1994.

[16] D Jeffrey Demanes and Michel I Ghilezan. High-dose-rate brachytherapy as monotherapy for prostate cancer. *Brachytherapy*, 13(6):529–541, 2014.

[17] Leah RM Dickhoff, Ellen M Kerkhof, Heloisa H Deuzeman, Carien L Creutzberg, Tanja Alderliesten, and Peter AN Bosman. Adaptive objective configuration in bi-objective evolutionary optimization for cervical cancer brachytherapy treatment planning. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1173–1181, 2022.

[18] Johannes CA Dimopoulos, Richard Pötter, Stefan Lang, Elena Fidarova, Petra Georg, Wolfgang Dörr, and Christian Kirisits. Dose–effect relationship for local control of cervical cancer by magnetic resonance image-guided brachytherapy. *Radiotherapy and oncology*, 93(2):311–315, 2009.

[19] Stefan Ecker, Christian Kirisits, Maximilian Schmid, Johannes Knoth, Gerd Heilemann, Astrid De Leeuw, Alina Sturdza, Kathrin Kirchheiner, Nina Jensen, Remi Nout, et al. EviGUIDE-a tool for evidence-based decision making in image-guided adaptive brachytherapy for cervical cancer. *Radiotherapy and Oncology*, 186:109748, 2023.

[20] Peter JB Hancock. An empirical comparison of selection methods in evolutionary algorithms. In *AISB workshop on evolutionary computing*, pages 80–94. Springer, 1994.

[21] John H Holland. *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. MIT press, 1992.

[22] Christian Horoba and Frank Neumann. Benefits and drawbacks for the use of epsilon-dominance in evolutionary multi-objective optimization. In *Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 641–648, 2008.

[23] Peter J Hoskin, Ana M Rojas, Peter J Bownes, Gerry J Lowe, Peter J Ostler, and Linda Bryant. Randomised trial of external beam radiotherapy alone or combined with high-dose-rate brachytherapy boost for localised prostate cancer. *Radiotherapy and Oncology*, 103(2):217–222, 2012.

[24] Johannes Karlsson, Ann-Charlotte Dreifaldt, Louise Bohr Mordhorst, and Bengt Sorbe. Differences in outcome for cervical cancer patients treated with or without brachytherapy. *Brachytherapy*, 16(1):133–140, 2017.

[25] Sourabh Katoch, Sumit Singh Chauhan, and Vijay Kumar. A review on genetic algorithm: Past, present, and future. *Multimedia tools and applications*, 80:8091–8126, 2021.

[26] Hamid Reza Khalkhali, Rasool Gharaaghaji, Rohollah Valizadeh, Zahra Kousehlou, and Haleh Ayatollahi. Ten years' survival in patients with cervical cancer and related factors in West Azerbaijan Province: Using of Cox proportion hazard model. *Asian Pacific journal of cancer prevention: APJCP*, 20(5):1345, 2019.

[27] Kathrin Kirchheiner, Remi A Nout, Jacob C Lindegaard, Christine Haie-Meder, Umesh Mahantshetty, Barbara Segedin, Ina M Jürgenliemk-Schulz, Peter J Hoskin, Bhavana Rai, Wolfgang Dörr, et al. Dose–effect relationship and risk factors for vaginal stenosis after definitive radio (chemo) therapy with image-guided brachytherapy for locally advanced cervical cancer in the EMBRACE study. *Radiotherapy and oncology*, 118(1):160–166, 2016.

[28] Christian Kirisits, Richard Pötter, Stefan Lang, Johannes Dimopoulos, Natascha Wachter-Gerstner, and Dietmar Georg. Dose and volume parameters for MRI-based treatment planning in intracavitary brachytherapy for cervical cancer. *International Journal of Radiation Oncology* Biology* Physics*, 62(3):901–911, 2005.

[29] Andreas C Koenig. A study of mutation methods for evolutionary algorithms. *University of Missouri-Rolla*, 2002.

[30] Jina Li, Gaoming Liu, Jiayou Luo, Shipeng Yan, Ping Ye, Jie Wang, and Miyang Luo. Cervical cancer prognosis and related risk factors for patients with cervical cancer: A long-term retrospective cohort study. *Scientific Reports*, 12(1):13994, 2022.

[31] Ruifeng Liu, XiaoHu Wang, JinHui Tian, KeHu Yang, Jun Wang, Lei Jiang, and Xiang Yong Hao. High dose rate versus low dose rate intracavity brachytherapy for locally advanced uterine cervix cancer. *Cochrane Database of Systematic Reviews*, (10), 2014.

[32] Ngoc Hoang Luong, Han La Poutré, and Peter AN Bosman. Multi-objective gene-pool optimal mixing evolutionary algorithms. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 357–364, 2014.

[33] Ngoc Hoang Luong, Anton Bouter, Marjolein C Van Der Meer, Yury Niatsetski, Cees Witteveen, Arjan Bel, Tanja Alderliesten, and Peter AN Bosman. Efficient, effective, and insightful tackling of the high-dose-rate brachytherapy treatment planning problem for prostate cancer using evolutionary multi-objective optimization algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pages 1372–1379, 2017.

[34] Ngoc Hoang Luong, Tanja Alderliesten, and Peter AN Bosman. Improving the performance of MO-RV-GOMEA on problems with many objectives using Tchebycheff scalarizations. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 705–712, 2018.

[35] Stefanus C Maree, Ngoc Hoang Luong, Ernst S Kooreman, Niek van Wieringen, Arjan Bel, Karel A Hinnen, Henrike Westerveld, Bradley R Pieters, Peter AN Bosman, and Tanja Alderliesten. Evaluation of bi-objective treatment planning for high-dose-rate prostate brachytherapy—A retrospective observer study. *Brachytherapy*, 18(3):396–403, 2019.

[36] Rafael Martínez-Monge and Alfonso Gómez-Iturriaga. High-dose-rate brachytherapy in lower eyelid cancer. *Brachytherapy*, 6(3):227–229, 2007.

[37] Renaud Mazeron, Lars U Fokdal, Kathrin Kirchheiner, Petra Georg, Noha Jastaniyah, Barbara Šegedin, Umesh Mahantshetty, Peter Hoskin, Ina Jürgenliemk-Schulz, Christian Kirisits, et al. Dose–volume effect relationships for late rectal morbidity in patients treated with chemoradiation and MRI-guided adaptive brachytherapy for locally advanced cervical cancer: Results from the prospective multicenter EMBRACE study. *Radiotherapy and oncology*, 120(3):412–419, 2016.

[38] Lucas C Mendez and Gerard C Morton. High dose-rate brachytherapy in the treatment of prostate cancer. *Translational andrology and urology*, 7(3):357, 2018.

[39] Subir Nag, Elmer R Cano, D Jeffrey Demanes, Ajmel A Puthawala, Bhadrasain Vikram, et al. The American Brachytherapy Society recommendations for high-dose-rate brachytherapy for head-and-neck carcinoma. *International Journal of Radiation Oncology* Biology* Physics*, 50(5):1190–1198, 2001.

[40] Subir Nag, Jeanne M Quivey, John D Earle, David Followill, James Fontanesi, Paul T Finger, and American Brachytherapy Society. The American Brachytherapy Society recommendations for brachytherapy of uveal melanomas. *International Journal of Radiation Oncology* Biology* Physics*, 56(2):544–555, 2003.

[41] Christel N Nomden, Astrid AC De Leeuw, Judith M Roesink, Robbert JHA Tersteeg, Marinus A Moerland, Petronella O Witteveen, Henk W Schreuder, Eleonore BL Van Dorst, and Ina Maria Jürgenliemk-Schulz. Clinical outcome and dosimetric parameters of chemo-radiation including MRI guided adaptive brachytherapy with tandem-ovoid applicators for cervical cancer patients: A single institution experience. *Radiotherapy and oncology*, 107(1):69–74, 2013.

[42] Firuza D Patel, Suresh C Sharma, Pritam S Negi, Sushmita Ghoshal, and Brahm D Gupta. Low dose rate vs. high dose rate brachytherapy in the treatment of carcinoma of the uterine cervix: A clinical trial. *International Journal of Radiation Oncology* Biology* Physics*, 28(2):335–341, 1994.

[43] Richard Pötter, Petra Georg, Johannes CA Dimopoulos, Magdalena Grimm, Daniel Berger, Nicole Nesvacil, Dietmar Georg, Maximilian P Schmid, Alexander Reinthaller, Alina Sturdza, et al. Clinical outcome of protocol based image (MRI) guided adaptive brachytherapy combined with 3D conformal radiotherapy with or without chemotherapy in patients with locally advanced cervical cancer. *Radiotherapy and Oncology*, 100(1):116–123, 2011.

[44] Richard Pötter, Kari Tanderup, Christian Kirisits, Astrid de Leeuw, Kathrin Kirchheiner, Remi Nout, Li Tee Tan, Christine Haie-Meder, Umesh Mahantshetty, Barbara Segedin, et al. The EMBRACE II study: The outcome and prospect of two decades of evolution within the GEC-ESTRO GYN working group and the EMBRACE studies. *Clinical and translational radiation oncology*, 9:48–60, 2018.

[45] Christian Ritz. Toward a unified approach to dose–response modeling in ecotoxicology. *Environmental toxicology and chemistry*, 29(1):220–229, 2010.

[46] Mark J Rivard, Bert M Coursey, Larry A DeWerd, William F Hanson, M Saiful Huq, Geoffrey S Ibbott, Michael G Mitch, Ravinder Nath, and Jeffrey F Williamson. Update of AAPM Task Group No. 43 Report: A revised AAPM protocol for brachytherapy dose calculations. *Medical physics*, 31(3):633–674, 2004.

[47] Sílvio Rodrigues, Pavol Bauer, and Peter AN Bosman. A novel population-based multi-objective CMA-ES and the impact of different constraint handling techniques. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 991–998, 2014.

[48] Maximilian P Schmid, Jacob C Lindegaard, Umesh Mahantshetty, Kari Tanderup, Ina Jürgenliemk-Schulz, Christine Haie-Meder, Lars U Fokdal, Alina Sturdza, Peter Hoskin, Barbara Segedin, et al. Risk factors for local failure following chemoradiation and magnetic resonance image–guided brachytherapy in locally advanced cervical

cancer: Results from the EMBRACE-I study. *Journal of Clinical Oncology*, 41(10): 1933–1942, 2023.

[49] Sofia Spampinato, Lars U Fokdal, Richard Pötter, Christine Haie-Meder, Jacob C Lindegaard, Maximilian P Schmid, Alina Sturdza, Ina M Jürgenliemk-Schulz, Umesh Mahantshetty, Barbara Segedin, et al. Importance of the ICRU bladder point dose on incidence and persistence of urinary frequency and incontinence in locally advanced cervical cancer: An EMBRACE analysis. *Radiotherapy and Oncology*, 158:300–308, 2021.

[50] Sofia Spampinato, Lars U Fokdal, Richard Pötter, Christine Haie-Meder, Jacob C Lindegaard, Maximilian P Schmid, Alina Sturdza, Ina M Jürgenliemk-Schulz, Umesh Mahantshetty, Barbara Segedin, et al. Risk factors and dose-effects for bladder fistula, bleeding and cystitis after radiotherapy with imaged-guided adaptive brachytherapy for cervical cancer: An EMBRACE analysis. *Radiotherapy and Oncology*, 158:312–320, 2021.

[51] Sofia Spampinato, Nina BK Jensen, Richard Pötter, Lars U Fokdal, Cyrus Chargari, Jacob C Lindegaard, Maximilian P Schmid, Alina Sturdza, Ina M Jürgenliemk-Schulz, Umesh Mahantshetty, et al. Severity and persistency of late gastrointestinal morbidity in locally advanced cervical cancer: Lessons learned from EMBRACE-I and implications for the future. *International Journal of Radiation Oncology* Biology* Physics*, 112(3):681–693, 2022.

[52] William M Spears. *Evolutionary Algorithms: The role of mutation and recombination*. Springer Science & Business Media, 2000.

[53] Stefan Strohmaier and Grzegorz Zwierzchowski. Comparison of 60 Co and 192 Ir sources in HDR brachytherapy. *Journal of contemporary brachytherapy*, 3(4):199–208, 2011.

[54] Thanatip Tantivatana and Kanisa Rongsriyam. Treatment outcomes of high-dose-rate intracavitary brachytherapy for cervical cancer: A comparison of Ir-192 versus Co-60 sources. *Journal of Gynecologic Oncology*, 29(5):e86, 2018.

[55] Peter FM Teunis, Nico JD Nagelkerke, and Charles N Haas. Dose response models for infectious gastroenteritis. *Risk Analysis*, 19(6):1251–1260, 1999.

[56] Dirk Thierens and Peter AN Bosman. Optimal mixing evolutionary algorithms. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 617–624, 2011.

[57] Anant J Umbarkar and Pranali D Sheth. Crossover operators in genetic algorithms: A review. *ICTACT journal on soft computing*, 6(1), 2015.

[58] U.S. Department of Health and Human Services, National Cancer Institute. Common Terminology Criteria for Adverse Events (CTCAE), November 2017.

URL https://ctep.cancer.gov/protocolDevelopment/electronic_applications/ctc.htm. Accessed: 2025-04-08.

[59] Gustavo A Viani, Gustavo B Manta, Eduardo J Stefano, and Ligia I de Fendi. Brachytherapy for cervix cancer: Low-dose rate or high-dose rate brachytherapy–A meta-analysis of clinical trials. *Journal of Experimental & Clinical Cancer Research*, 28:1–12, 2009.

[60] Henrike Westerveld, Kathrin Kirchheiner, Remi A Nout, Kari Tanderup, Jacob C Lindegaard, Sofia Spampinato, Alina Sturdza, Nicole Nesvacil, Kjersti Bruheim, Taran P Hellebust, et al. Dose-effect relationship between vaginal dose points and vaginal stenosis in cervical cancer: An EMBRACE-I sub-study. *Radiotherapy and Oncology*, 168:8–15, 2022.

Appendix A

Additional plots

This appendix provides additional plots for several figures shown throughout this thesis. Figures A.1, A.2, A.3, and A.4 show more cases for the tuning results to support the decisions made. Furthermore, an overview of the golden-corner improvement heatmaps of the best-performing method is provided in Figure A.5, showing all cases not included in the main text. Lastly, extra dose distribution comparisons are included in Figure A.6, giving additional insights into their differences and demonstrating the observed effects over more patients.

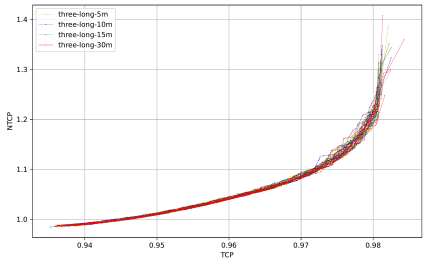

(a) Patient 1: Time limit comparison for normalization method ”long”

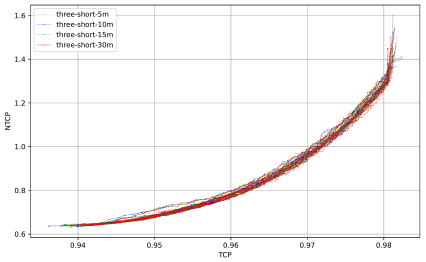

(b) Patient 1: Time limit comparison for normalization method ”short”

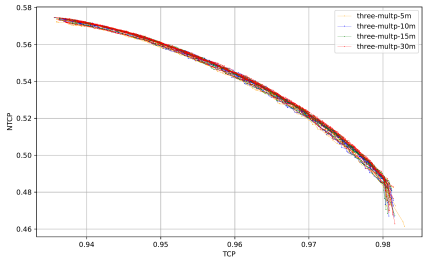

(c) Patient 1: Time limit comparison for aggregation method ”multp”

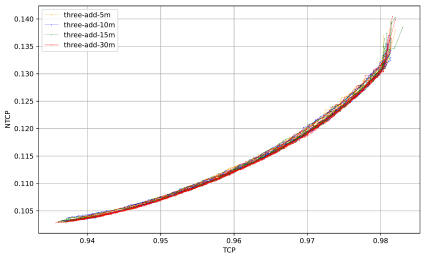

(d) Patient 1: Time limit comparison for aggregation method ”add”

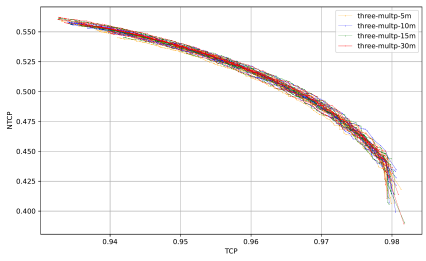

(e) Patient 2: Time limit comparison for aggregation method ”multp”

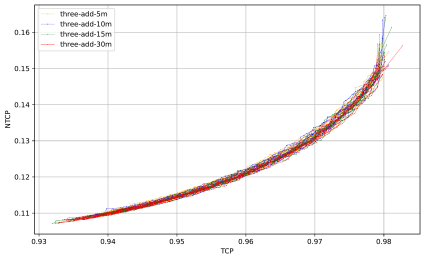

(f) Patient 2: Time limit comparison for aggregation method ”add”

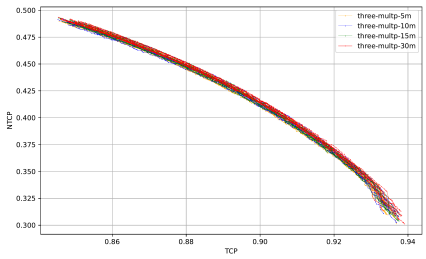

(g) Patient 3: Time limit comparison for aggregation method ”multp”

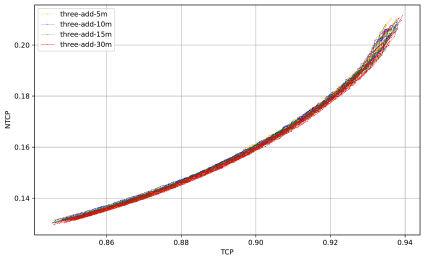

(h) Patient 3: Time limit comparison for aggregation method ”add”

Figure A.1: Additional plots showing the multi-objective time limit differences for the dose-response objective formulations in their own TCP-NTCP spaces.

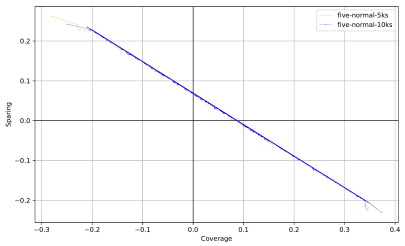

(a) Patient 1: Elitist archive size comparison for run type ”normal”, clusters ”30c”, and time limit ”10m” in coverage-sparing space

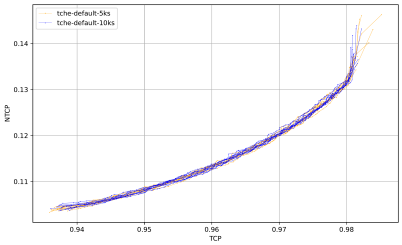

(b) Patient 1: Elitist archive size comparison for run type ”default”, clusters ”20c”, and time limit ”30m” in TCP-NTCP space

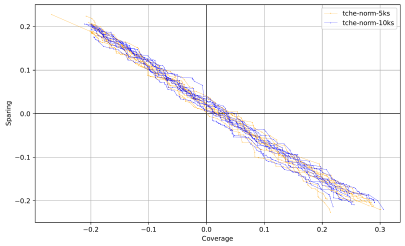

(c) Patient 2: Elitist archive size comparison for run type ”norm”, clusters ”20c”, and time limit ”5m” in coverage-sparing space

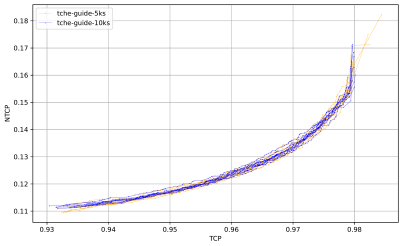

(d) Patient 2: Elitist archive size comparison for run type ”guide”, clusters ”30c”, and time limit ”10m” in TCP-NTCP space

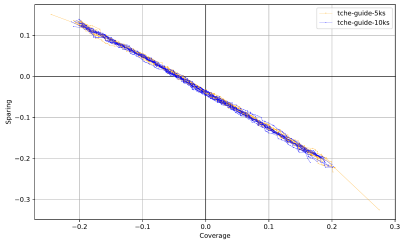

(e) Patient 3: Elitist archive size comparison for run type ”guide”, clusters ”20c”, and time limit ”5m” in coverage-sparing space

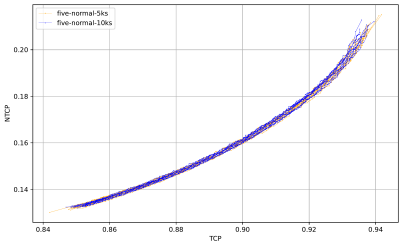

(f) Patient 3: Elitist archive size comparison for run type ”normal”, clusters ”30c”, and time limit ”15m” in TCP-NTCP space

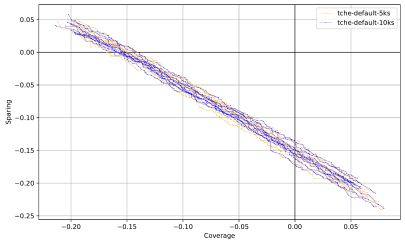

(g) Patient 4: Elitist archive size comparison for run type ”default”, clusters ”20c”, and time limit ”15m” in coverage-sparing space

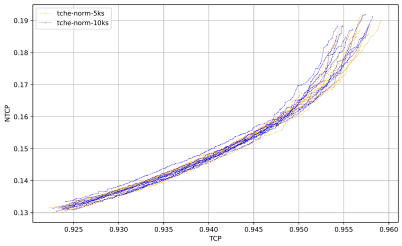

(h) Patient 4: Elitist archive size comparison for run type ”norm”, clusters ”30c”, and time limit ”30m” in TCP-NTCP space

Figure A.2: Additional plots showing the elitist archive size differences for the many-objective techniques in the coverage-sparing and TCP-NTCP spaces.

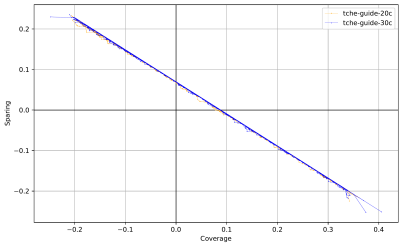

(a) Patient 1: Number of clusters comparison for run type ”guide”, archive size ”5ks”, and time limit ”5m” in coverage-sparing space

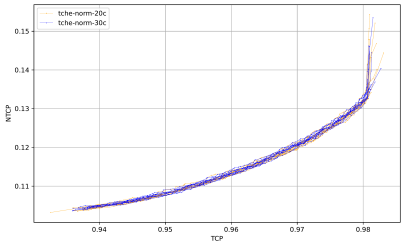

(b) Patient 1: Number of clusters comparison for run type ”norm”, archive size ”5ks”, and time limit ”15m” in TCP-NTCP space

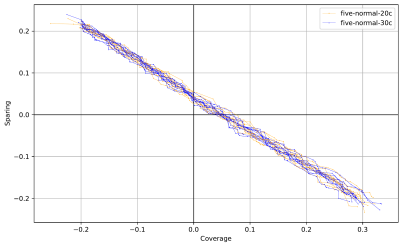

(c) Patient 2: Number of clusters comparison for run type ”normal”, archive size ”5ks”, and time limit ”30m” in coverage-sparing space

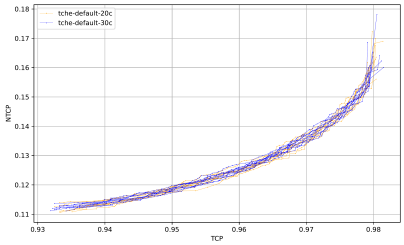

(d) Patient 2: Number of clusters comparison for run type ”default”, archive size ”5ks”, and time limit ”10m” in TCP-NTCP space

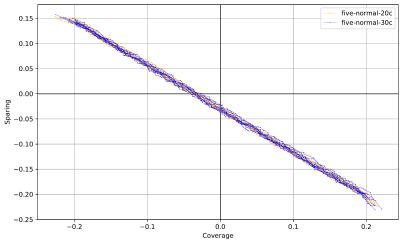

(e) Patient 3: Number of clusters comparison for run type ”normal”, archive size ”5ks”, and time limit ”10m” in coverage-sparing space

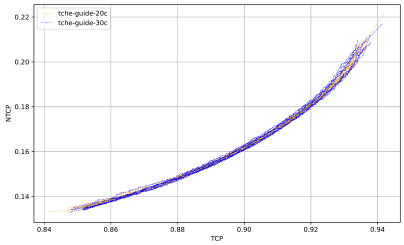

(f) Patient 3: Number of clusters comparison for run type ”guide”, archive size ”5ks”, and time limit ”15m” in TCP-NTCP space

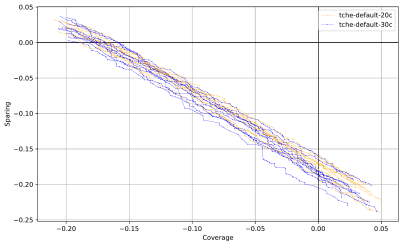

(g) Patient 4: Number of clusters comparison for run type ”default”, archive size ”5ks”, and time limit ”5m” in coverage-sparing space

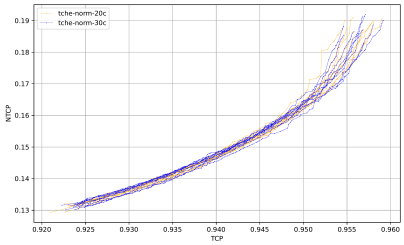

(h) Patient 4: Number of clusters comparison for run type ”norm”, archive size ”5ks”, and time limit ”30m” in TCP-NTCP space

Figure A.3: Additional plots showing the number of clusters’ differences for the many-objective techniques in the coverage-sparing and TCP-NTCP spaces.

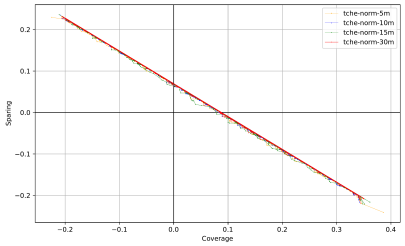

(a) Patient 1: Time limit comparison for run type ”norm”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

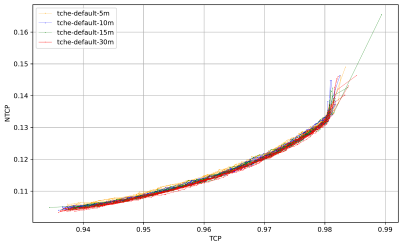

(b) Patient 1: Time limit comparison for run type ”default”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

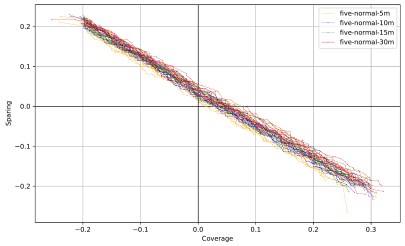

(c) Patient 2: Time limit comparison for run type ”normal”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

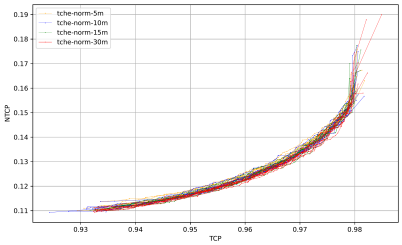

(d) Patient 2: Time limit comparison for run type ”norm”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

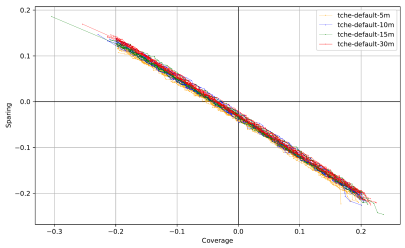

(e) Patient 3: Time limit comparison for run type ”default”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

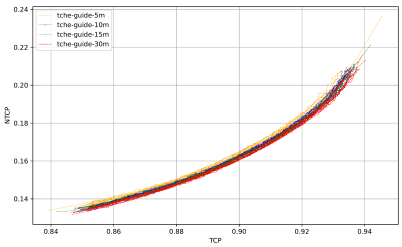

(f) Patient 3: Time limit comparison for run type ”guide”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

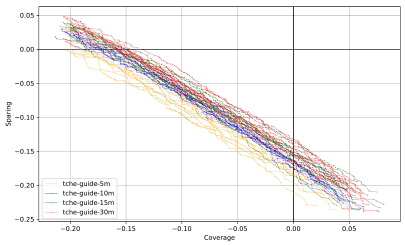

(g) Patient 4: Time limit comparison for run type ”guide”, clusters ”20c”, and archive size ”5ks” in coverage-sparing space

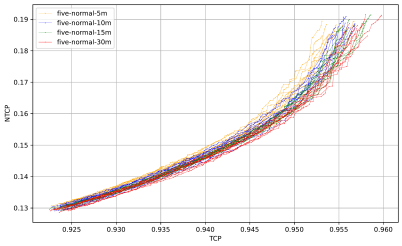

(h) Patient 4: Time limit comparison for run type ”normal”, clusters ”20c”, and archive size ”5ks” in TCP-NTCP space

Figure A.4: Additional plots showing the time limit differences for the many-objective techniques in the coverage-sparing and TCP-NTCP spaces.

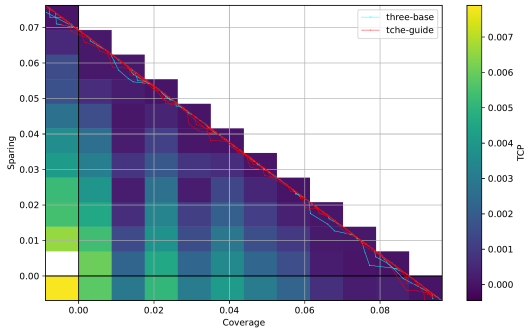

(a) Patient 1: Golden-corner TCP differences between run types ”base” and ”guide”

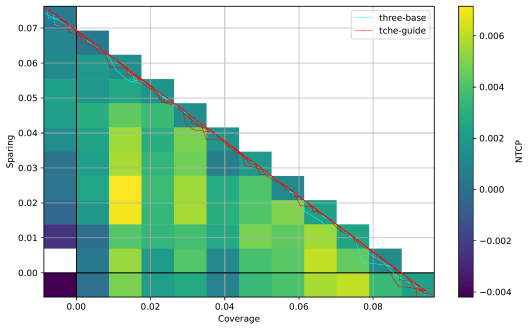

(b) Patient 1: Golden-corner NTCP differences between run types ”base” and ”guide”

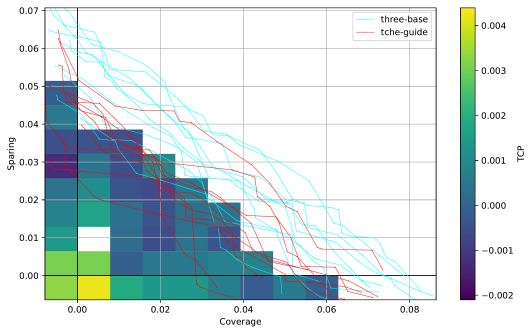

(c) Patient 2: Golden-corner TCP differences between run types ”base” and ”guide”

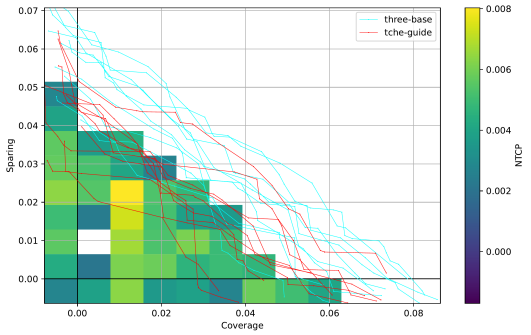

(d) Patient 2: Golden-corner NTCP differences between run types ”base” and ”guide”

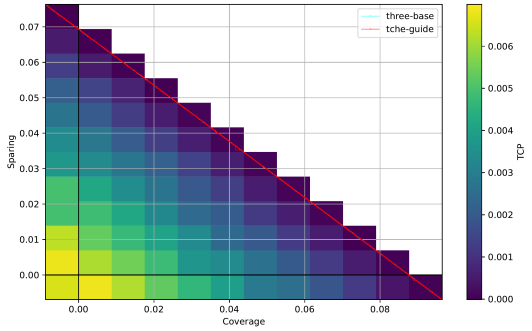

(e) Patient 5: Golden-corner TCP differences between run types ”base” and ”guide”

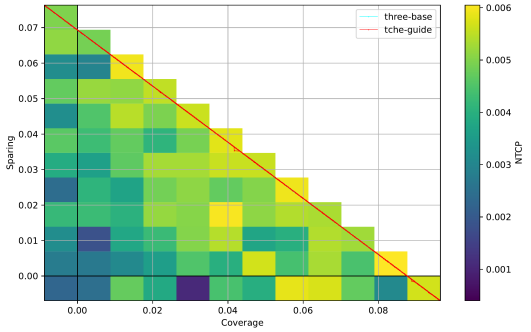

(f) Patient 5: Golden-corner NTCP differences between run types ”base” and ”guide”

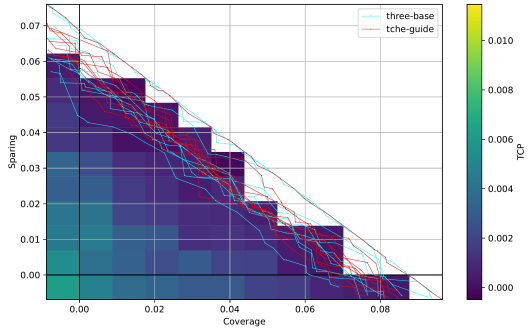

(g) Patient 6: Golden-corner TCP differences between run types ”base” and ”guide”

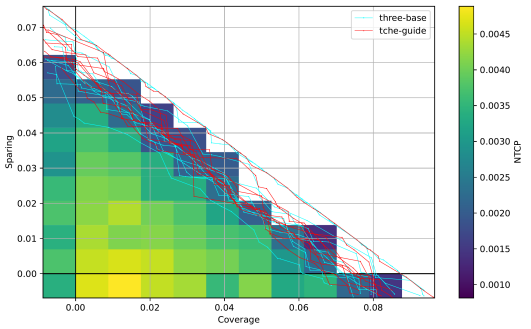

(h) Patient 6: Golden-corner NTCP differences between run types ”base” and ”guide”

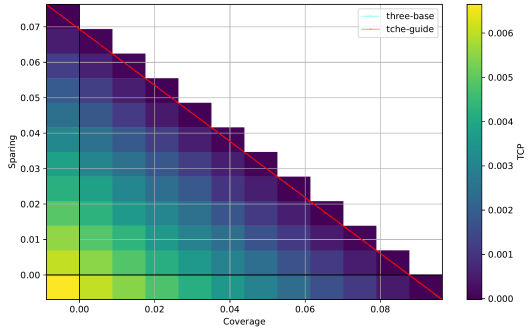

(i) Patient 10: Golden-corner TCP differences between run types ”base” and ”guide”

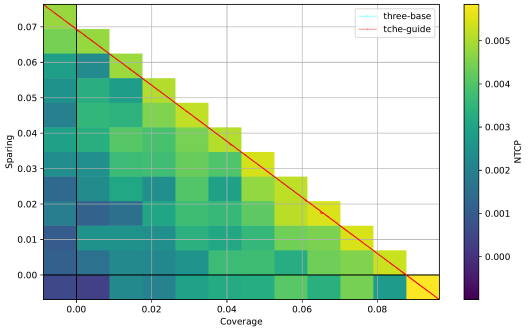

(j) Patient 10: Golden-corner NTCP differences between run types ”base” and ”guide”

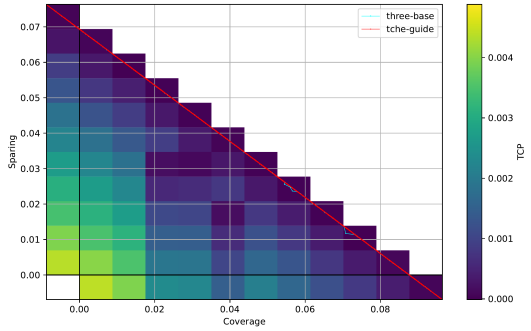

(k) Patient 11: Golden-corner TCP differences between run types ”base” and ”guide”

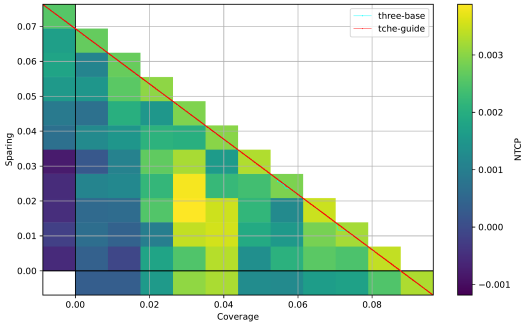

(l) Patient 11: Golden-corner NTCP differences between run types ”base” and ”guide”

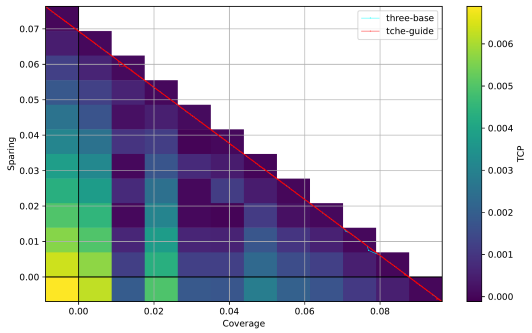

(m) Patient 13: Golden-corner TCP differences between run types ”base” and ”guide”

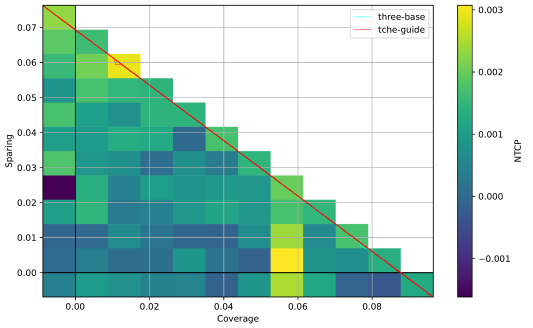

(n) Patient 13: Golden-corner NTCP differences between run types ”base” and ”guide”

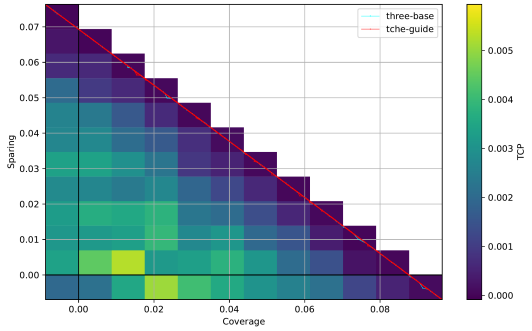

(o) Patient 14: Golden-corner TCP differences between run types ”base” and ”guide”

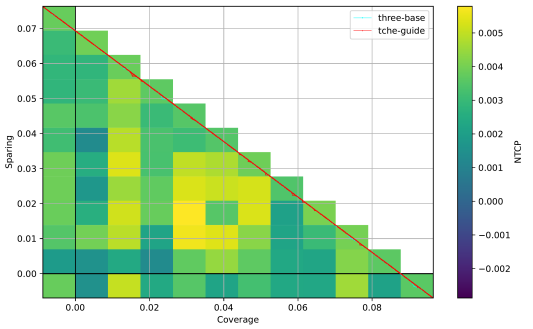

(p) Patient 14: Golden-corner NTCP differences between run types ”base” and ”guide”

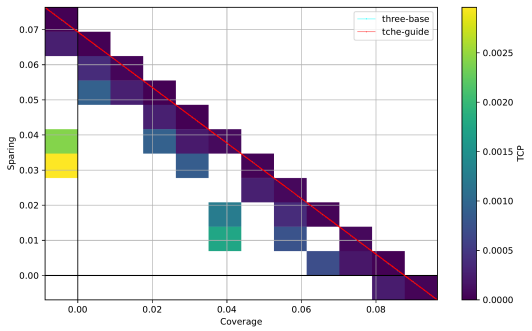

(q) Patient 15: Golden-corner TCP differences between run types ”base” and ”guide”

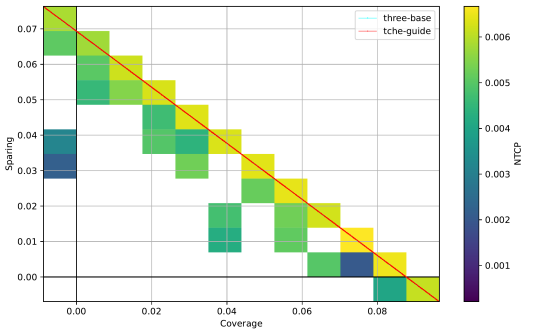

(r) Patient 15: Golden-corner NTCP differences between run types ”base” and ”guide”

Figure A.5: Additional plots showing heatmap TCP and NTCP golden-corner differences for baseline ”base” and best-performing candidate ”guide” for all patients not already shown in the main text. The results are displayed in the coverage-sparing space.

(a) Dose distribution comparison on sparing-oriented random-front ”base” plans and NTCP-preselected ”guide” for patient 2.

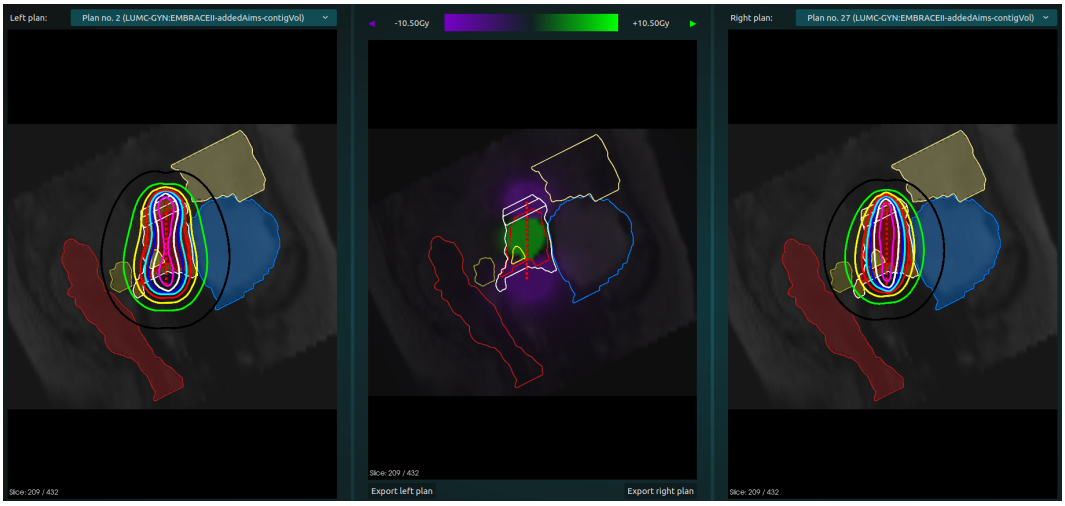

(b) Dose distribution comparison on coverage-oriented random-front ”base” plans and NTCP-preselected ”guide” for patient 5.

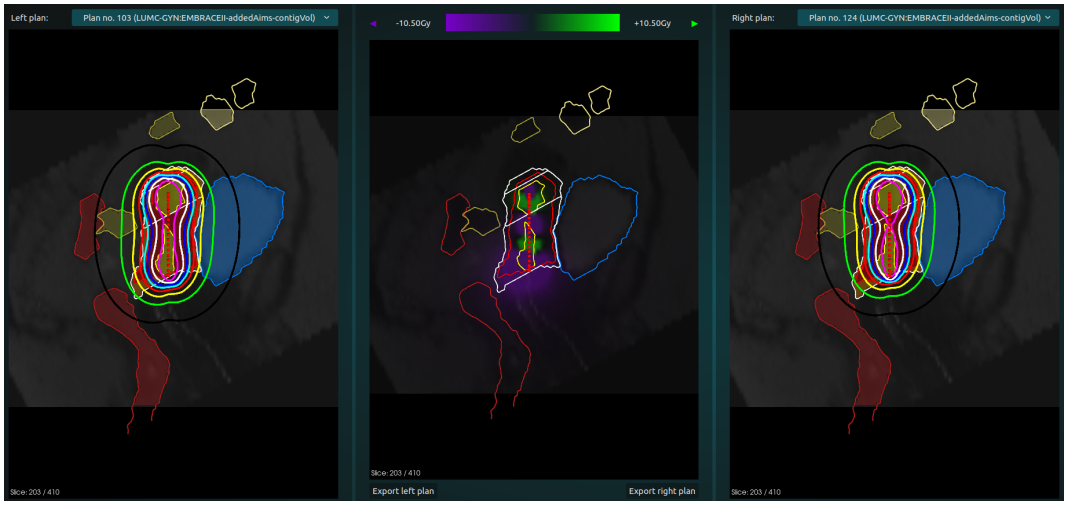

(c) Dose distribution comparison on sparing-oriented random-front ”base” plans and NTCP-preselected ”guide” plans for patient 15.

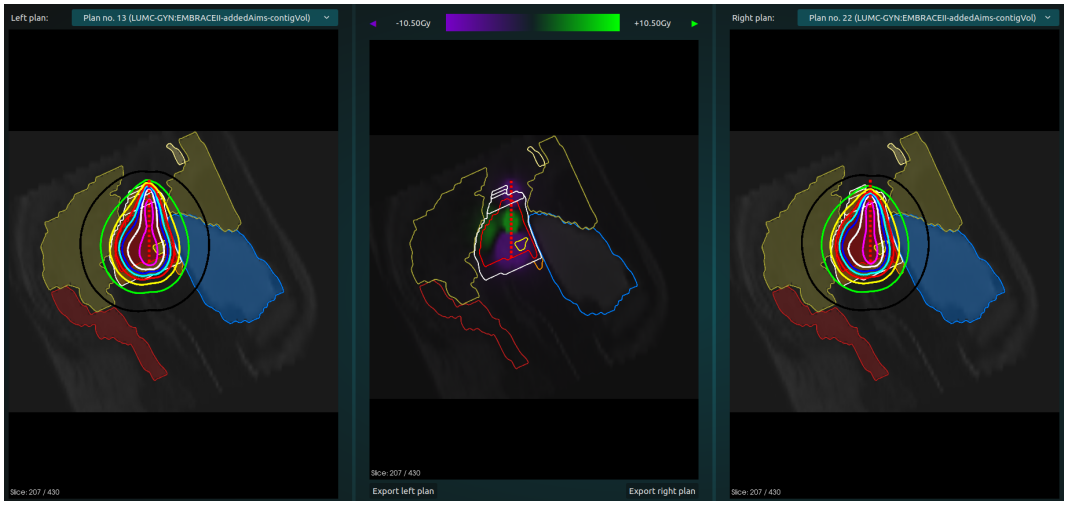

(d) Dose distribution comparison on coverage-oriented NTCP-preselected ”base” and ”guide” plans for patient 6.

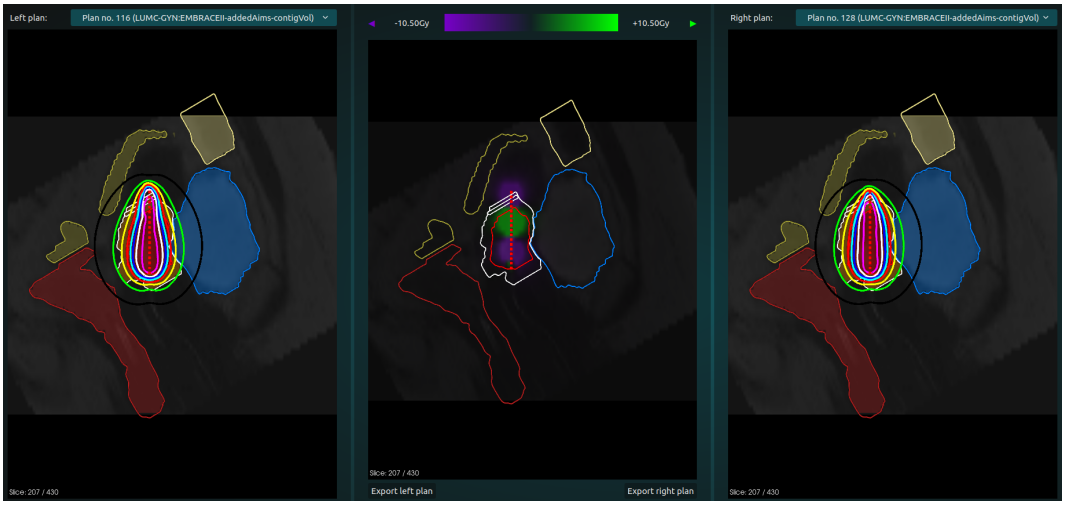

(e) Dose distribution comparison on sparing-oriented NTCP-preselected ”base” and ”guide” plans for patient 11.

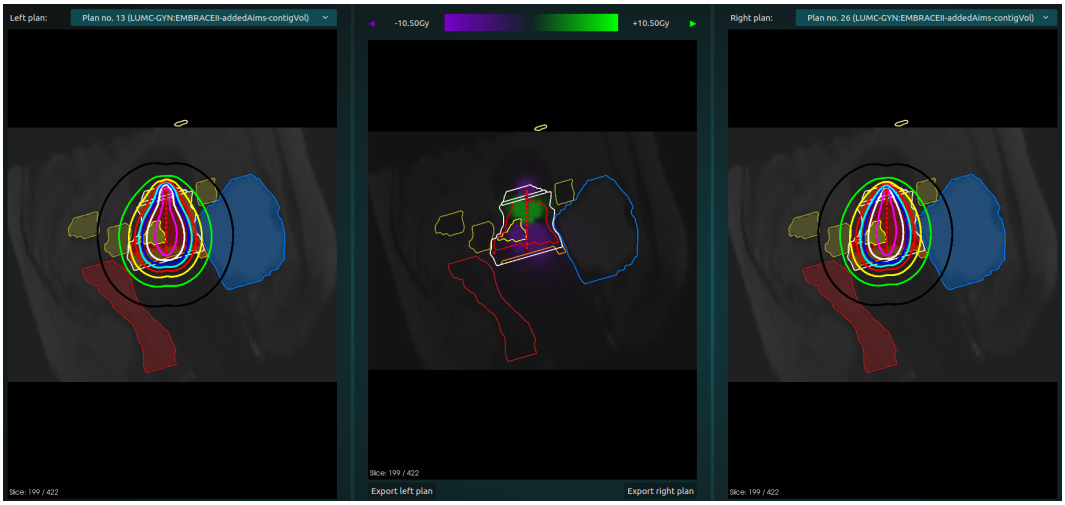

(f) Dose distribution comparison on coverage-oriented NTCP-preselected ”base” and ”guide” plans for patient 12.

Figure A.6: Additional dose distribution comparisons of NTCP for the BRIGHT-base plans on the left and guide-candidate plans on the right. The plans show their isodose lines, and the middle highlights their differences. The purple areas indicate a higher BRIGHT-base dose, and the green areas have a higher candidate-guide dose.