Evolutionary Optimization for Breast Cancer Brachytherapy Treatment Planning using BRIGHT

MO-RV-GOMEA in optimizing treatment plans for Internal Irradiation of Breast Tumors

Master Thesis

Daniela Toader

Delft University of Technology

TU Delft

Evolutionary Optimization for Breast Cancer Brachytherapy Treatment Planning using BRIGHT

MO-RV-GOMEA in optimizing treatment plans for Internal Irradiation of Breast Tumors

by

Daniela Toader

Student number

4904648

Thesis Committee:	Prof. dr. Peter A.N. Bosman, CWI, TU Delft, Chair
	Dr. Tanja Alderliesten, LUMC, Core member
	Dr. Annibale Panichella, TU Delft, Core member
	Dr. Anton Bouter, CWI
	MSc. Leah R.M. Dickhoff, LUMC
Faculty:	Faculty of Electrical Engineering, Mathematics and Computer Science, TU Delft

TU Delft

Abstract

This thesis utilizes Evolutionary Algorithms (EAs) within the BRIGHT framework for developing breast cancer brachytherapy treatment plans. We use expert knowledge and state-of-the-art EAs to formulate treatment planning as a multi-objective optimization problem whose solutions can be applied to actual patients. We propose four novel 2- and 3-objective formulations of this problem, which we implement within BRIGHT and empirically validate against anonymized data from 9 real-world patient cases. We demonstrate that all four formulations, under reasonable computational and time budgets, are capable of generating plans that match or exceed the properties of reference treatment plans. To verify the clinical relevance of our contributions, we rely on the expertise of a clinical expert who assesses whether the generated plans can be used in clinical treatment planning. The results of our empirical analysis show that 17 of the 18 plans presented to the expert are clinically acceptable and of immediate value to practitioners within the field of breast brachytherapy today.

Contents

1 Introduction

This research employs Artificial Intelligence (AI), and more specifically, Evolutionary Algorithms (EAs), to optimize treatment plans for patients undergoing partial breast irradiation (PBI). By leveraging the adaptive and robust nature of EAs, we can explore a vast solution space to identify the most effective treatment plan configurations. This integration of EAs in breast brachytherapy (BT) can enhance the efficiency and efficacy of radiation therapy by providing personalized treatment plans that improve patient outcomes and reduce the risk of complications [1].

This chapter builds up to the research questions that we seek to address within this work. To motivate this study, we begin by introducing how clinicians treat breast cancer and why brachytherapy is an appealing solution for treatments. Following this, we describe the conditions a treatment plan must meet to be clinically relevant. Finally, we explain how we aim to advance the state-of-the-art through automated EA plan generation.

1.1. Breast cancer and possible treatments

Breast cancer is a prevalent and potentially life-threatening condition characterized by the uncontrolled growth of cells in the breast tissue [2]. Effective treatment of breast cancer often requires a multi-faceted approach tailored to the specific characteristics of the tumor and the patient's overall health [3]. The primary treatment options for breast cancer include surgery, chemotherapy, radiation therapy, hormone therapy, and targeted therapy, each playing a critical role in managing the disease [3].

Surgery is often the first line of defense. It can involve either a lumpectomy, where only the tumor and a small margin of surrounding tissue are removed, or a mastectomy, which involves the removal of one or both breasts. Post-surgical treatments typically include radiation therapy and chemotherapy to eliminate any remaining cancer cells and reduce the risk of recurrence [3].

Radiation therapy is a key component of breast cancer treatment, delivered in various forms to maximize efficacy. External Beam Radiation Therapy (EBRT) is the most common modality, using high-energy X-rays or protons to target the tumor site while sparing surrounding healthy tissue [4]. Intraoperative Radiation Therapy (IORT) directly provides a single, concentrated dose of radiation to the tumor bed during surgery, effectively targeting residual cancer cells and reducing the need for prolonged postoperative treatment [4]. Brachytherapy, an internal radiation technique, plays a particularly significant role in partial breast irradiation, where the goal is to concentrate the radiation dose on the tumor while sparing surrounding healthy tissue. By placing radioactive sources directly within or near the tumor, brachytherapy allows for precise dose delivery, reduced treatment times, and a highly localized therapeutic effect. It is particularly effective for Accelerated Partial Breast Irradiation (APBI) and can be delivered in different dose rates, such as low-dose rate (LDR) or high-dose rate (HDR) forms. The real-world application we tackle in this study concerns APBI with interstitial multi-catheter brachytherapy (IMB). Studies, such as the Phase III multicenter trial comparing interstitial brachytherapy with external beam radiation therapy, have demonstrated the efficacy of brachytherapy in providing comparable local control and reduced toxicity [5]. This approach minimizes the risk of damage to adjacent tissues and organs while enhancing patient convenience [6, 7].

Chemotherapy is another key modality, employing cytotoxic drugs to target rapidly dividing cancer

cells. It can be administered as neoadjuvant therapy to shrink tumors before surgery, as adjuvant therapy to eliminate residual cancer cells after surgery, or as palliative therapy to reduce symptoms in advanced stages. Chemotherapy is typically delivered intravenously or orally and often involves combinations of drugs to improve efficacy and reduce the risk of resistance [8].

Among these treatment options, brachytherapy stands out for its precision and effectiveness, particularly when enhanced with automated treatment planning. This research focuses on optimizing breast brachytherapy using EAs to create individualized treatment plans. By leveraging EAs, this approach enables comprehensive exploration of potential treatment configurations, ensuring each patient receives a plan suited to their unique tumor characteristics and reducing potential toxicities.

Before delving into the details of how EAs generate suitable treatment plans, we first describe the status quo of formulating treatment plans and the criteria such plans must meet before being applied in the clinic.

1.2. Treatment for brachytherapy patients in the clinic

To better understand the problem, we discern two primary components of the brachytherapy treatment in the clinic: treatment planning and delivery. The treatment planning consists of a series of steps: the delineation of tumors and organs at risk (OARs) using imaging modalities such as CT and/or MRI, reconstruction of applicators and catheters, and dose calculation and computation of dose-volume histograms (DVHs) [9].

Advanced planning tools and algorithms are employed to delineate the breast volumes, reconstruct the applicators, and accurately predict dose distributions. Effective planning for breast brachytherapy is essential in clinical settings to ensure that the target area receives a uniform and adequate dose while reducing the risk of radiation-induced side effects. Such tools become invaluable in reducing the delay between imaging and treatment delivery [9].

The treatment planning process involves creating detailed clinical plans that specify radioactive sources' optimal placement and delivery time. These plans are designed to achieve precise dose distributions that concentrate the therapeutic dose on the tumor site and create steep dose gradients to protect nearby organs and tissues, such as the skin, lungs, and heart. The placements and delivery times are named dwell positions and dwell times, respectively. There are different fixed dwell positions in each of the inserted catheters, and, within one treatment, the longer the radioactive source resides at a position, the more dose is delivered at that site. Thus, treatment planning entails finding the best possible set of dwell times for each dwell position.

1.2.1. Clinical acceptance

For a treatment plan to be applied in the clinic, it must meet stringent clinical acceptance criteria, which are critical to ensuring the efficacy and safety of the treatment. These criteria include several Dose Volume Indices (DVIs), also called dose-volume histogram parameters, specific to breast APBI. Key DVIs encompass the coverage of the target volume (typically defined as the area of the tumor with a margin) by the prescribed dose, ensuring that a sufficient dose is delivered to the entire tumor region. Additionally, it is crucial to limit the maximum dose to critical and healthy structures, such as the skin, chest wall, and underlying organs, to minimize the risk of radiation damage and associated complications [10].

Homogeneity within the target volume is another critical factor for clinical acceptance. Dose homogeneity ensures that the radiation is evenly distributed throughout the tumor, reducing the risk of both underdosing (which can lead to ineffective treatment) and overdosing (which can cause unnecessary damage to healthy tissues). The formation of large blobs of high radiation doses, named hotspots, can commonly cause tissue damage. This research applies hotspot registration techniques while targeting reducing their volume during optimization. Furthermore, achieving homogeneity involves carefully balancing the radiation dose across the tumor volume, often requiring iterative adjustments during treatment plan optimization.

Clinical acceptance also involves evaluating the treatment plan against established protocols and guidelines, such as those provided by the American Brachytherapy Society (ABS) [10, 11] and other professional organizations. These guidelines offer benchmarks for acceptable dose distributions, critical structure sparing, and overall treatment quality. By adhering to these standards, clinicians can ensure that the treatment plans not only meet the technical requirements but also provide the best possible therapeutic outcomes for patients, maximizing the effectiveness of the treatment while minimizing

risks. In what follows, we formalize the procedure of automatically generating clinically acceptable PBI treatment plans through evolutionary optimization and formulate the main research questions we seek to answer in this thesis.

1.3. Problem statement

The complexity of creating optimal treatment plans for PBI poses significant challenges due to the intricate interdependencies between numerous variables, including tumor size, location, patient-specific anatomical considerations, and catheter placement.

Current manual and semi-automated planning methods often fall short of fully exploring the vast solution space, potentially leading to suboptimal treatment configurations. These limitations can result in less effective radiation delivery, higher risk of complications, and overall diminished therapeutic outcomes. There is a critical need for advanced methodologies to enhance the speed, efficiency, and personalization of treatment plans in breast brachytherapy.

This research addresses this gap by employing AI, specifically EAs, to optimize treatment plans for patients undergoing PBI. Treatment planning is a nuanced problem in which clinicians attempt to balance the conflicting objectives of delivering sufficient doses of radiation to cancer tissues while simultaneously sparing surrounding healthy organs. This makes treatment planning an inherently multi-objective (MO) problem. The motivation behind choosing EAs for treatment planning stems from over 3 decades of research showing that EAs are practical tools for solving such MO problems [12]. EAs are well-suited for this task due to their adaptive and robust nature, which allows for extensive exploration and exploitation of the solution space. Based on a state-of-the-art EA, we employ BRIGHT ("BRachytherapy via artificial Intelligent GOMEA-Heuristic based Treatment planning") to generate radiation therapy treatment plans. By integrating BRIGHT into the planning process, this study aims to identify the most effective treatment configurations, thereby enhancing the accuracy and efficacy of radiation therapy. The ultimate goal is to provide personalized treatment plans that significantly improve patient outcomes, shorten treatment times, and reduce the risk of complications, advancing the field of breast cancer treatment.

1.4. Research questions

To address the current problems in breast brachytherapy treatment planning, this research focuses on three main research questions:

Research questions

RQ1. How can we use BRIGHT for breast brachytherapy treatment planning?

RQ2. How do core parameter settings influence BRIGHT brachytherapy treatment planning outcomes?

RQ3. How do clinical experts judge the breast brachytherapy treatment plans optimized with BRIGHT?

Before addressing these research questions in detail, we first provide the necessary background context in Chapter 2. In Chapter 3, we divide each question into subquestions and present our methodology for addressing each subquestion. Chapter 4 displays the results of our empirical study on clinical patients in various optimization settings. Chapter 5 summarizes the research and discusses its most insightful findings. Finally, Chapter 6 concludes this work.

2 Background

To optimize radiation therapy treatment plans, radiation oncologists must dedicate significant time and effort to configuring dwell positions and times accurately. These parameters are critical for delivering precise and effective doses to the target areas while minimizing exposure to surrounding healthy tissues and organs at risk. This process is both time-intensive and complex in clinical practice, requiring the integration of anatomical, pathological, and dosimetric considerations [9, 13].

This chapter introduces the foundation on which we built our solution. We introduce the metrics that serve as bases for optimization criteria that evaluate the quality of generated plans. We follow this with an introduction of how EAs, and, more specifically, their multi-objective variants, can use these metrics to evolve clinically relevant treatment plans. Finally, we highlight one particular feature of treatment plans that current EAs struggle with, which we aim to improve in this work.

2.1. Automated Treatment Planning in the Clinic

This section outlines the clinical practices and aims within breast brachytherapy shared by our clinical expert. We define the breast target organs and the OARs and finally describe the DVIs accounted for during the treatment planning in the clinic.

2.1.1. Current Clinical Practice

In the clinic, following the catheter implants, CT scans of the patients are taken. Further, catheters are reconstructed within the Varian BrachyVision System, and ROIs are contoured. While some standard ROIs are automatically generated, others require separate delineations or adjustments. Experts create the so-called *ring structure*, a restricted volume of healthy tissue surrounding the CTVs to guide the optimizer. We define this region more rigorously in the following chapter. The clinical plans are obtained through dose calculation and optimization using Varian BrachyVision. The radiation source used during treatment is GammaMed Ir192 HDR Plus, Varian.

2.1.2. Breast Targets and OARs

The clinical target volumes (CTVs) for breast brachytherapy are designed to encompass the tumor bed with appropriate margins for microscopic disease spread. Various studies support the definition of the CTV as the normal breast tissue within 1–2 cm from the lumpectomy cavity edge limited by breast tissue extent [10, 14, 15, 16, 17]. Two primary target volumes are considered:

- Clinical Target Volume (CTV) at 1 cm (CTV_{1cm}): This target volume includes the tumor bed with a 1 cm margin to ensure the prescribed dose covers the area at risk of harboring microscopic disease.
- Clinical Target Volume (CTV) at 1.5 cm ($CTV_{1.5cm}$): This extended target volume encompasses the tumor bed with a 1.5 cm margin, accounting for potential microscopic disease spread beyond the immediate vicinity of the tumor bed.

It is critical to protect organs at risk from excessive radiation exposure to minimize potential complications. The primary OARs considered in this study include the skin, chest wall/pectoralis, lung, and

heart. Protecting the skin is essential to avoid radiation-induced dermatitis and ulceration [18]. The chest wall and pectoralis muscles need protection to prevent pain and functional impairment. Minimizing radiation exposure to the lungs reduces the risk of radiation pneumonitis. Similarly, avoiding excessive radiation to the heart is crucial to prevent cardiotoxicity, particularly in cases of left-sided breast cancer [10]. To this end, in clinical practice, experts create a ring structure that margins the target volumes. The ring becomes an effective technique for sufficiently constraining the dose to a prescribed volume, which in turn protects the OARs that fall outside the ring [11].

2.1.3. Clinical Aims and DVIs

Dose-volume indices (DVI) are essential metrics in evaluating the effectiveness of the brachytherapy treatment plan. The primary target of DVIs is ensuring that the tumor receives a dose sufficient to achieve therapeutic effectiveness while maintaining the doses to OARs within acceptable limits [9].

DVIs are metrics used to quantify the volume of a target or OAR receiving a specific radiation dose. As we further describe in more detail in Section 3.2, V_x is the volume percentage of the target or OAR receiving at least x% of the prescribed dose. D_x is the most irradiated x% of the target or OAR volume. The essential aims that make up our experts' clinical protocol are the following:

- $V_{95}^{CTV_{1cm}} > 95\%$;
- $V_{90}^{CTV_{1.5cm}} > 90\%$;
- Minimize V_{150} and V_{200} for the breast tissue;
- Constrain the PD not to exceed the so-called *ring structure*.

The experts we consulted minimize the V_{150}^{breast} and V_{200}^{breast} in combination with the coverage targets as a proxy for the Dose Homogeneity Index (DHI) in Equation 2.1.

$$DHI_{organ} = 1.0 - \frac{V_{150}^{organ}}{V_{100}^{organ}} \tag{2.1}$$

Experts can only calculate the DHI and assess the plan's clinical suitability after the optimization process has concluded. Through practical experience and joint research, experts have observed that maximizing the DHI reduces the chance of late toxicities. To this end, they have learned how to use their available optimization methods best to gauge this metric [19].

Regarding organs at risk, the most effective way found by the experts to spare the surrounding organs to the target is to create a ring structure around the $CTV_{1.5cm}$ [10]. They constrain the dose outside this ring to be less than the prescription dose.

2.2. Evolutionary Algorithms

Before delving into techniques that can help increase the quality of care patients receive in the clinic, we first need to establish an adequate framework to express them. Fundamentally, the problem of finding optimal treatment plans is a mathematical optimization problem. The scope of the problem encompasses *solutions* (treatment plans) that aim to treat the area affected by cancer while simultaneously minimizing damage to surrounding organs. To aid clinicians in finding adequate solutions, we require a strategy that is able to automatically explore the space of possible solutions and uncover high-quality plans. Over the years, researchers have developed many ways of tackling such problems, ranging from algorithms that consider inverse planning methods [13, 20] to more sophisticated methods that exploit the underpinning properties of optimization such as simulated annealing, linear programming and mixed integer linear programming [21, 22], and evolutionary algorithms [23]. Among these numerous options, we focus on one particularly promising paradigm inspired by nature itself – evolutionary algorithms (EAs).

EAs are computational methods inspired by the principles of natural selection and genetic evolution. These algorithms mimic nature's evolutionary biological processes, such as selection, crossover, and mutation. Through this process, EAs evolve a population of candidate solutions toward optimal individuals for complex problems. While EAs are rooted in biological inspiration, state-of-the-art EAs incorporate advanced techniques from fields such as statistical inference to enhance their performance. In this thesis, we focus on such a subclass of EAs that leverages statistical methods in its recombination operator. These methods involve estimating probability distributions over subsets of solution

variables, allowing for a more refined and efficient search of the solution space. This choice is motivated by recent research showing that this method works particularly well for the treatment planning problem [24]. We employ EAs through BRIGHT to automatically create a set of qualitative treatment plans for each patient, from which the clinical experts can select their preferred plan [19]. BRIGHT and EAs have also previously proven effective for prostate and cervix brachytherapy applications [25, 26], being clinically introduced at the Amsterdam University Medical Centers (Amsterdam UMC, location Academic Medical Center) for prostate cancer patients [25]. This work seeks to advance its application to breast cancer treatment.

2.2.1. GOMEA

Optimization strategies can be broadly categorized by the degree to which they exploit the underlying problem's structure. This taxonomy leads to three distinct cases. At one end of the spectrum, white-box (WB) approaches require a complete understanding of the problem's facets, which can be exploited to generate better solutions. At the opposite end, black-box (BB) methods assume only a rudimentary grasp of the problem and thus pose far less stringent constraints. Nestled between the two extremes is the class of so-called gray-box (GB) methods, where distinct components of the problem's structure can be separately exploited without requiring a whole (often untractable) model.

One of the leading advancements in this field is the Gene-Pool Optimal Mixing Evolutionary Algorithm (GOMEA). GOMEA introduces an innovative GB approach to recombination, focusing on mixing genes from a pool of solutions to create offspring that inherit statistically successful traits from their parents. In this case, the gray-box categorization stems from a statistical model that is inferred based on outstanding individuals in the population. In addition, the realization of the GOMEA framework is flexible and allows for the inclusion of domain-specific knowledge into its linkage model data structure.

The gene-pool optimal mixing routine [27] is performed for each solution in the population as a standalone variation operator and works as follows. Each solution is first cloned into an identical copy. Then, following a model of variable correlation called a *linkage model*, different sets of variables are iteratively considered. For each set, GOM selects a parent at random from the previous population and swaps the selected variables of the offspring with that of the *donor*. Based on whether the change of variables improves the quality of the solution, GOM either accepts or discards the swap. Though GOM alone does not guarantee solution convergence [27], further auxiliary enhancements such as Forced Improvement (FI) can provide this property [28]. Overall, GOM enhances the exploration and exploitation capabilities of the algorithm, making it particularly powerful for solving complex optimization problems with intricate variable dependencies.

In this work, we employ the Multi-Objective Real-Valued Gene-Pool Optimal Mixing Evolutionary Algorithm (MO-RV-GOMEA) [24], a state-of-the-art algorithm that targets multi-objective problems, estimating probability distributions over subsets of solution variables and exploiting interdependencies between them. We describe multi-objective real-valued optimization and MO-RV-GOMEA in further detail in Section 2.2.2.

2.2.2. Multi-Objective Optimization

Multi-objective (MO) optimization describes a class of mathematical optimization problems in which multiple objective functions are to be optimized simultaneously. These problems arise in many areas of science and engineering, where complex real-life problems can be modeled as multiple functions with often conflicting objectives. Before surveying the landscape of EAs that target MO optimization, we first establish the foundational properties of such problems.

Formally and without loss of generality, a multi-objective optimization problem can be formulated as described by Equation 2.2. That is, the optimal solution(s) x belong(s) to the set of all possible solutions $\mathcal{X}$ whose objective values are the real-valued vector $\mathbf{f}(x)$, where each position of the vector is the numerical value of one objective. Naturally, for the problem to be multi-objective, the number of objectives $n \geq 2$.

$$\min_{x \in \mathcal{X}} \left(f_i(x) \right), \; f_i \in \mathbb{R}. \tag{2.2}$$

Unlike their single-objective counterparts, MO optimization problems seldom have a single solution that simultaneously excels in all objectives. For this reason, the comparison of competing solutions is often carried out in terms of *Pareto dominance* and *Pareto (approximation) fronts*. Given any two

solutions x_1 and x_2, x_1 is said to *dominate* x_2 if and only if both cases of Equation 2.3 hold. More intuitively, x_1 dominates x_2 if it is *at least as good* as x_2 in all objectives and *strictly better* than x_2 in at least one objective. A solution $x \in \mathcal{X}$ is said to be Pareto-optimal if and only if $\nexists x' \in \mathcal{X}$ such that x' dominates x. The set of all Pareto-optimal solutions is called the Pareto *front* or Pareto *frontier*. Obtaining a good approximation of this set is most often the goal of MO optimization strategies.

$$\begin{cases} \forall i \in \{1, \ldots, n\},\ f_i(x_1) \leq f_i(x_2), \\ \exists i \in \{1, \ldots, n\},\ f_i(x_1) < f_i(x_2). \end{cases} \tag{2.3}$$

Over the years, several multi-objective evolutionary algorithms have emerged that exploit problem knowledge differently. Discrete optimization MOEAs include NSGA-II [29] and MO-GOMEA [30], while MAMaLGaM [31] addresses continuous optimization. More recently, Bouter et al. [24] have introduced MO-RV-GOMEA, an extension of the discrete MO-GOMEA that adapts elements of prior algorithms to better suit continuous optimization problems. We briefly describe the core attributes of MO-RV-GOMEA before discussing the work of Bouter et al. [24].

2.2.3. Real-Valued Optimization

MO-RV-GOMEA [24] proceeds in a loop that consists of selection, clustering, model learning, and variation. Following a standard truncation selection procedure, solutions are bundled into distinct clusters to improve the spread of individuals across the Pareto approximation front. The algorithm first establishes single-objective clusters that outwardly expand the scope of the Pareto approximation front before heuristically selecting cluster centers that increase sparsity. After partitioning the population into non-overlapping clusters, a Family of Subsets (FOS) linkage model is learned for each cluster. Assigning models per cluster accounts for differences in the dependency structure of underlying problem variables, which are likely to emerge in MO optimization. Once established, the linkage model serves as a basis for variation. This borrows techniques for Estimation of Distribution Algorithms (EDAs) but is tightly integrated into the GOM routine.

The RV-GOM step estimates a multivariate Gaussian distribution that provides samples, which are the basis of a partial update to each solution in the cluster. Like MAMaLGaM [31], variation scaling techniques control the rate at which solutions evolve. These mechanisms include adaptive variance scaling and anticipated mean shifts. Finally, the algorithm includes a forced improvement mechanism to update solutions that have not shown Pareto improvement in several generations. To achieve this, MO-RV-GOMEA attempts to shift individuals towards the best solutions associated with their respective clusters.

2.3. Automated Treatment Planning in BRIGHT

BRIGHT is an optimization tool developed to improve the quality and efficiency of brachytherapy treatment planning. It uses a multi-objective real-valued evolutionary algorithm (MO-RV-GOMEA) to generate high-quality treatment plans that balance tumor coverage and sparing of healthy tissues [1].

The DVI aims concern both the target volumes and the organs at risk and are defined to ensure adequate coverage and sparing. These aims are incorporated into the optimization model through the BRIGHT objective functions to guide the generation of treatment plans. To determine the quality of a plan, we define the Golden Corner – a specific region where all protocol aims are satisfied. Figure 2.1 illustrates the Golden Corner principle with a two-objective approximation front example.

The weights w_c and w_s are assigned to prioritize the most violated DVIs, ensuring that optimization efforts focus on the most critical areas. The least violated DVI receives a weight of 1, and the weights increase exponentially by a factor of 10 for consecutive violations, normalized between 0 and 1. The optimization of the DVIs continues past the point when the aims are reached – the stopping criterion is the time budget. Further optimizing the indices allows BRIGHT to improve the plan configurations as much as possible within the computational budget.

2.3.1. Tumor Coverage

The first objective of BRIGHT is to ensure adequate tumor coverage. This is quantified using the Least Coverage Index (LCI), which measures how well the dose covers the clinical target volumes. The LCI for a plan p is defined as in Equation 2.4:

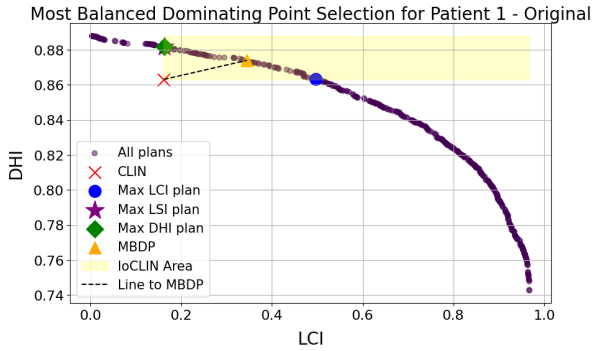

Figure 2.1: Pareto approximation front example with two objectives: Least Coverage Index (LCI) and DHI (Dose Homogeneity Index). We consider the plans with LCI>0.17 and DHI>0.86 to be in the Golden Corner.

$$LCI_w(p) = \sum_{a \in \text{coverage aims}} w_c \delta_c(DVI^a) \tag{2.4}$$

where $\delta_c(DVI^a)$ represents the difference between the current dose-volume index value and its target, and w_c is the weight assigned to each coverage aim. The weights adapt during optimization and are designed to prioritize the most violated DVIs, ensuring that optimization focuses on achieving the necessary coverage.

2.3.2. Healthy Organ Sparing

The second objective focuses on sparing healthy organs from excessive radiation. This is quantified using the Least Sparing Index (LSI), which measures how well the plan avoids unnecessary radiation to OARs. The LSI for a plan p is defined as in Equation 2.5:

$$LSI_w(p) = \sum_{a \in \text{sparing aims}} w_s \delta_s(DVI^a) \tag{2.5}$$

where $\delta_s(DVI^a)$ represents the difference between the current DVI value and its sparing aim, and w_s is the weight assigned to each sparing aim. Similar to the LCI, the weights are assigned to prioritize the most violated DVIs, ensuring that the optimization also considers the sparing of healthy tissues.

2.4. Homogeneity in Treatment Planning

The distribution of radiation doses throughout the target organ can have significant consequences for the patient's well-being. To ensure that the risk of tissue damage is minimized, radiation delivery techniques have to meet certain criteria. One such crucial criterion is *homogeneity*. Homogeneity refers to how sharp the changes in radiation doses are between neighboring regions. Homogeneous treatment plans deliver radiation in doses that resemble uniformity, with smooth transitions between close points. By contrast, heterogeneous plans have sharp transitions between adjacent regions. These differences lead to so-called *hotspots*. In this section, we elaborate on why homogeneous plans are important for treatment plan quality and link their underlying properties to the underpinning evolutionary algorithm.

2.4.1. Challenges of Achieving Homogeneous Treatment Plans

Non-homogeneous treatment plans can lead to the formation of hotspots, which are contiguous regions receiving excessively high radiation doses. Hotspots pose significant risks, including the potential for radio-necrosis, where radiation irreversibly damages healthy tissue, which can result in severe side effects. For cancers such as prostate, the damage can result in bladder dysfunction [32], while for breast cancer cases, radio necrosis can cause chronic pain, potentially difficult to manage, significantly impacting the patient's quality of life [33]. Long-term effects of breast radiation can lead to fibrosis, where the breast tissue becomes stiff and less flexible. This can cause chronic pain and limit the movement of the affected area [34]. Moreover, hotspots compromise the overall safety and effectiveness of the treatment.

Another challenge associated with heterogeneous treatment plans is their susceptibility to disturbances and uncertainties, such as patient movement, anatomical variations, and inaccuracies in catheter placement or organ delineation. Heterogeneous plans with pronounced peaks and valleys in dose distribution are more vulnerable to these uncertainties, potentially leading to suboptimal treatment outcomes [35].

2.4.2. Current Optimization Techniques and Their Limitations

Traditional optimization techniques in HDR brachytherapy include manual adjustments and the use of linear penalty models (LPMs) to optimize dwell times and positions of radioactive sources. Methods like HIPO (hybrid inverse treatment planning optimization algorithm) [13] and IPSA (Inverse Planning Simulated Annealing) [20] automate and improve the planning process [13, 20]. These methods aim to meet clinical protocol objectives by achieving prescribed DVIs for target and healthy tissues. However, these models use a simplification of the definition of DVIs since they use gradients. Therefore, they do not directly optimize them. Furthermore, DVIs are aggregate measures that may not adequately capture dose contiguity, leading to the potential formation of hotspots even in plans that meet DVI criteria [36].

2.4.3. Contributions to Enhancing Homogeneity

Commandeur's research [35] builds on BRIGHT by addressing treatment plan homogeneity. Commandeur discusses metrics such as the Hotspot Size Index (HSI) and the sum of extra V indices to measure and reduce contiguous high-dose regions within treatment plans. By incorporating these metrics into BRIGHT's optimization process, Commandeur's work aimed to generate more homogeneous treatment plans that minimize hotspots while adhering to clinical protocol goals.

To detect hotspots effectively, Commandeur developed a hotspot registration method based on connected component analysis using the graph-based Afforest algorithm. This method identifies contiguous high-dose regions (hotspots) by treating the dose distribution as a graph, where nodes represent dose calculation points and edges represent dose contiguity.

The Afforest algorithm is a scalable, efficient method for detecting connected components in a graph. Each dose calculation point (DCP) within the target and sparing volumes is considered a node in the graph for hotspot registration. Edges between nodes are established based on spatial proximity and dose levels, ensuring that nodes with high spatially close doses form a connected component. The algorithm proceeds with initialization: each DCP is initially its own component. Further, random edges are sampled and then processed to merge the connected components of the nodes they link. The merging step involves iteratively uniting the nodes that share high dose values and are spatially close into larger connected components. Once all relevant edges are processed, the remaining components represent the hotspots.

This algorithm efficiently identifies large contiguous high-dose regions, enabling precise quantification of hotspots. The Hotspot Size Index (HSI) is then calculated to guide the optimization process, ensuring that the algorithm targets and mitigates these hotspots, leading to more homogeneous dose distributions.

2.4.4. Broader Implications

The improvements in dose uniformity achieved through Commandeur's work [35] have implications beyond prostate HDR brachytherapy. Similar challenges in achieving dose uniformity exist in other cancer treatments, such as breast brachytherapy. The techniques developed can also be adapted and applied to improve treatment outcomes in these contexts.

Improving dose homogeneity in breast brachytherapy treatment plans is essential for maximizing therapeutic effectiveness and minimizing adverse side effects. In this research, we adopt the same hotspot detection algorithm and analyze its suitability for breast cancer treatment planning.

3

Research Method

Automated treatment planning in breast brachytherapy leverages algorithms to optimize the placement and duration of radiation sources at given sites. The primary objective is to speed up the planning process, potentially improve the dose distribution calculation, and finally speed up treatment time. These goals concern clinicians, as their workflow can be eased and sped up, and patients, as the time spent in a brachytherapy treatment setting is reduced, respectively.

While placement optimization can be done, and related research has explored catheter placement optimization and dwell positions autoactivation [37, 38], this work focuses on an equally important and complementary problem: duration optimization. This chapter outlines the steps necessary to apply BRIGHT to the breast case to achieve these goals in a practical and advantageous way. We address the overarching method in five steps. Firstly, we describe the anonymized patient data and the techniques that prepare it for optimization in Section 3.1. Secondly, we report the clinical aims defined by the breast protocols devised by clinical experts in Section 3.2. We use these targets to construct and adjust the BRIGHT optimization functions to further optimize with GOMEA within BRIGHT. Third, we outline the experimental setup employed to evaluate the results of applying BRIGHT to the breast case in Section 3.3. To prune the results GOMEA generates, we define a plan selection scheme that we use to reduce the number of plans we analyze together with clinician experts in Section 3.4. Finally, we use the form displayed in Section 3.5 to assess the clinical feasibility of the generated plans.

3.1. Data and Expert Knowledge

The data in this study consist of anonymized breast brachytherapy cases belonging to nine patients with differing anatomical topologies. Each case pertains to a size category based on the breast volume within the scan.

Each case receives a unique identifier based on case export time used throughout this study: p1, p2, p3, p4, p5, p6, p7, p8, and p9 (Patients 1 through 9). The prescription dose used during the radiation treatment for p1 is 4.3 Gy, while for all other patients, the prescription dose is 7.5 Gy.

We also note that Patient 8 and Patient 9 have suboptimal implant configurations, according to our clinical expert. This observation suggests that their cases will be more challenging to optimize. Each patient has associated CTV_{1cm} and $CTV_{1.5cm}$ volumes and the number of dwell positions, which vary significantly across the dataset. Table 3.1 shows the aforementioned metrics for each patient in our dataset.

For CTV_{1cm}, the median volume is 51.6 cm^3, with a range of 73.1 cm^3, spanning from 28.9 cm^3 (Patient 3) to 102 cm^3 (Patient 5). For $CTV_{1.5cm}$, the median volume is 87.4 cm^3, with a range of 110.6 cm^3, spanning from 53.4 cm^3 (Patient 3) to 164 cm^3 (Patient 5). The number of dwell positions also demonstrates variability, with a median of 134 positions and a range of 103 positions, spanning from 87 (Patient 3) to 190 (Patient 5). The metrics highlight the differing tumor sizes, corresponding to differing anatomical considerations and treatment complexities. Such diversity stresses the importance of individualized treatment planning, which we enable with BRIGHT and describe in Section 3.2.

This research was conducted in collaboration with an expert from Virginia Commonwealth University, whose contribution covered evaluating breast brachytherapy treatment plans generated by BRIGHT.

Table 3.1: Summary of CTV volumes and dwell positions for all patients.

Patient	CTV_{1cm} Volume	$CTV_{1.5cm}$ Volume	Number of Dwell Positions
Patient 1	95.1 cm^3	138 cm^3	175
Patient 2	37.4 cm^3	69.1 cm^3	102
Patient 3	28.9 cm^3	53.4 cm^3	87
Patient 4	51.6 cm^3	87.4 cm^3	111
Patient 5	102 cm^3	164 cm^3	190
Patient 6	58.2 cm^3	99.5 cm^3	134
Patient 7	50.4 cm^3	87.1 cm^3	144
Patient 8	59.4 cm^3	92.6 cm^3	143
Patient 9	42.1 cm^3	77.1 cm^3	102

The grounds for the evaluations are part of the experimental setup in Section 3.3.

3.2. Enabling BRIGHT to Optimize Breast Plans

This section addresses the first main research question: *How can we use BRIGHT for breast brachytherapy treatment planning?*

Research questions

RQ1. How can we use BRIGHT for breast brachytherapy treatment planning?

- RQ1.1. What particularities of the data should we account for prior to integrating with BRIGHT? How is the dataset different from the other instances previously tackled with BRIGHT?
- RQ1.2. What steps are necessary to enable BRIGHT to generate treatment plans for breast brachytherapy?
- RQ1.3. How can we obtain clinically acceptable dose distributions using BRIGHT for breast?
- RQ1.4. How can we obtain homogeneous dose distributions when optimizing with BRIGHT?

In the following subsections, we describe the implementation details and additions that help answer the sub-questions and, ultimately, RQ1.

3.2.1. Defining the Breast ROIs

Different ROIs within the breast and surrounding organs are identified through high-resolution imaging modalities, in this case, CT scans. These imaging techniques provide detailed anatomical data for delineating the tumor bed and the surrounding critical structures.

Defining the Breast Targets and OARs

The clinical target volume is outlined with a margin, typically 1 to 1.5 centimeters, to ensure thorough coverage of potential microscopic disease spread. For the breast case, the clinical aims concern the (CTV_{1cm}) and ($CTV_{1.5cm}$), which we use as targets in BRIGHT. In the clinic, the dwell activation area is a volume that identifies the region where dwell positions can be activated automatically in the BrachyVision system during planning. The activation area is necessary for the system's optimizer to enable dwells. However, BRIGHT does not directly require this structure to activate dwells. We initially allow BRIGHT to consider all dwell positions as active. This approach works well for all patients except Patient 2, for whom a few dwell positions were activated too far from the target. To this end, we deactivate certain dwell positions that exceed a certain margin from our breast ROIs.

Moreover, a ring structure is created around the activation area to protect OARs, such as the skin, chest wall, and pectoralis. This ring acts as a buffer zone, constraining the radiation dose outside this region to acceptable levels. As it has proven effective in clinical practice to ensure dose conformity, we import this structure into BRIGHT.

Defining the Breast Avoidance ROI

To answer RQ1.1, we note that the breast organ is significantly larger than other organs for which BRIGHT has been previously used. An example of a significantly smaller organ BRIGHT has been used to optimize plans for is the prostate [39]. The optimization relies on DVIs, which are:

- $D_v^o =$ dose index; the minimum dose received by the most irradiated sub-volume v of ROI o.
- $V_d^o =$ volume index; the sub-volume of ROI o that receives *at least dose* d.

To drive the optimization process, we utilize methods that can efficiently and accurately estimate DVIs through sampling. This procedure works as follows. First, we discretize the ROI into a set of randomly generated *dose calculation* (DC) points. These points are then subject to a Monte Carlo sampling process that approximates the ROI and computes a vector of values D with entries d_i that encode the total radiation dose received at the DC point with index i. We then link each discrete point to the ROIs it belongs to, aggregating points into collections D_o for each ROI o. We evenly divide the sampled points across ROIs to avoid biasing particular regions. An exception to this is the avoidance area ROI, which receives double the number of DCPs due to its larger size when computing its hotspot volume. This choice allows more precise connected component analysis during the hotspot volume calculations.

The discretization and sampling processes are crucial for estimating the radiation profile of the solution. Based on the sampled points and their distribution between ROIs, we compute two values describing this profile: the volume and dose indices.

Equation 3.1 defines the Volume Index (VI), with the S, O superscript indicating that the points are sorted on dose. The VI of a given ROI is o is the fraction of DC points in ROI o that receives at least a dose d. This quantity describes the degree to which the radiation covers the region, estimated on the evenly sampled DC points.

$$V_d^o = \frac{\|\{\boldsymbol{d}_i \in \boldsymbol{D}^o | \boldsymbol{d_i} \geq d\}\|}{|\boldsymbol{D}^o|} \tag{3.1}$$

Equation 3.2 describes the Dose Index (DI). Intuitively, the dose index tracks the minimum dose received at any sampled set for a given ROI o. This metric indicates whether the amount of radiation received by a specific ROI is appropriate, insufficient, or excessive.

$$D_v^o = \frac{\boldsymbol{d}^{S,O}_{\lfloor (v/v^o) \rfloor}}{PD} \tag{3.2}$$

We define the breast avoidance region of interest to improve the optimization results and the computational costs. The avoidance ROI structure is accomplished by merging the ring structure with the dwell activation area. The merging process retains all relevant geometrical information, thus rigorously delineating the avoidance area for subsequent optimization steps.

Integrating the contours of the activation area and the ring structure into a complete avoidance ROI allows BRIGHT to focus the optimization efforts on regions of utmost importance without compromising solution quality. This is possible as the metrics we compute across the avoidance region will be valid for the entirety of the breast structure, thanks to the strict ring-sparing constraint. For instance, no hotspots should occur within the breast outside the ring, as very high radiation doses should not pass this delineation. Constraining the V100% within the ring would result in the dose outside of it being too low for a point to be considered a hotspot. While this constraint is respected during optimization, V100% within the ring slightly higher values can result after re-evaluation on more DCPs. We used a limit of 1 cm^3 to achieve values similar to those observed in clinical plans.

We introduce a margin to the ring and the avoidance ROI of 15cm to tackle the aforementioned dwell activation outside of the avoidance ROI. This method allows for the deactivation of three dwell positions deemed too far from the avoidance structure.

Another particularity of the breast data is its potential lack of breast delineations. In practice, this makes the boundaries of the organ unclear and makes demarcation challenging. This observation strengthens the necessity of the merged avoidance area since other means of determining the optimization area can become very costly in computation efforts. For example, the optimization region could be determined by computing the area between the skin organ and the pectoralis, which requires

sampling dose calculation points across all mentioned structures. To circumvent the extra computation effort, we take advantage of the avoidance areas, which clinical experts generate routinely, unlike breast delineations.

Furthermore, upon automatically generating the avoidance structures, experts follow up by adjusting the structures so that the optimizer results match their expectations and goals best. This common practice can introduce a lot of manual labor. For this reason, we also allow optimization aims to be defined in terms of the unadjusted contours we have been provided with Patient 2.

3.2.2. BRIGHT Protocols for Breast

To address diverse clinical requirements, we define different protocols. These protocols determine the treatment plan aims and can be standardized clinical constraints or consensus measures based on expert groups' recommendations such as the American Brachytherapy Society [10].

Such protocols vary in their approach to sparing certain regions and optimizing specific indices such as $V_{95}^{CTV_{1cm}}$. For instance, the protocol in Table 3.2 focuses on sparing the ring region while optimizing the DHI, aiming to minimize radiation exposure to the ring structure and for a homogeneous dose distribution within the target volume. Another protocol may prioritize optimizing the DHI without specifically emphasizing sparing the ring region, focusing solely on achieving the most homogeneous dose distribution possible.

Optimizing metrics other than DVIs in the clinic is not generally an option through the generic optimizers embedded into the radiation machinery in clinically used treatment planning software. For this reason, experts attempt to achieve the best DHI by proxy by maximizing the $V100\%$ (i.e., the volume receiving PD) in the target volume while minimizing $V150\%$, (i.e., the volume receiving 150% of the PD). In BRIGHT, we can directly calculate and optimize the clinical goal metric, DHI, without needing a proxy. We note that although promising, this scalarization approach is imperfect. One drawback of the DHI is its limited capacity to capture nuanced problem details, such as spatial properties and structure.

The initial protocol aims used in BRIGHT for the breast experiments are as follows:

Targets (" $>$ ")	Organs at risk (" $<$ ")
$V_{95}^{CTV_{1cm}} > 95\%$ $V_{90}^{CTV_{1.5cm}} > 90\%$	$D_{0.1}^{ring} < 100\%$

Table 3.2: Initial protocol used in the objective function calculations with a prescribed dose of 7.5 Gy (4.3 Gy for p1)

Targets (" $>$ ")	Organs at risk (" $<$ ")
$V_{95}^{CTV_{1cm}} > 95\%$ $V_{90}^{CTV_{1.5cm}} > 90\%$	$V_{100}^{ring} < 1cm^3$

Table 3.3: Reworked protocol used in the objective function calculations with a prescribed dose of 7.5 Gy (4.3 Gy for p1). The targets comprise the LCI objective, while the OARs comprise either a hard constraint or a third objective (LSI) during the optimization. In addition, the second objective consists of either DHI or hotspot volume.

Following consultations with our clinical expert, we adjusted the OAR aim to constrain the volume receiving V100% in the ring to $1cm^3$. Further iterations uncovered that this limit would vary across patients, thus requiring a new approach to determine a reasonable bound. The protocol used in the expert evaluation is shown in Table 3.3. The targets aid in computing the LCI objective, while the OAR aids in computing the LSI objective or becomes a hard constraint based on the choice of 3- or 2-objective optimization (choice made when declaring the protocol).

3.2.3. Dose Homogeneity Index

A homogeneous dose distribution aims to prevent localized regions of excessive radiation that could lead to increased toxicity. A homogeneous dose distribution in the context of radiation therapy, includ-

ing breast brachytherapy, refers to the even distribution of the prescribed radiation dose throughout the target volume. Ideally, the goal is to deliver a uniform dose to the tumor or treatment area while protecting OARs.

The DHI objective aims to achieve dose homogeneity within the target area, which is essential for avoiding hotspots that can lead to adverse side effects. For this purpose, we use the definition in Equation 2.1. The goal is to maximize the DHI, with an ideal value of 1.0, indicating perfect homogeneity. In practice, a DHI of 0.75 or higher is clinically acceptable, with values above 0.85 being preferred according to clinical standards [10].

We follow the method of [40] and define two added targets in BRIGHT that allow for computing the DHI: V_{100}^{avoid} and V_{150}^{avoid}. The $avoid$ region of interest consists of the merged avoidance regions (the dwell activation area and the ring structure).

3.2.4. Contiguous Volumes

While DHI provides a general assessment of dose uniformity, it may be insufficient in specific cases, particularly for high-dose-rate brachytherapy in breast cancer treatment. The DHI treats dose variations uniformly across the entire volume without considering the spatial arrangement of these dose variations. This lack of spatial information makes assessing the clinical risks associated with the dose distribution challenging. For instance, a plan with a high DHI might have an acceptable level of overall homogeneity. However, it could still have small, intense hotspots near sensitive structures that have all of the high-dose regions concentrated in close proximity, increasing the risk of tissue damage or necrosis. The Contiguous Volume Analysis (CVA) method, as proposed by Commandeur [35], addresses this issue by detecting connected high-dose regions (hotspots) and minimizing their size and volume, which the DHI alone cannot account for.

We incorporated a previously developed hotspot volume detection algorithm into our optimization process to minimize regions of excessive radiation dose within treatment plans. The algorithm is outlined in the pseudocode in Figure 3.1 and consists of two key stages: label initialization and connected component analysis. The `LabelInitialization` stage (line 2) iterates over DCPs in a parallel fashion and marks them with one of two labels. DCPs exceeding a predefined dose threshold for a given OAR are labeled with their corresponding indices (line 9), while those that fall beneath the threshold are marked with a placeholder value (e.g., -1 line 11), which establishes the initial identification of potential hotspots based on radiation dose thresholds.

The connected component analysis stage takes place in the `Afforest` routine (line 14) and leverages a parallelized union-find structure to group contiguous DCPs into clusters corresponding to distinct hotspots. The algorithm efficiently identifies connected regions of excessive dose by iteratively merging and compressing these labels across neighboring DCPs (lines 21, 28) by means of the `FindAndCompress` procedure (lines 35-42). For each cluster, the volume of the hotspot is computed by summing the contributions of all associated DCPs. Hotspots that meet a predefined minimum volume threshold (e.g., clinically significant regions larger than 0.1 cm^3) are considered for further evaluation.

The hotspot volumes (HVs), or contiguous volumes (contVs), are then integrated into the treatment plan optimization process in the objective function. Minimizing the HV during optimization incentivizes the algorithm to reduce hotspot sizes while maintaining adequate dose coverage for the CTVs. Incorporating this algorithm allows the optimization framework to directly address dose homogeneity and prioritize spatial properties of the dose distribution.

Managing the dose distribution within contiguous volumes around the target is essential to ensure dose homogeneity and effective treatment. We employ the method from [35] to detect contiguous volumes (contVs or HVs) of high doses through connected component analysis and hotspot detection as shown in 3.1. We experiment with different minimum hotspot doses, lower bounds for hotspot volumes, and factor edge length. The factor edge length requires a careful configuration, as this parameter can influence what makes or breaks a hotspot. It is considered separately in an experiment from Appendix B. In the given clinical setting, the minimum dose for considering hotspots for the breast case is generally 150% of the prescription dose. We further compare this heuristic to different contiguous volume configurations in our experiments.

Algorithm 2 Hotspot Detection

```
1:
2:  procedure LabelInitialization(D, R_lb,o)        // D = Set of DCPs per treatment plan
3:                                                  // R_lb,o = hotspot radiation lower bound per organ o
4:      L ← []                                      // Array of DCP labels per solution
5:      for all D_x ∈ D  in parallel do
6:          for all d ∈ D  in parallel do
7:              if R(d) ≥ R_lb,o then               // check against radiation lower bound of the organ
8:                                                  // that the DCP is sampled in
9:                  L_x[d_i] ← d_i                  // d_i = index of DCP d
10:             else
11:                 L_x[d_i] ← −1
12:     return L
13:
14: procedure Afforest(L, E, #rounds)               // L = labels of DCPs per treatment plan
15:                                                 // E = set of edges between DCPs
16:     for i ← 1, #rounds do
17:         for all L_x ∈ L  in parallel do         // L_x = labels for treatment plan x
18:             for all l ∈ L_x : i ≤ ||N(l)|| in parallel do
19:                 union_async(l, N(l)_i, L_x)     // N(l) = neighbours of DCP l
20:             for all l ∈ L_x  in parallel do
21:                 FindAndCompress(l, L_x)
22:     c_x ← most_most_frequent_component(L_x)
23:     for all L_x ∈ L  in parallel do
24:         for all l ∈ L_x : l ≠ c_x  in parallel do
25:             for i ← #rounds, ||N(l)|| in parallel do
26:                 union_async(l, N(v)_i, L_x)
27:             for l ∈ L_x  in parallel do
28:                 FindAndCompress(l, L_x)
29: procedure union_async(u, v, L_x)                // p_u = parent label of DCP u
30:     p_u ← FindAndCompress(u, L_x), p_v ← FindAndCompress(v, L_x)
31:     while p_u ≠ p_v do
32:         if p_u = L_x[p_u] and CAS(&P[u], p_u, p_v) then
33:             return
34:         p_u ← FindAndCompres(u, L_x), p_v ← FindAndCompress(v, L_x)
35: procedure FindAndCompress(u, L_x)
36:     r ← u
37:     if L_x[r] = r then
38:         return r
39:     while r ≠ L_x[r] do
40:         r ← L_x[r]
41:     while j ← L_x[u] > r do
42:         L_x[u] ← r, u ← j
```

Figure 3.1: Hotspot Detection method, taken from [35]. Used to calculate and minimize the hotspot volume during optimization with BRIGHT.

3.3. Experimental Setup

This study uses a computational analysis approach to determine suitable parameter configurations for BRIGHT. Additionally, it uses a clinical analysis approach to determine whether treatment plans created by BRIGHT can be considered clinically acceptable. We describe the experimental setup in terms of BRIGHT parameter settings and the chosen objective functions. This section's goal is to address RQ2 and RQ3, for which we divide them into sub-questions:

Research questions

RQ2. How do core parameter settings influence BRIGHT brachytherapy treatment planning outcomes?

- RQ2.1. How much optimization time do we need to obtain protocol-acceptable plans?
- RQ2.2. How many DCPs do we need to obtain plans with accurate dose calculations?
- RQ2.3. How do different optimization functions compare in terms of objective values and protocol conformity?

The specifications of the servers used in our experiments are detailed in Table 3.4. **2O** refers to the server used for experiments with two objectives, while **3O** refers to the server used for experiments with three objectives.

Component	**Server 2O**	**Server 3O**
CPU	Intel Xeon Bronze 3206R, 1.90 GHz, 16 cores	Intel Xeon E5-2630 v4, 2.20 GHz, 20 cores
RAM	780 GB	1055 GB
GPU Model	NVIDIA RTX A5000	NVIDIA TITAN X (Pascal)
GPU Memory	24564 MiB	12288 MiB
Driver Version	525.89.02	560.35.03
CUDA Version	12.0	12.6

Table 3.4: Hardware specifications of the servers used for the experiments.

3.3.1. Computational Experiments

BRIGHT utilizes several optimization parameters to achieve its aims effectively, including MO-RV-GOMEA-specific and brachytherapy-specific parameters. A significant number of parameters are set to the defaults in BRIGHT, while some have been varied or modified. These parameters include the optimization time budget and the number of dose calculation points. When enabling the contiguous volumes technique to minimize the hotspot volume, we vary the minimum hotspot dose, lower bound for hotspot volumes, and the factor for the edge length. To test the robustness of the BRIGHT results, we put the consistency of our resulting elitist solutions to the test. We run each experiment configuration ten times to visualize the solution variance across the approximation front.

We analyze the solutions generated by BRIGHT using the different protocols and optimization functions under various time constraints and illustrate the used settings in Table 3.5 and Table 3.6.

Optimization	**Time (s)**	**DCPS**
`2O minV150`	600	250k
`2O DHI`	30, 60, 180, 300, 600, 900	250k
`3O DHI`	30, 60, 180, 300, 600, 900	250k
`2O contV`	30, 60, 180, 300, 600, 900	10k, 20k, 50k, 100k, 200k, 250k, 300k
`3O contV`	30, 60, 180, 300, 600, 900	10k, 20k, 50k, 100k, 200k, 250k, 300k

Table 3.5: Optimization setup parameters: Time and DCPS for all optimization techniques.

Based on the results of the computational experiments, we choose the most suitable technique (i.e., objectives and algorithm settings), which we employ in the following steps. We apply a selection scheme so we can analyze our technique together with our experts.

Optimization	Min Hotspot Dose (%)	LB Hotspot Volume (cc)	Factor Edge Length
`20 contV`	110, 130, 150, 170, 200	0.05, 0.1, 0.2, 0.5	1, 2, 3, 5, 10
`30 contV`	110, 130, 150, 170, 200	0.05, 0.1, 0.2, 0.5	1, 2, 3, 5, 10

Table 3.6: Optimization setup parameters: Min Hotspot Dose, LB Hotspot Volume, and Factor Edge Length for contV optimizations.

3.4. Plan Selection Scheme

To make the clinical analysis feasible, given the study time frame, we select plans based on predetermined criteria. The plan selection scheme is divided into data processing, normalization, and visualization steps, executed through a multi-threaded Python script.

The analysis is based on a randomly chosen run of BRIGHT given the selected optimization technique and setup parameters (in our case, 2-objective DHI and 3-objective contV, each using 250,000 DCPs and 600s). The clinical and BRIGHT output files are processed for a run, and the filenames follow specific templates. The primary metrics used for analysis include the LCI and DHI, which reflect the quality of each plan. The selection can be performed on any of the available objectives (LCI, LSI, contV, DHI) or metrics such as V_{150}^{avoid}. These become the grounds of our selection method. The data is then normalized to fall within the $[0, 1]$ range, enabling direct comparison between the objectives and the clinical and resulting plans. Once the data is loaded and normalized, each treatment plan is compared against the clinical plan. Plans with better performance than the clinical plan are identified, and the most balanced plan, representing the most balanced option within the better-than-clinical group, is selected. Additionally, the plans at the extremes of the better-than-clinical group are extracted from the results.

Firstly, we highlight three different plans based on their objective function values and balance between objectives. We focus the preselection on the area of plans better than the clinical plan in terms of our optimization objectives. Thus, we filter on plans that dominate the clinical plan and extract the following three plans:

- Highest LCI;
- Highest DHI or lowest contV;
- Most balanced dominating point (MBDP) in terms of all optimization objectives.

The third criterion is shown in Equation 3.3. It corresponds to the most balanced dominating point and aims to maximize the improvement in terms of all optimization objectives to provide a reasonable trade-off between them. Figure 3.2, with its highlighted plans, shows an example plot of a plan chosen for clinical evaluation.

$$\begin{aligned} p &= \text{argmax}_{p \in \{\text{dominates clinical}\}}(\min\left(obj_1(p), obj_2(p), obj_3(p)\right)) \\ obj_1(p) &= LCI(p), \\ obj_2(p) &= DHI(p) \text{ or } -contV(p) \text{ by choice of metric if minimizing } contV, \text{ otherwise } DHI(p), \\ obj_3(p) &= LSI(p) \text{ if 3-objective optimization, otherwise } None \end{aligned} \tag{3.3}$$

3.5. Clinical Acceptance

We perform an offline evaluation of two preselected BRIGHT plans for each patient to understand how experts view the results obtained with BRIGHT and answer RQ3. We aim to know whether the BRIGHT plans are viable solutions for clinical use and potential further feedback from a clinical standpoint.

Research questions

RQ3. How do clinical experts judge the breast brachytherapy treatment plans optimized with BRIGHT?

To perform offline evaluations that uncover the expert opinions on the resulting treatment plans, we compose a brief survey illustrated in Figure 3.3, based on the BRIGHT validation study for cervix

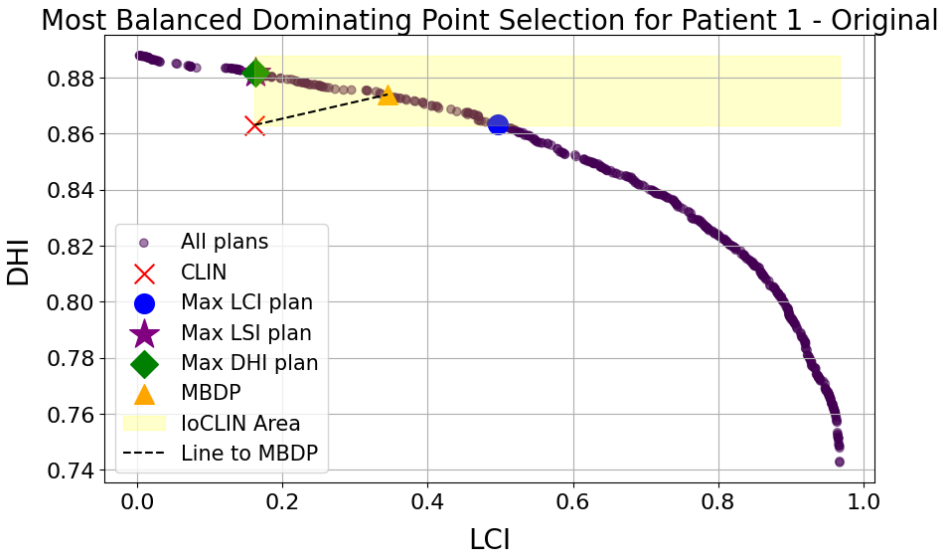

Figure 3.2: Selection method example showing the clinical plan (CLIN), highest LCI, and most balanced dominating point (MBDP), highlighted in yellow is the improvement over CLIN area (IoCLIN).

brachytherapy at Leiden University Medical Center (LUMC) for plan comparison.

We perform the first preselection phase to find the first set of plans (Plan 1) for all patients using the 2O DHI technique (maximizing DHI and LCI). We consider this a suitable aim as it involves DHI maximization, a measure of significant importance to clinical practice.

Following the first preselection phase, we discover that certain patients compromise too much on homogeneity due to the hard constraint on the ring. This suggests we could further reduce ring sparing in favor of homogeneity. Thus, adding the LSI objective can improve trade-offs between the clinical aims. Having found an opportunity to optimize a third objective, we introduce the 3O DHI and 3O contV techniques.

For the second preselection phase, we focus on the 3O contV technique. This approach aims to reduce the ring sparing compared to the 2-objective approaches while minimizing the hotspot volume within the avoidance ROI. In this case, we hypothesize that by using connected component analysis, we can explore spatial structures and, ultimately, reduce hotspots.

Lastly, once we select the most balanced plan for the two chosen techniques, we take screenshots of the BRIGHT interface. Each shared plan contains at least 10 and up to 50 screenshots of slices with isodose lines, based on the number of slices per patient and the isodose line changes across slices (i.e., very similar slices are omitted). Moreover, the screenshots contain the DVI tables within BRIGHT. Across patients, we summarize the main DVIs in Table 4.1. Finally, our expert fills out each plan's questionnaires Figure 3.3.

BRIGHT validation study for BREAST BT at CWI: Plan Evaluation

Plan Quality

The plan is satisfactory in terms of the DVH parameter values.

☐ Strongly disagree ☐ Disagree ☐ Neutral ☐ Agree ☐ Strongly agree

Why?
Answer:

The plan is satisfactory in terms of homogeneity.

☐ Strongly disagree ☐ Disagree ☐ Neutral ☐ Agree ☐ Strongly agree

Why?
Answer:

Is the plan clinically acceptable?

☐ yes ☐ no

If not, why?
Answer:

If you like, please share any unaddressed feedback:
Answer:

Figure 3.3: BRIGHT validation study questionnaire for plan evaluation. Based on the BRIGHT validation study for cervix brachytherapy at LUMC

4

Results

This chapter showcases the results of our experimental analysis of two MO optimization setups for nine different patients. Sections 4.2 and 4.1 each describe two optimization setups that use 2 and 3 objectives, respectively. The former is primarily concerned with optimizing DHI, while the latter substitutes it for contV. Section 4.3 describes the criteria we employed to pick individual plans from the Pareto approximation front generated by the EA. Finally, Section 4.4 summarizes our clinical expert's feedback for the selected plans. Based on the analysis and the expert's feedback, we make general recommendations in Chapter 5.

4.1. Experimental Results for DHI

In this section, we present the results of the optimization experiments focused on maximizing the DHI and maximizing the LCI. In this scenario, we employ DHI with the expectation that generated solutions would be favorable for clinical practice, given the clinical importance of DHI in relation to late toxicities. We include the LCI objective as it is key for ensuring that plans effectively cover the ROIs. The optimization time budget and the number of dose calculation points varied, and results were aggregated over ten runs for each patient. In this section, we primarily aim to address RQ2 and its subquestions, stated in Section 3.3, and analyze the evidence relevant to answering RQ1 of Section 3.2.

4.1.1. Results for 2-Objective Optimization: DHI + LCI

Effects of optimization time: RQ2.1

We begin by analyzing the effects of varying the time budget. Figure 4.1 displays the results. We highlight that all experiments generate solutions that dominate the clinical treatment plan for computational budgets above 180 seconds for patients 1, 2, and patients 4 to 7. For patients 1, 2, 4, 6, and 7, all computational budget experiments with a budget of 60 seconds or higher result in strict improvements over the reference solution, whereas higher budgets are required for the other patients. For patients 3 and 9, the experiments could not generate solutions that dominated the clinical plan. In addition to this observation, we highlight the following trends:

- **Trade-offs between DHI and LCI**: As seen in Figures 4.1a through 4.1i, there is a trade-off between DHI and LCI. Optimizing for higher LCI tends to reduce DHI, especially in the regions where LCI is maximized, and vice-versa.
- **Longer time budgets lead to better results**: As we increase the time budget from 30 seconds to 900 seconds, the quality of solutions improves significantly. The Pareto approximation fronts expand and shift toward higher DHI values while maintaining a high LCI.
- **Impact on patients at extremes and suboptimal implants**: Patients at extreme (with smallest/largest CTVs), such as Patient 5, or patients with suboptimal implants, such as patients 8 and 9, require more computational effort to optimize effectively (as the experiments with contV and LCI also observe). The results for patients at extremes and/or with suboptimal implants show broader distributions compared to patients like Patient 2 or Patient 4. Furthermore, we note the different shapes of the Pareto approximation front of patients 5, 8, and 9 for computational budgets of

180 seconds or lower. For these time budgets, their Pareto approximation fronts are significantly scarcer the higher the LCI. This hints at the difficulty of optimizing coverage for problems with volumes at extremes or suboptimal implants.

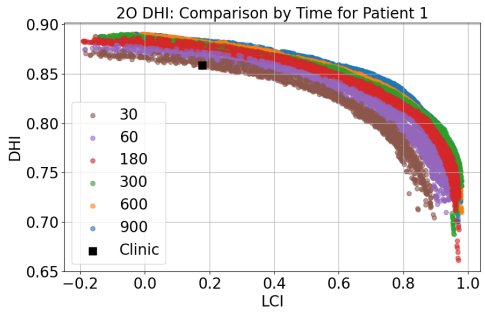

(a) Patient 1: Comparison by time (DHI + LCI)

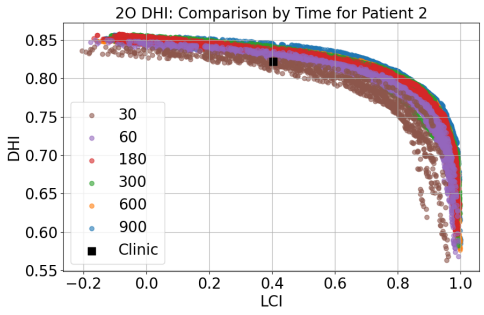

(b) Patient 2: Comparison by time (DHI + LCI)

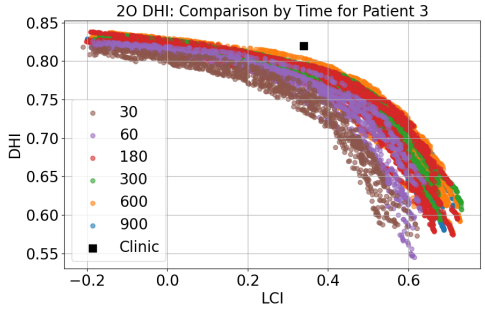

(c) Patient 3: Comparison by time (DHI + LCI)

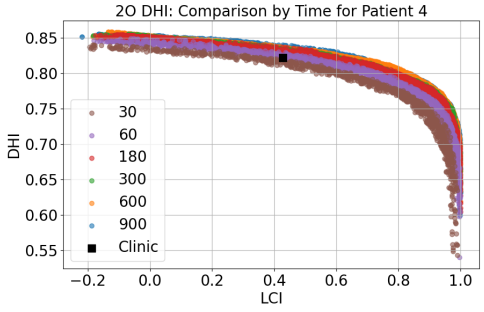

(d) Patient 4: Comparison by time (DHI + LCI)

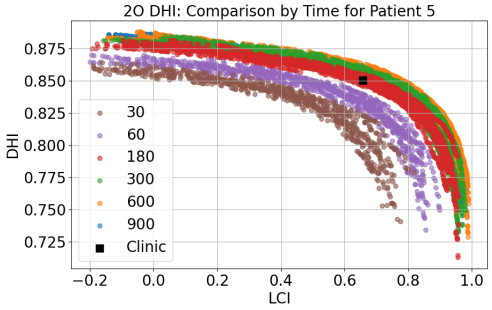

(e) Patient 5: Comparison by time (DHI + LCI)

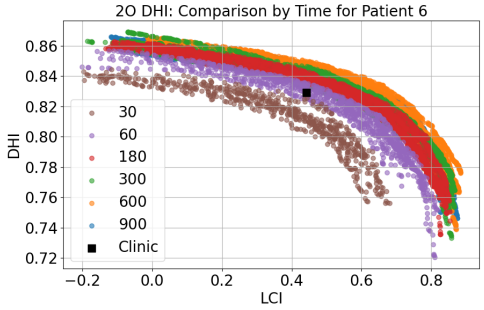

(f) Patient 6: Comparison by time (DHI + LCI)

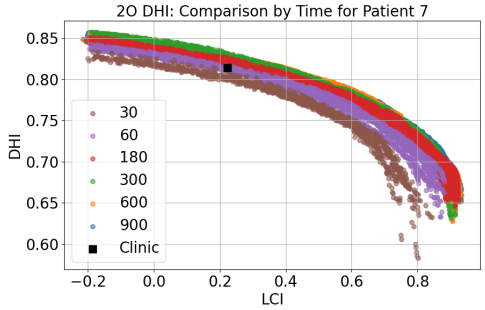

(g) Patient 7: Comparison by time (DHI + LCI)

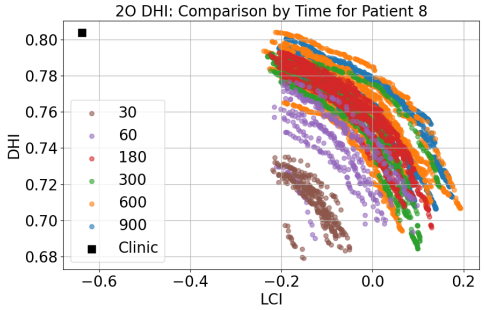

(h) Patient 8: Comparison by time (DHI + LCI)

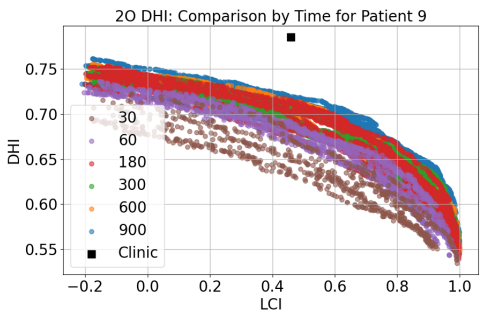

(i) Patient 9: Comparison by time (DHI + LCI)

Figure 4.1: Comparison of time budgets for different patients based on the 2-objective optimization (DHI + LCI). Each experiment was run ten times. Each plot includes the clinical plan as a reference point.

Effects of dose calculation points: RQ2.2
We analyze the effects of varying the number of dose calculation points on the generated solutions. We use an experimental setup with between 10,000 and 300,000 DCPs and 600-second runs. Figure 4.2 displays the results. We note the following features:

- **Clinical plan comparison**: For all except patients 3, 8, and 9, runs of BRIGHT with 50,000 or more DCP consistently generate Pareto approximation fronts that supersede the clinical reference point. The results demonstrate that this configuration can produce Pareto-dominating solutions with respect to plans produced in the clinic.
- **Objective trade-offs**: As in the contV + LCI configuration, DHI + LCI also leads to results with significantly higher LCI values. While the upper bound for LCI in this experiment is above 0.9 for all patients but Patient 9, increasing the number of DCPs tends to extend the breadth of the front across the LCI axis. This is especially visible in the transition between 10,000 and 50,000 dose calculation points and is less pronounced as the number of DCPs increases.
- **Diminishing returns**: The results suggest that increasing the number of DCPs beyond 200,000 yields little to no improvements for a fixed time budget of 600 seconds, and the front even decays in some instances with 300,000 dose calculation points. This is especially visible for the suboptimal implant patients 8 and 9, as well as the extreme volume case of Patient 3. While this result is not generalized, it highlights the importance of performing preliminary studies that aim to find a balanced number of DCPs for the computational budget available to clinicians.

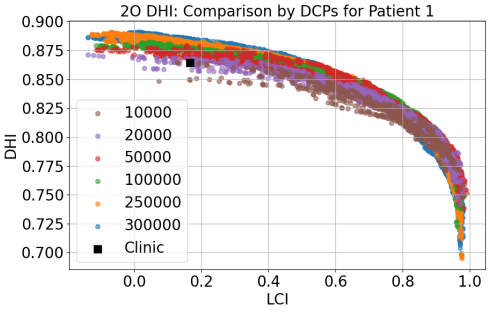

(a) Patient 1: Comparison by number of DCPs (DHI + LCI)

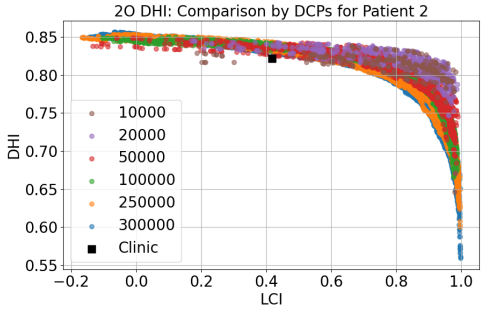

(b) Patient 2: Comparison by number of DCPs (DHI + LCI)

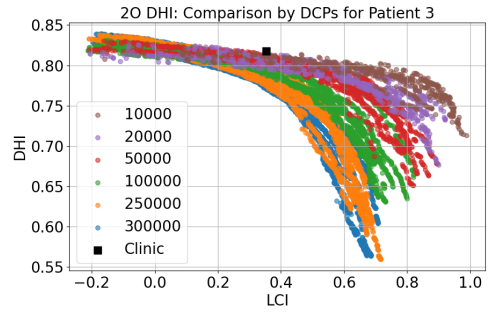

(c) Patient 3: Comparison by number of DCPs (DHI + LCI)

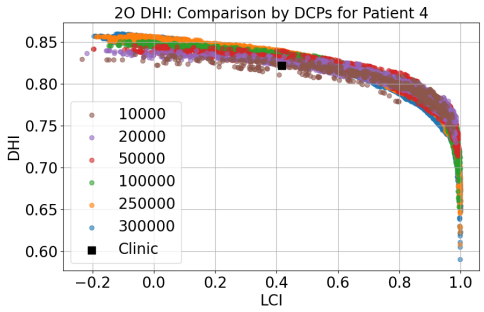

(d) Patient 4: Comparison by number of DCPs (DHI + LCI)

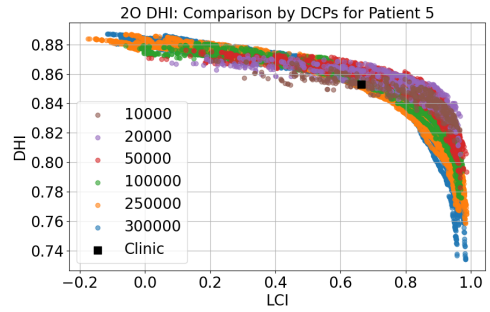

(e) Patient 5: Comparison by number of DCPs (DHI + LCI)

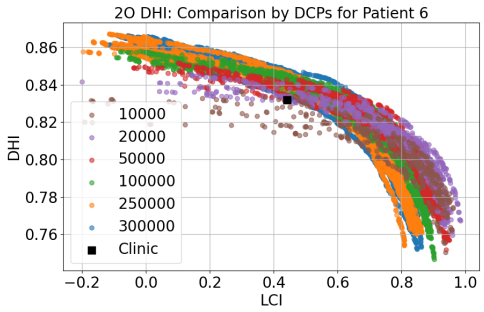

(f) Patient 6: Comparison by number of DCPs (DHI + LCI)

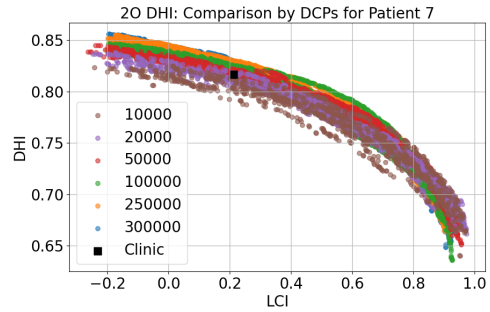

(g) Patient 7: Comparison by number of DCPs (DHI + LCI)

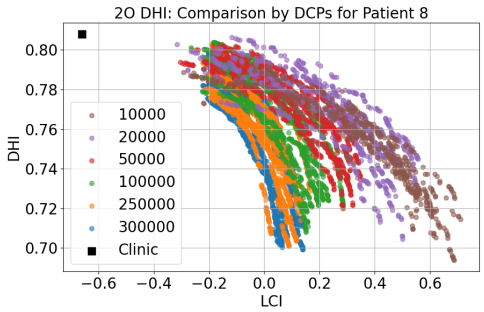

(h) Patient 8: Comparison by number of DCPs (DHI + LCI)

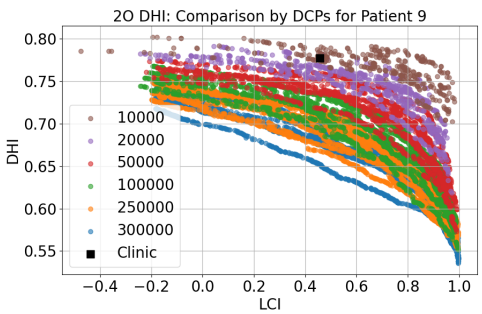

(i) Patient 9: Comparison by number of DCPs (DHI + LCI)

Figure 4.2: Comparison of the number of DCPs for different patients based on the 2-objective optimization (DHI + LCI). Each setting was run ten times as in previous experiments. Each plot includes the clinical plan as a reference point.

4.1.2. Results for 3-Objective Optimization: DHI + LCI + LSI

In this subsection, we explore the effects of adding the Least Sparing Index as a third objective, next to DHI and LCI. The rationale is that adding a third objective to the optimization problem may prove valuable for clinicians as this could alter the profile of the evolved plans. By embracing trade-offs in terms of LSI instead of discarding solutions through hard constraints, we can generate clinically appealing plans in potentially different areas of the optimization landscape. In this case, LSI seeks to minimize the radiation to surrounding organs, which is a conflicting (and therefore compelling) objective.

Effects of optimization time: RQ2.1

Figure 4.3 shows the effects of the time budget on the 3-objective optimization problem. For legibility, we show the 3D approximation front split into one comparison between the DHI and LCI components and one between LSI and LCI and assess the effect of the optimization problem on the three objectives. We highlight the following features:

- **Higher time budgets do not always lead to improved solutions:** We specifically highlight the along sections of the Pareto approximation front, in which runs with a budget of 180 is sufficient to supersede the clinical point across all patients. This behavior is more accentuated than in the 2-objective techniques. This indicates that lower time budgets (around 180 seconds) may already suffice and reach a point of saturation, where increasing the computation time leads to only marginal improvements.
- **Profile similarity:** We previously highlighted more complicated cases, such as Patient 5 and Patient 8, which display similar behavior to the 2-objective case, though notably, the trend is lesser. Such patients require a more extended computational budget to dominate the clinical plans.

Together, these two observations lead to a compelling trade-off. The improvement in solution quality is lesser for the 3-objective formulation than its 2-objective counterpart. Despite its higher computational overhead, the 3-objective algorithm provides more and higher quality solutions for low computational budgets. In the case of DHI, adding the LSI objective shows considerable improvements (for less overhead than contV, shown in Section 4.2.2). This can be explained by the presence of the hotspot detection mechanism of contV and the lack thereof in the case of DHI.

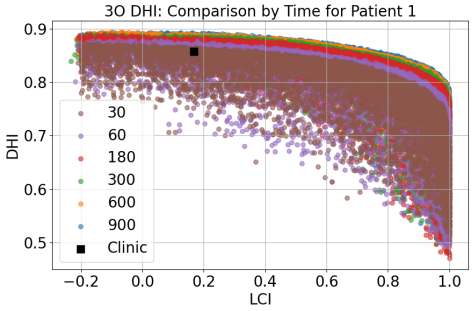

(a) Patient 1: Comparison by time (DHI + LCI)

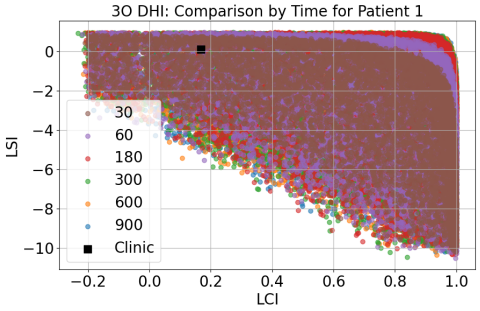

(b) Patient 1: Comparison by time (LSI + LCI)

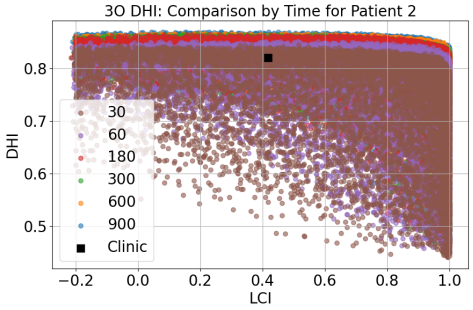

(c) Patient 2: Comparison by time (DHI + LCI)

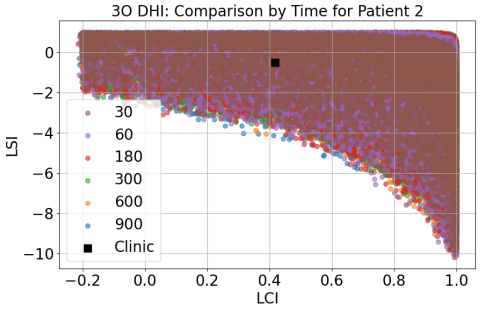

(d) Patient 2: Comparison by time (LSI + LCI)

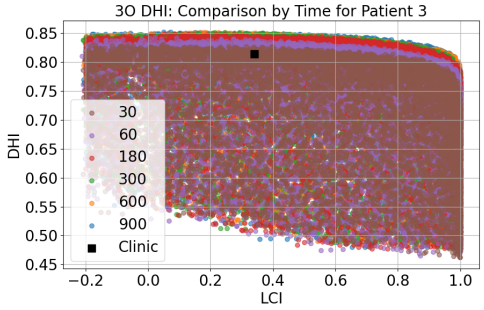

(e) Patient 3: Comparison by time (DHI + LCI)

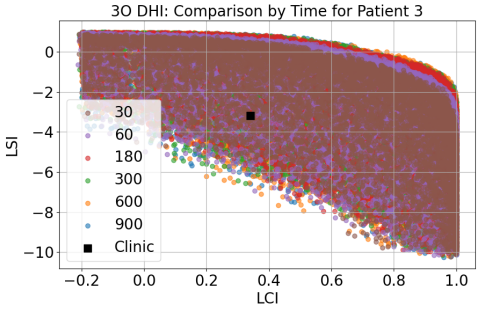

(f) Patient 3: Comparison by time (LSI + LCI)

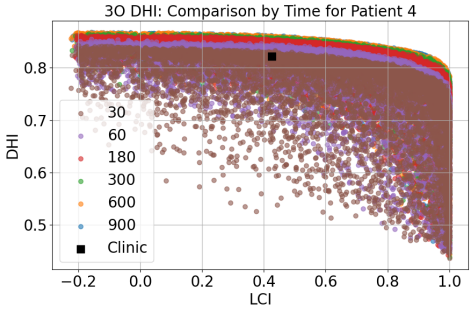

(g) Patient 4: Comparison by time (DHI + LCI)

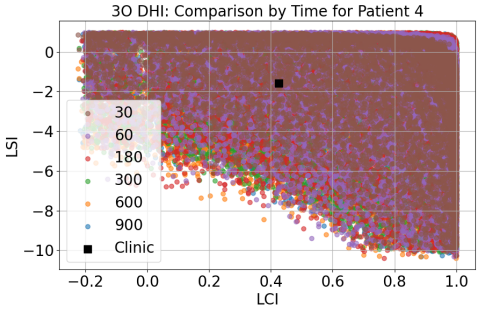

(h) Patient 4: Comparison by time (LSI + LCI)

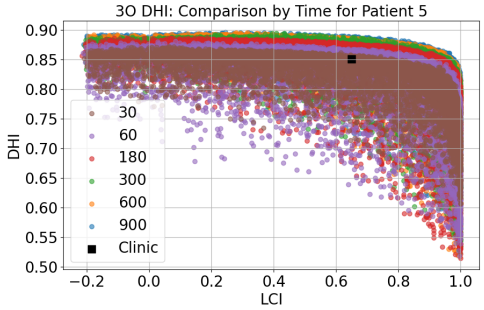

(i) Patient 5: Comparison by time (DHI + LCI)

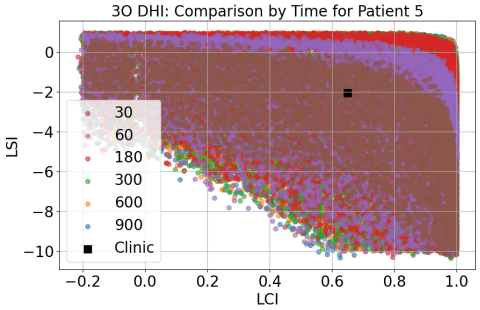

(j) Patient 5: Comparison by time (LSI + LCI)

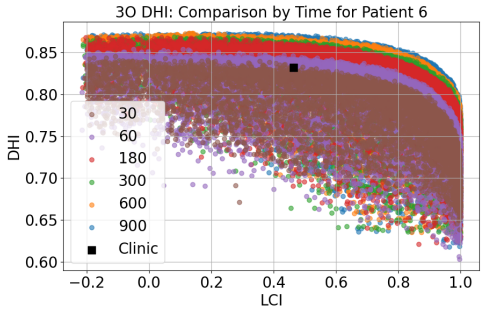

(k) Patient 6: Comparison by time (DHI + LCI)

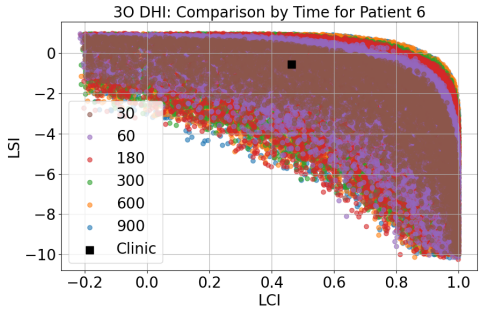

(l) Patient 6: Comparison by time (LSI + LCI)

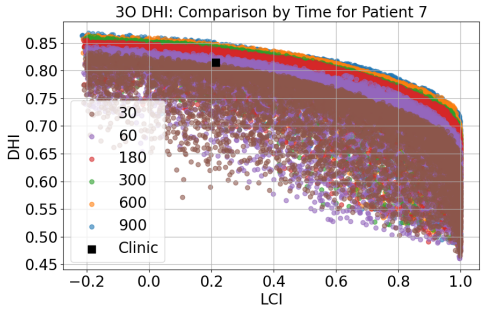

(m) Patient 7: Comparison by time (DHI + LCI)

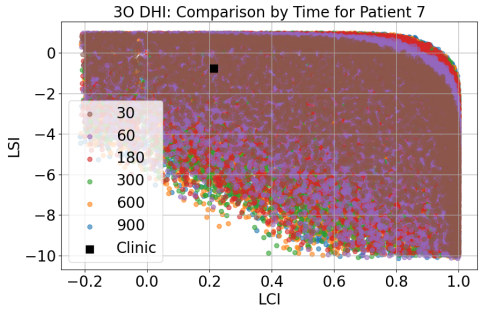

(n) Patient 7: Comparison by time (LSI + LCI)

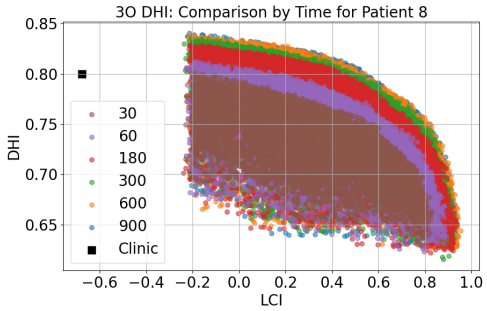

(o) Patient 8: Comparison by time (DHI + LCI)

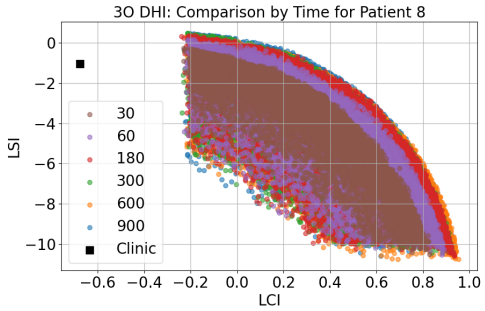

(p) Patient 8: Comparison by time (LSI + LCI)

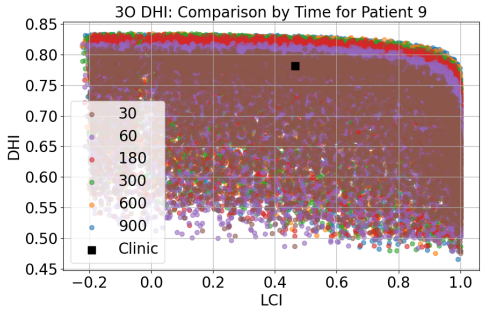

(q) Patient 9: Comparison by time (DHI + LCI)

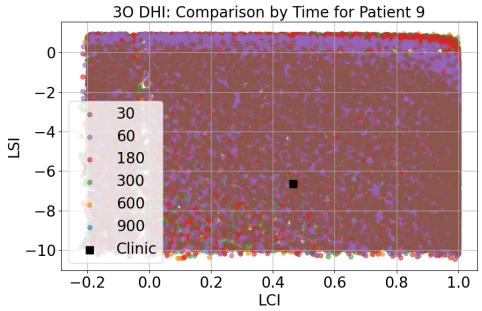

(r) Patient 9: Comparison by time (LSI + LCI)

Figure 4.3: Comparison of time budgets for different patients based on the 3-objective optimization (DHI + LCI + LSI). Each plot includes the clinical plan as a reference point.

Effects of dose calculation points: RQ2.2
As in previous experiments, we vary the number of DCPs from 10,000 to 300,000 and display the results in Figure 4.4. For clarity, we present the LCI objective against DHI and then LSI in parallel to observe all three objectives without a 3-dimensional visualization overhead. We highlight the following observations:

- **Clinical plan comparison**: Across all nine patients, the DHI + LCI + LSI configuration generates solutions that exceed the objective values of the clinical plan in terms of DHI, LCI, and LSI. We note that introducing LSI, in addition to the fronts on the (narrower) 2-objective front, the 3-objective formulation generates more solutions with similar DHI and LCI features to the reference.
- **Broader objective range**: In addition to producing more solutions in the vicinity of the reference clinical plan, the DHI + LCI + LSI technique also includes more plans towards the ”edges” of the two objectives for both LCI, DHI and LCI, LSI pairs. The Pareto approximation fronts include more solutions overall and more solutions towards the extreme ranges of the front (LCI ≥ 0.9, DHI ≥ 0.8, LSI ≥ 0).

We further note the correlation between the shape of the fronts and the patients' features. Figures 4.2c and 4.4e show a similar trend – the addition of LSI effectively makes the 3D Pareto approximation front span almost the entirety of both DHI + LCI, and LSI + LCI spaces. While this may prove overwhelming for clinicians in practice, it could provide additional insight into which underlying features of patients like Patient 3 underpin this trend, which then, in turn, can provide further explanations as to which (2- or 3-objective) implementation of BRIGHT might be preferable in practice.

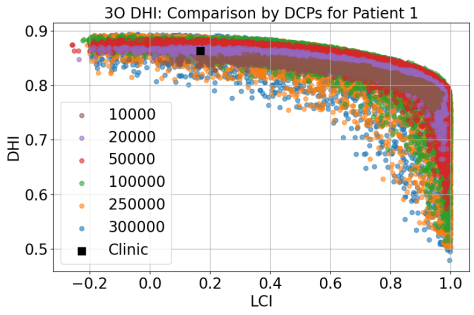

(a) Patient 1: Comparison by DCPs (DHI + LCI)

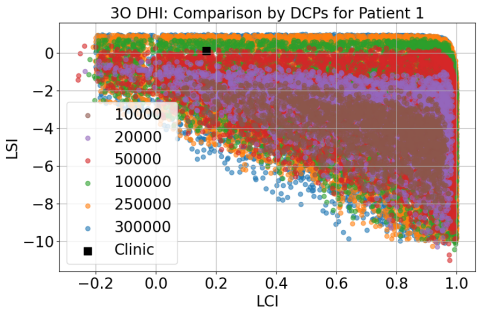

(b) Patient 1: Comparison by DCPs (LSI + LCI)

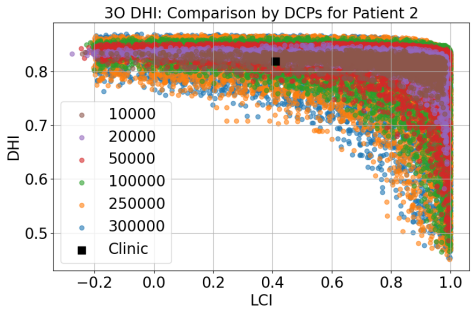

(c) Patient 2: Comparison by DCPs (DHI + LCI)

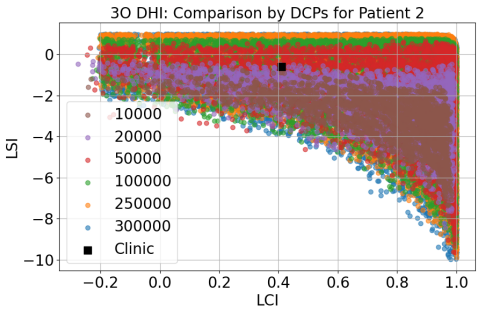

(d) Patient 2: Comparison by DCPs (LSI + LCI)

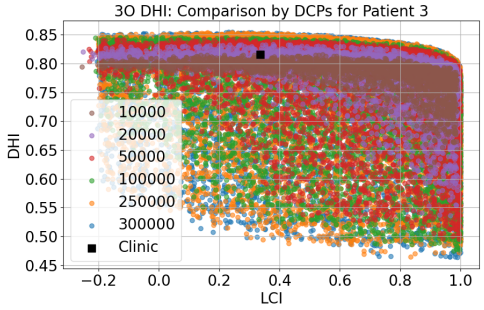

(e) Patient 3: Comparison by DCPs (DHI + LCI)

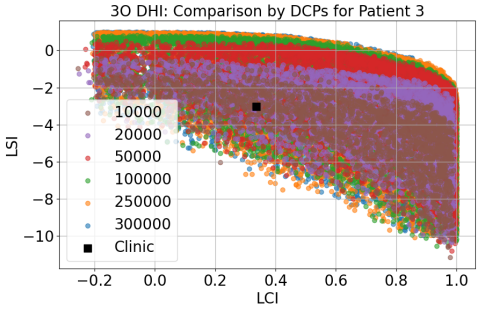

(f) Patient 3: Comparison by DCPs (LSI + LCI)

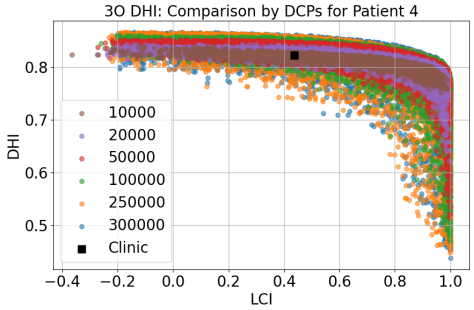

(g) Patient 4: Comparison by DCPs (DHI + LCI)

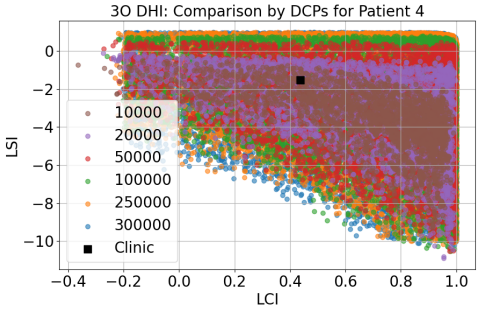

(h) Patient 4: Comparison by DCPs (LSI + LCI)

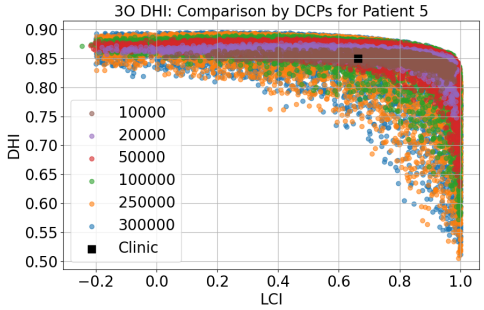

(i) Patient 5: Comparison by DCPs (DHI + LCI)

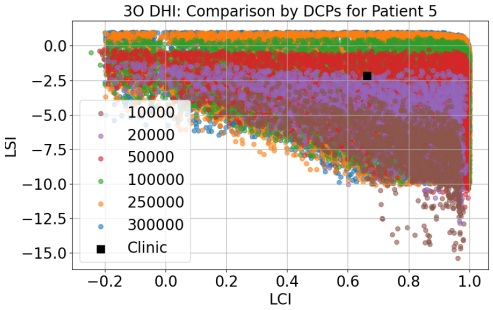

(j) Patient 5: Comparison by DCPs (LSI + LCI)

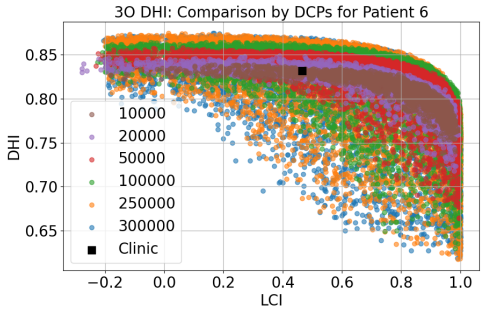

(k) Patient 6: Comparison by DCPs (DHI + LCI)

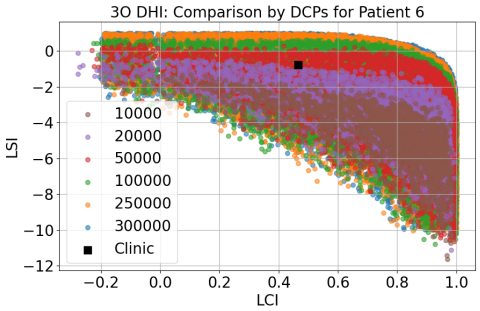

(l) Patient 6: Comparison by DCPs (LSI + LCI)

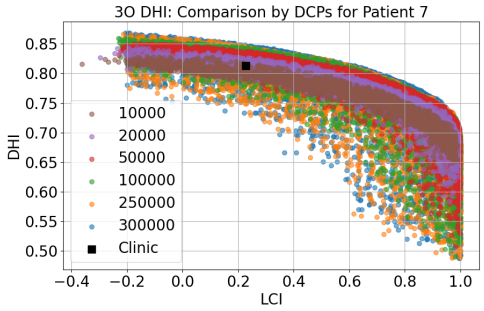

(m) Patient 7: Comparison by DCPs (DHI + LCI)

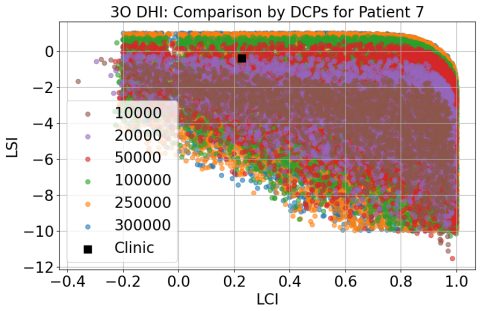

(n) Patient 7: Comparison by DCPs (LSI + LCI)

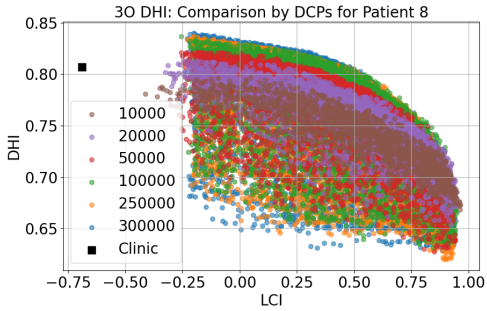

(o) Patient 8: Comparison by DCPs (DHI + LCI)

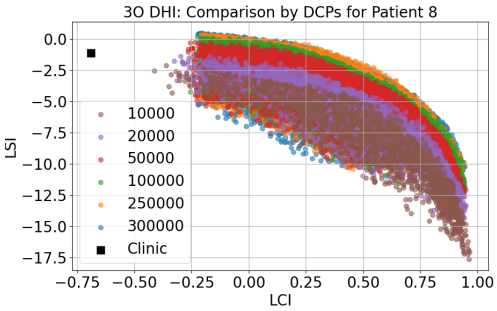

(p) Patient 8: Comparison by DCPs (LSI + LCI)

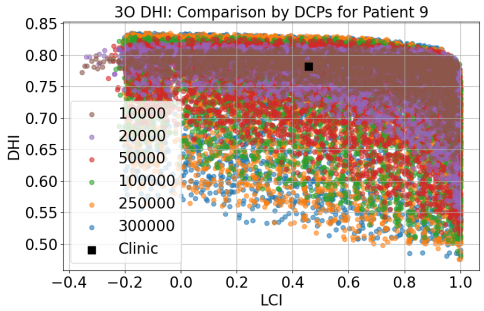

(q) Patient 9: Comparison by DCPs (DHI + LCI)

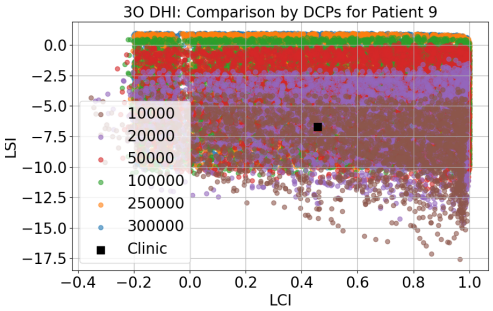

(r) Patient 9: Comparison by DCPs (LSI + LCI)

Figure 4.4: Comparison of DCP budgets for different patients based on the 3-objective optimization (DHI + LCI + LSI). Each setting was run ten times. Each plot includes the clinical plan as a reference point.

4.2. Experimental Results for contV (HV)

This section presents the results of running BRIGHT with 2 and 3 multiobjective formulations. Section 4.2.1 provides an overview of the 2-objective formulation, while 4.2.2 analyzes the 3-objective counterpart. For both instances, we consider the effects of 2 variables – the number of dose calculation points optimization time. For each instance, we report the DHI and LCI of the solutions BRIGHT retains, as these metrics are more widely considered in clinical practice. The results of this section are primarily concerned with RQ2 and its subquestions, defined in Section 3.3. Together with Section 4.1, we aim to formulate answers for the remainder of the sub-questions of RQ1.

Three additional experiments that consider three supplementary variables in factor edge length, bound hotspot volume, and minimum hotspot dose are available in Appendix B. We found the most advantageous configuration to correspond to the combination of expert knowledge and contV research parameters previously employed in [35]. Namely, `factor_edge_length` = 2, `lower_bound_hotspot_volume` = 0.1 cm^3, and `min_hotspot_dose` = 150% of PD.

4.2.1. 2-Objective Optimization: HV + LCI

This section presents the results of the optimization experiments that focus on minimizing the hotspot volume and maximizing the LCI. This technique is 2-objective and contains the ring-sparing hard constraint. The optimization time budget and the number of dose calculation points were varied, and the results include comparisons between different numbers of DCPs and time budgets for each patient.

Effects of optimization time: RQ2.1

We analyze the results of experiments varying the optimization time for the five patients. Each experiment was run ten times, and the results were aggregated. The number of dose calculation points for this round of experiments was fixed at 250,000, a significant factor affecting computation time. These experiments aimed to assess the effect of increasing the optimization time budget on the algorithm's performance regarding the two primary objectives: LCI and HV.

Figure 4.5a through 4.5i correspond to different patients, showing the Pareto approximation front of different solutions with varying optimization times. The optimization time budgets used in the experiments are 30, 60, 180, 300, 600, and 900 seconds. The clinical plan is included as a baseline for comparison. We note the following trends:

- **Clinical plan comparison**: Across all patients, we observe that increasing the optimization time generally improves solution quality, particularly in the higher time ranges (600-900 seconds). This is most evident as the solutions shift further right along the LCI axis. For 6 out of 9 patients, the Pareto approximation fronts dominate the clinical plan for all instances with a computational budget of over 600 seconds. For shorter time budgets (e.g., 30, 60 seconds), the results tend to cluster closer to the clinical plan or show either limited improvement or results with lesser properties than the clinical plan. This suggests the computational time was insufficient for exploring the solution space well. The distribution of these points is broader in terms of hotspot volume values, indicating more variance in contiguous volume performance with limited time for optimization. These points often fall beneath or close to the clinical plan, particularly in the left region where LCI is minimized. We note that, especially for Patient 8, the clinical plan has a different profile than the solutions generated by the BRIGHT. This could prove useful clinically, as experts might get a more diverse palette to select from.
- **Improvement with longer time budgets**: For the higher time budgets (300-900 seconds), the algorithm has more time to refine solutions, resulting in clear improvements in both LCI and HV. In all patients, the longer time runs push the solutions toward a better Pareto approximation front compared to shorter runs. Notably, the Pareto approximation front's slope steepens as LCI improves with increasing time. This indicates that the algorithm can produce solutions that dominate the clinical plan. Patient 1 and Patient 2 show significant improvements as time increases, particularly for runs of 600 and 900 seconds, which dominate the clinical plan more strongly. Patient 5 exhibits a similar trend, although the variance in contV is slightly higher for shorter runs, and the front is more dispersed. Patient 3 and Patient 4 also follow this trend, with Patient 3 showing the least variance overall, which may be attributed to the problem's smaller search space.
- **Contiguous volume**: An interesting observation is the trade-off between LCI and HV. Solutions with stronger LCI objective values may lead to a slight increase in contV, as increasing target

organ coverage can imply higher radiation doses, thus also larger hotspots. This trade-off is more prominent in patients with larger volumes, such as Patient 5, where larger contiguous regions also require more coverage.

- **Computational complexity**: It is important to note that increasing the number of dose calculation points makes the problem more computationally intensive, and in this case, the fixed number of 250,000 points amplifies this effect. More extended time budgets provide more opportunities to explore the solution space but also require substantially more computational resources. The connected component analysis used for calculating contV further complicates this. This leads to the observed clustering for shorter runs and more refined fronts for longer runs.

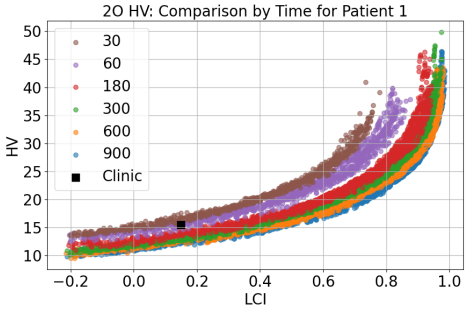

(a) Comparison by time for **Patient 1** (HV + LCI).

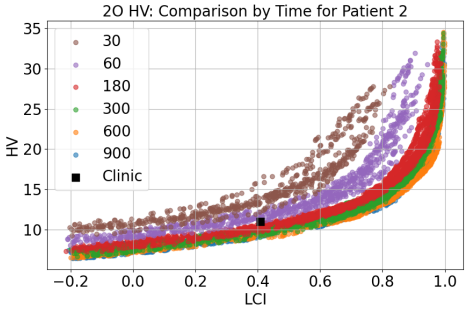

(b) Comparison by time for **Patient 2** (HV + LCI).

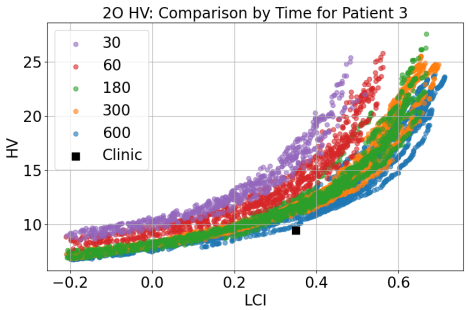

(c) Comparison by time for **Patient 3** (HV + LCI).

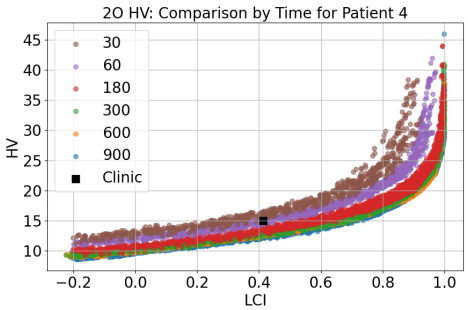

(d) Comparison by time for **Patient 4** (HV + LCI).

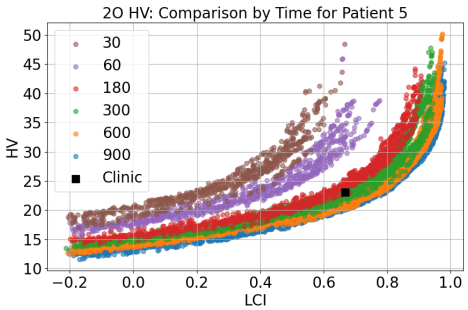

(e) Comparison by time for **Patient 5** (HV + LCI).

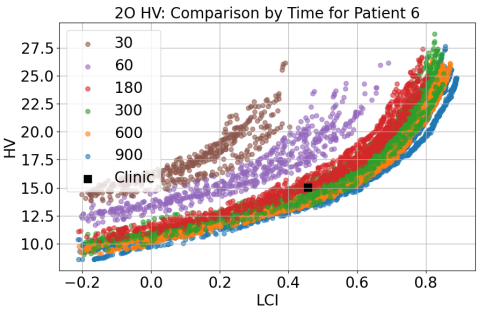

(f) Comparison by time for **Patient 6** (HV + LCI).

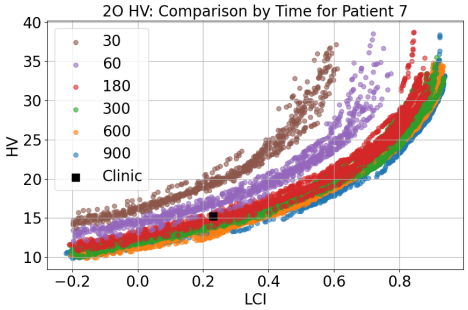

(g) Comparison by time for **Patient 7** (HV + LCI).

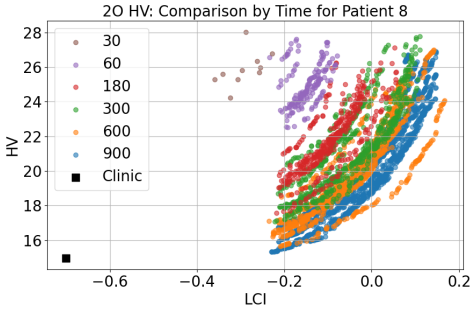

(h) Comparison by time for **Patient 8** (HV + LCI).

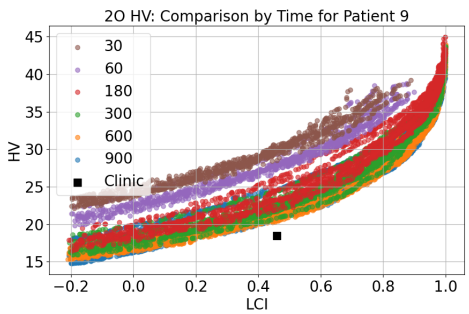

(i) Comparison by time for **Patient 9** (HV + LCI).

Figure 4.5: Comparison of time for different patients based on the 2-objective optimization (HV + LCI) with varying time budgets. Each setting has been run ten times to observe variance across the front. Each plot includes the clinical plan as a reference point.

Effects of dose calculation points: RQ2.2
Figure 4.6 shows how varying the number of dose calculation points affects the trade-off between the HV and the LCI. The results for each dose calculation point number have been aggregated over ten runs. Each run had an optimization time budget of 600 seconds. We highlight the following trends:

- **Clinical plan comparison**: The clinical plan serves as a reference point in each plot, helping to visually compare how the DCP-based optimizations perform relative to existing clinical standards. The results show that all tested instances generate Pareto approximation fronts that would either contain or dominate the reference clinical plan for all patients. In practice, this means that all optimization settings lead to improvements concerning the optimization criteria over the clinical plans. In some cases, higher DCPs significantly improve upon the clinical plan, while in others, the improvements are more modest.
- **Higher DCPs generally lead to better results**: As the number of DCPs increases, the solutions tend to improve. More sampling points allow for better granularity and precision in the dose calculations. However, for more difficult optimization cases, especially for patients 3, 8, and 9 (for small CTV and suboptimal implants), we note a decrease in HV with an increasing number of DCPs. We hypothesize that in small CTV cases, a higher number of dose calculation points can bias the optimization by overestimating the HV metric by discovering additional points with a dose higher than the hotspot lower bound. A robust analysis of such out-of-distribution cases likely requires a broader data set and more clinical resources and, as such, is left to future work.
- **Computational complexity increases**: However, the computational requirements also grow significantly as the number of DCPs increases. Each additional DCP introduces more complexity, especially in the case of contV optimization, which involves connected component analysis. This analysis is computationally expensive and more challenging with higher resolution (more DCPs).

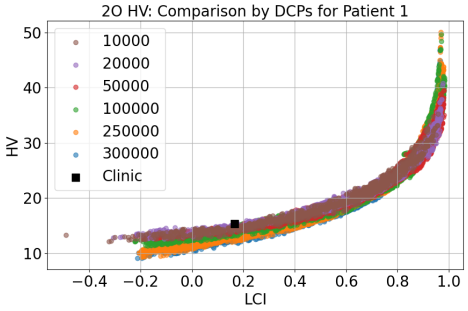

(a) Patient 1: Comparison by DCPs for **Patient 1** (HV + LCI).

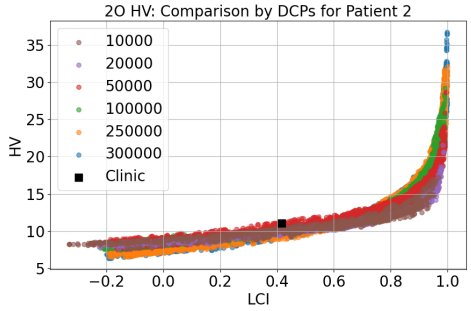

(b) Patient 2: Comparison by DCPs for **Patient 2** (HV + LCI).

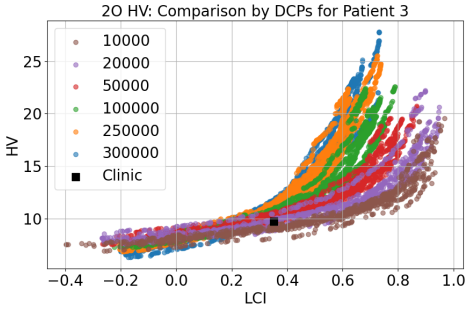

(c) Patient 3: Comparison by DCPs for **Patient 3** (HV + LCI).

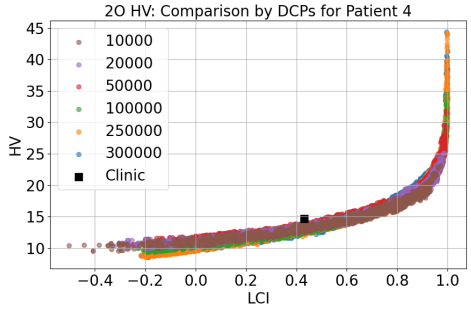

(d) Patient 4: Comparison by DCPs for **Patient 4** (HV + LCI).

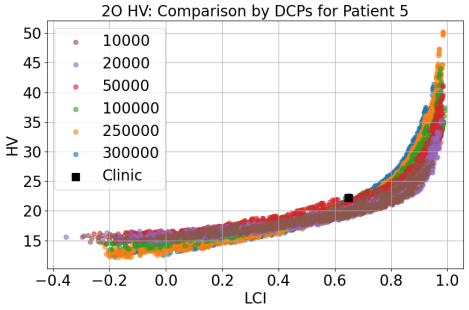

(e) Patient 5: Comparison by DCPs for **Patient 5** (HV + LCI).

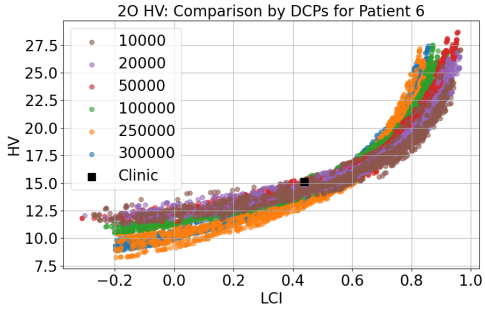

(f) Patient 6: Comparison by DCPs for **Patient 6** (HV + LCI).

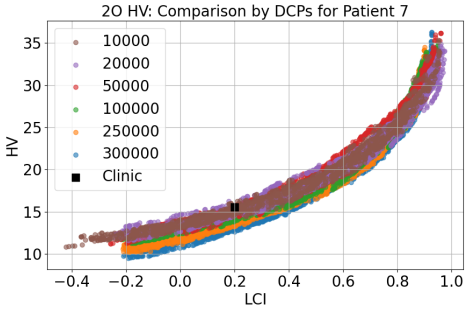

(g) Patient 7: Comparison by DCPs for **Patient 7** (HV + LCI).

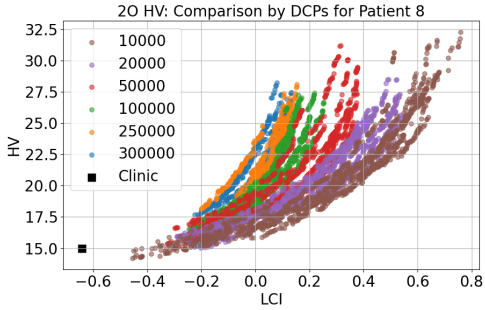

(h) Patient 8: Comparison by DCPs for **Patient 8** (HV + LCI).

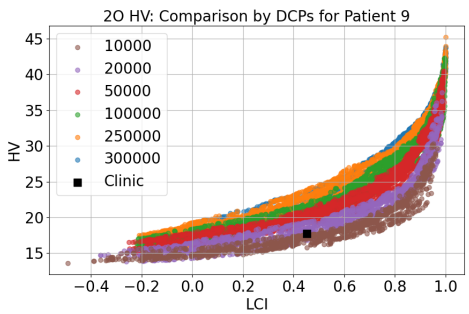

(i) Patient 9: Comparison by DCPs for **Patient 9** (HV + LCI).

Figure 4.6: Comparison of DCPs for different patients based on the 2-objective optimization (HV + LCI) with varying DCP numbers. Each setting has been run ten times to observe variance across the front. Each plot includes the clinical plan as a reference point.

4.2.2. 3-Objective Optimization: HV + LCI + LSI

In this section, we analyze the results of BRIGHT on a 3-objective formulation of the problem. In addition to HV and LCI, we supplement the optimization problem with an objective function that seeks to minimize the Least Sparing Index. We hypothesize that the inclusion of LSI can help protect OARs from excessive radiation and, in doing so, alter the structure of generated solutions and explore different areas of the search space. Again, we consider the effects of increasing the computational time and number of DCPs on the quality and distribution of obtained solutions.

Effects of optimization time: RQ2.1

In this section, we analyze the results of 3-objective optimization focusing on contV, LCI, and **LSI** with time budgets ranging from 30 seconds to 900 seconds. The experiments were run on the same five patients, with 250,000 dose calculation points, and the results were aggregated over ten runs. Each figure includes the clinical plan as a baseline for comparison.

- **Improvement with time budget**: Increasing the optimization time improves solution quality across all patients, particularly in the higher time ranges (600-900 seconds). Solutions tend to shift rightward along the LCI axis, showing improved coverage.
- **Time-performance relation**: Shorter time budgets (30-60 seconds) cluster closer to the clinical plan, offering modest improvements. Longer budgets (180-900 seconds) lead to Pareto-dominant solutions with better LCI and more balanced contV.

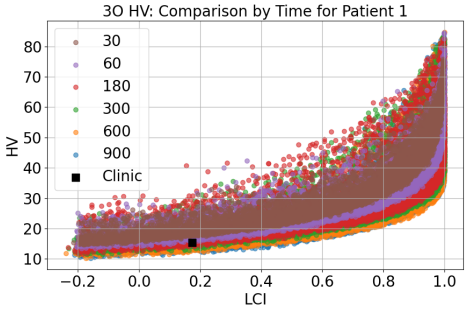

(a) Comparison by time for **Patient 1** (HV + LCI).

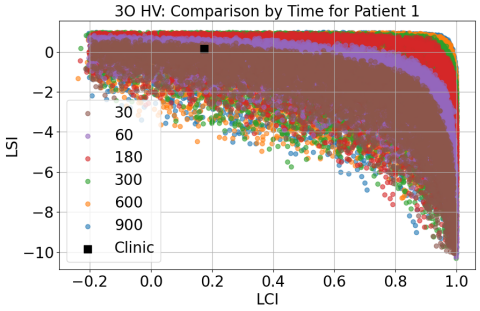

(b) Comparison by time for **Patient 1** (LSI + LCI).

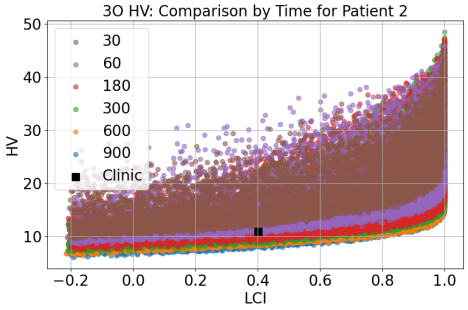

(c) Comparison by time for **Patient 2** (HV + LCI).

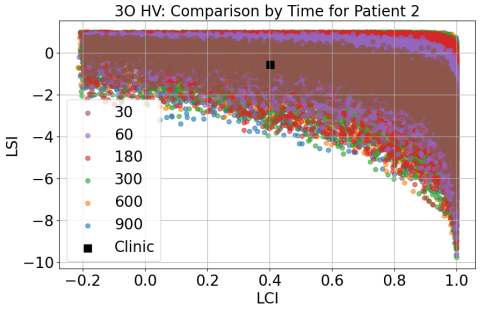

(d) Comparison by time for **Patient 2** (LSI + LCI).

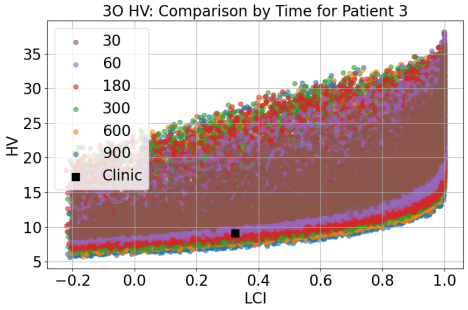

(e) Comparison by time for **Patient 3** (HV + LCI).

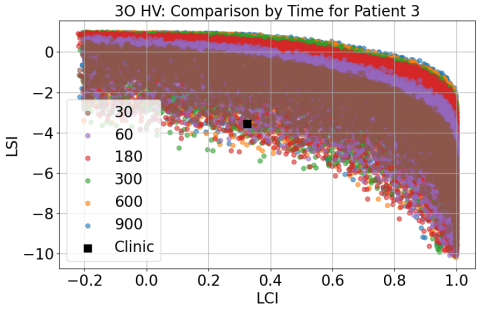

(f) Comparison by time for **Patient 3** (LSI + LCI).

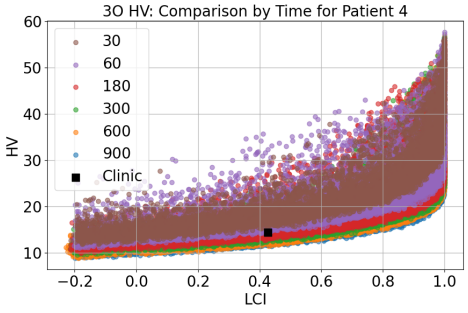

(g) Comparison by time for **Patient 4** (HV + LCI).

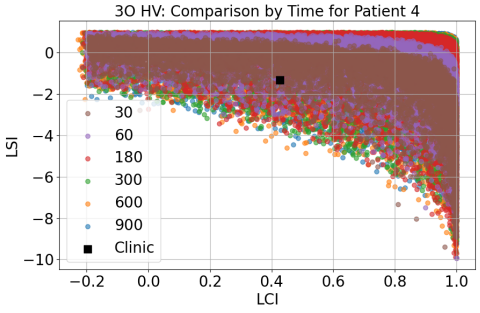

(h) Comparison by time for **Patient 4** (LSI + LCI).

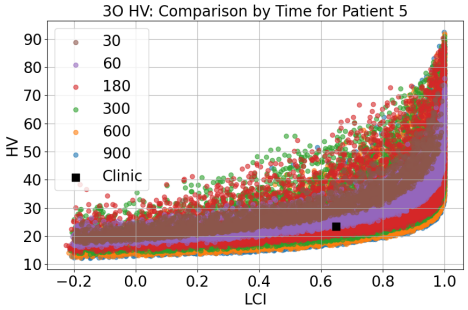

(i) Comparison by time for **Patient 5** (HV + LCI).

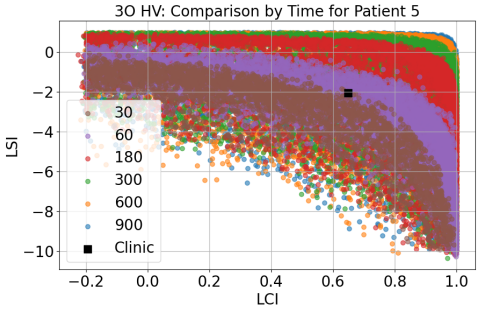

(j) Comparison by time for **Patient 5** (LSI + LCI).

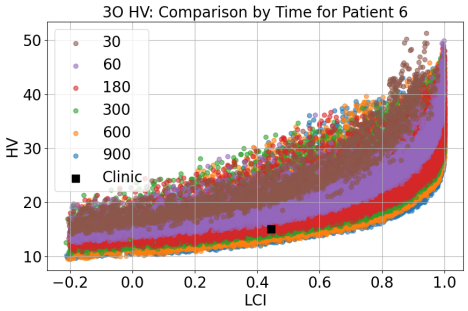

(k) Comparison by time for **Patient 6** (HV + LCI).

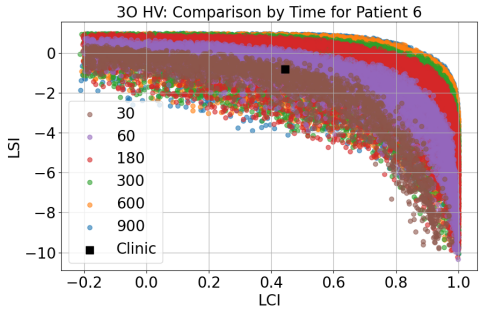

(l) Comparison by time for **Patient 6** (LSI + LCI).

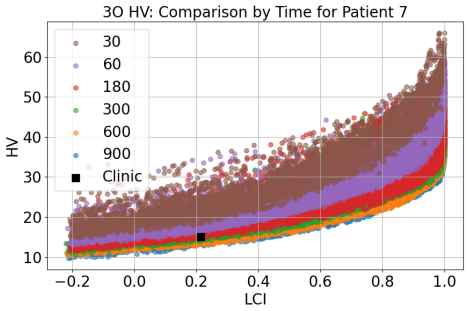

(m) Comparison by time for **Patient 7** (HV + LCI).

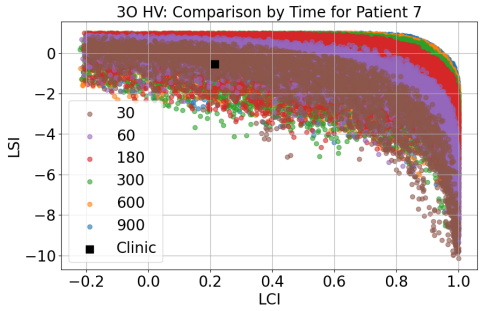

(n) Comparison by time for **Patient 7** (LSI + LCI).

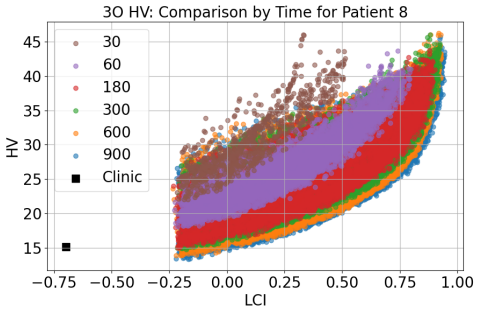

(o) Comparison by time for **Patient 8** (HV + LCI).

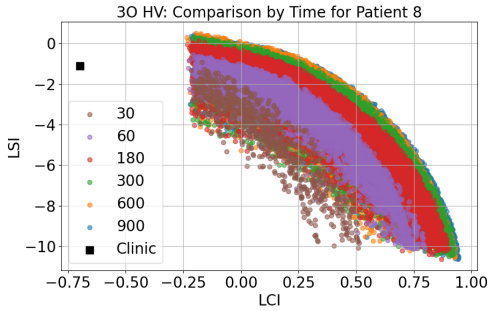

(p) Comparison by time for **Patient 8** (LSI + LCI).

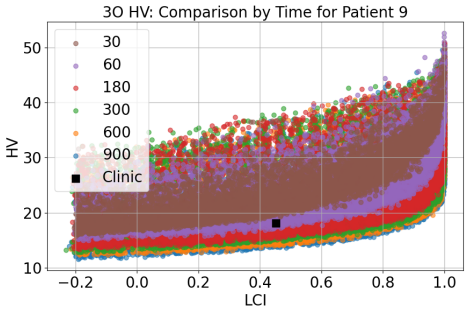

(q) Comparison by time for **Patient 9** (HV + LCI).

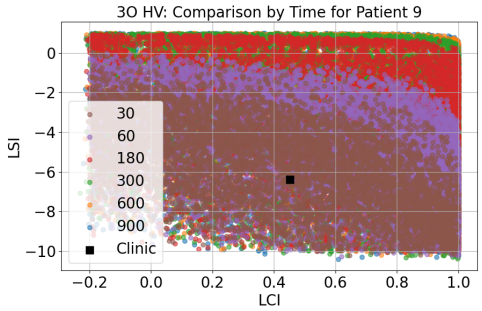

(r) Comparison by time for **Patient 9** (LSI + LCI).

Figure 4.7: Approximation fronts for plans generated by BRIGHT for breast. We vary the number of seconds to optimize and understand the algorithm convergence.

Effects of dose calculation points: RQ2.2
Figure 4.8 shows the results of varying the number of DCPs for a fixed budget of 600 seconds. For conciseness and ease of visualization, we show plots of HV and LCI and LSI and LCI in parallel. We highlight the following two key trends:

- **Clinical plan comparison**: across all nine patients, BRIGHT generates plans that improve the HV, LCI, and LSI metrics with respect to the reference solution. Compared to the 2-objective counterpart, we note that the 3-objective approach can create plans that satisfy the clinical aims for out-of-distribution patients 3, 8, and 9. In addition to this, we observe that the 3-objective formulation generates significantly more solutions that are similar to the clinical reference in terms of HV and LCI. This behavior stems from including the third LSI objective, which allows for a broader range of solutions with similar LCI and contV values.
- **Broader objective range**: In addition to producing more solutions in the vicinity of the reference clinical plan, the HV + LCI + LSI technique also includes more plans towards the ”edges” of the two objectives, especially HV. This observation holds across all DCP settings but is readily apparent when analyzing scenarios with fewer points. Consider, for instance, Patient 3, whose plans are analyzed in Figures 4.6c and 4.8e, for the 2- and 3-objective formulations, respectively. When comparing the plans that have the highest LCI values, for instance, with 10000 DCPs, the former is only able to produce results with an LCI objective of ≈ 0.85, while the latter exceeds 0.9. This comes at the cost of HV, whose values for the latter exceed 30 cm^3, while no amount of DCPs in the former reaches this range. This behavior is consistent across all patients.

These observations suggest that the 3-objective formulations may provide a wider set of choices for clinicians to consider than the 2-objective counterpart. The 3-objective BRIGHT algorithm produces more overall solutions and more varied solutions that better cover the extremes of the 2-objective Pareto approximation front. This also leads to more solutions in the immediate neighborhood of the clinical plan in terms of contV and LCI, which is especially visible for Patient 5 (Figures 4.6e and 4.8i). This, however, comes at a cost, as producing plans that dominate the clinical plan takes considerably longer optimization time. While 180 seconds was sufficient for the 2-objective technique to surpass the clinical plan in terms of objective values visibly, this is less noticeable in the 3-objective case. This suggests that the 2-objective formulation can reach a similar search space with less computational resources.

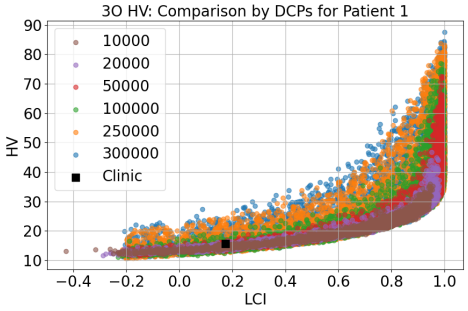

(a) Comparison by DCPs for **Patient 1** (HV + LCI).

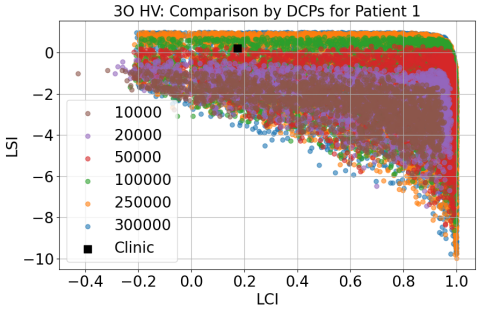

(b) Comparison by DCPs for **Patient 1** (LSI + LCI).

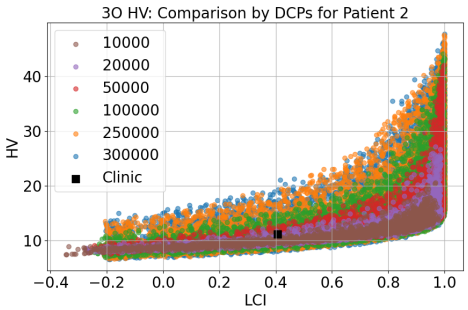

(c) Comparison by DCPs for **Patient 2** (HV + LCI).

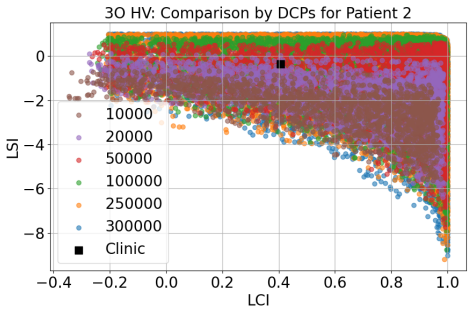

(d) Comparison by DCPs for **Patient 2** (LSI + LCI).

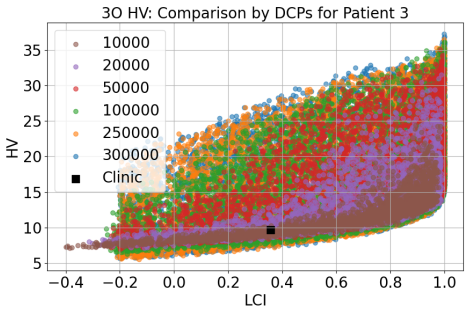

(e) Comparison by DCPs for **Patient 3** (HV + LCI).

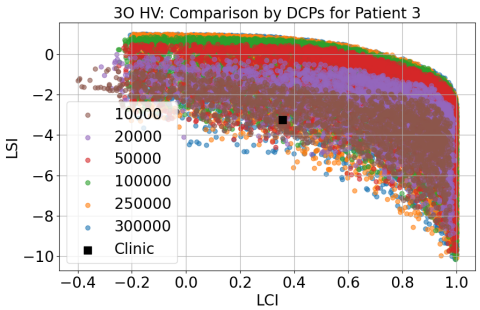

(f) Comparison by DCPs for **Patient 3** (LSI + LCI).

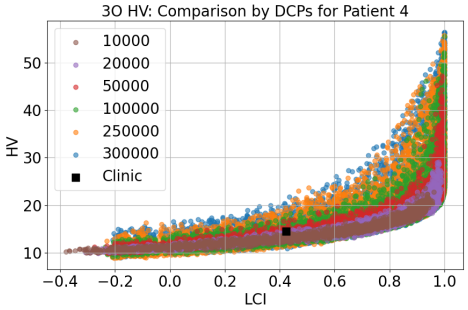

(g) Comparison by DCPs for **Patient 4** (HV + LCI).

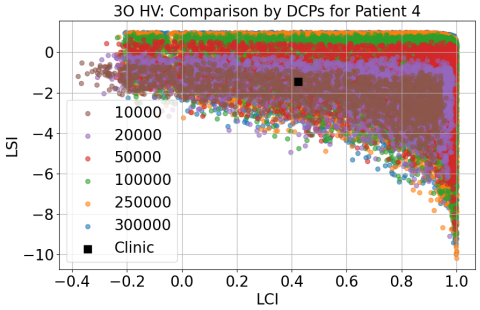

(h) Comparison by DCPs for **Patient 4** (LSI + LCI).

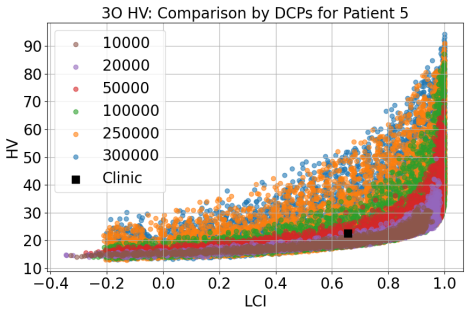

(i) Comparison by DCPs for **Patient 5** (HV + LCI).

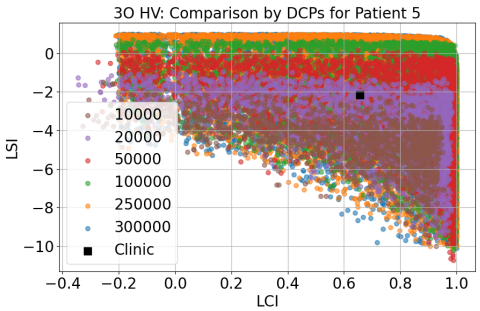

(j) Comparison by DCPs for **Patient 5** (LSI + LCI).

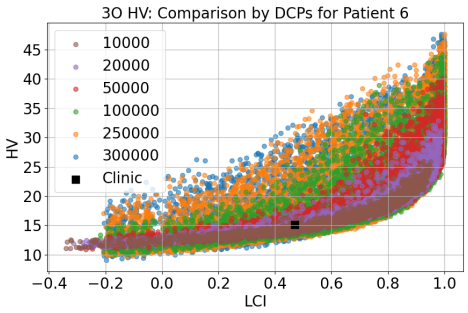

(k) Comparison by DCPs for **Patient 6** (HV + LCI).

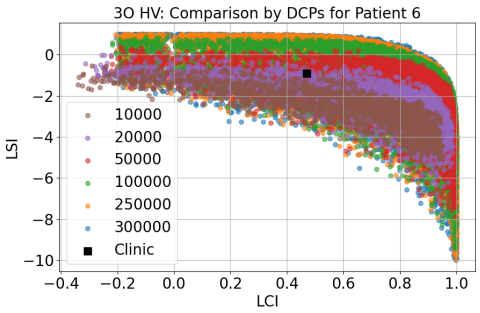

(l) Comparison by DCPs for **Patient 6** (LSI + LCI).

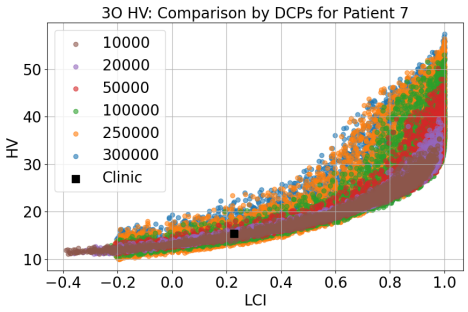

(m) Comparison by DCPs for **Patient 7** (HV + LCI).

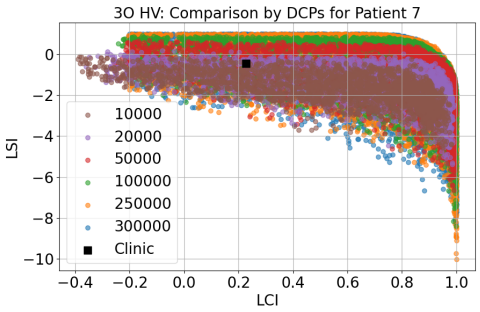

(n) Comparison by DCPs for **Patient 7** (LSI + LCI).

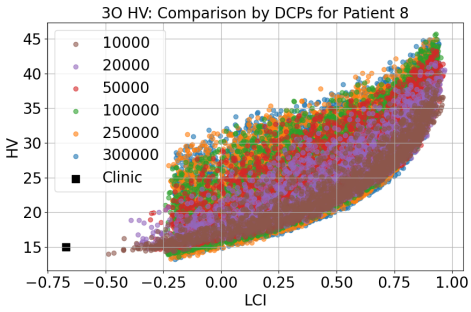

(o) Comparison by DCPs for **Patient 8** (HV + LCI).

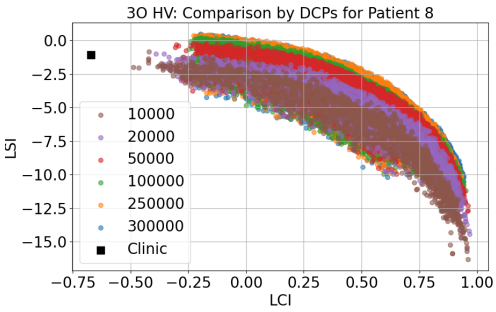

(p) Comparison by DCPs for **Patient 8** (LSI + LCI).

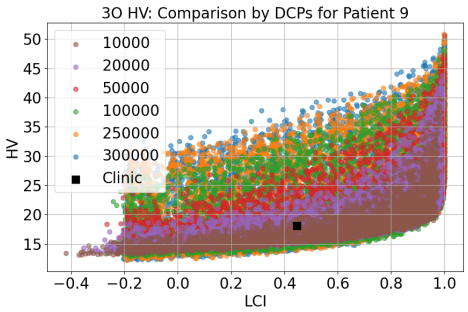

(q) Comparison by DCPs for **Patient 9** (HV + LCI).

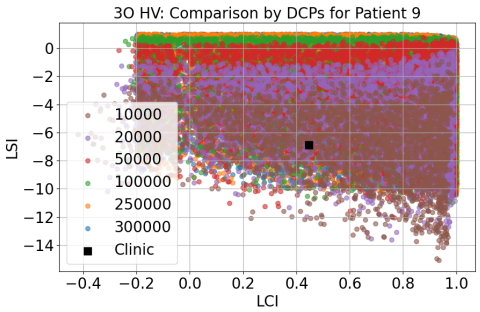

(r) Comparison by DCPs for **Patient 9** (LSI + LCI).

Figure 4.8: Approximation fronts for plans generated by BRIGHT for breast. We vary the number of dose calculation points to understand DCPs trade-offs

.

4.3. Plan Selection Results

Following the computational experiments, the next step within our methodology is plan selection, in order for the expert to assess the clinical acceptance of two plans per patient, with the aim of answering RQ3 formulated in Section 3.5. The visualizations of the plan selection results can be found in Appendix A.

In the context of the 2O DHI + LCI approach, DHI is used as a surrogate for minimizing hotspots by penalizing dose distributions that exhibit high-dose regions (e.g., those exceeding 150% of the prescribed dose). While this approach aligns with the clinical objective of minimizing hotspots, it remains an indirect method. For this reason, we take both the 2O DHI + LCI and the 3O contV + LCI + LSI approaches. The former targets an essential clinical metric, DHI, while the latter directly tackles the spatial distribution of dose, enabling more precise control over the formation of hotspots. As previously described, we enhance the optimization with an additional objective to uncover better trade-offs.

Table 4.1 summarizes the DVIs for the plans chosen as a result of our selection procedure outlined in Section 3.4.

Table 4.1: Plan-specific DVIs for all patients.

Patient	Plan	LCI (Protocol > 0)	DHI	CTV_{1cm} V95%	$CTV_{1.5cm}$ V90%	Ring V100%
Patient 1	P1 Plan 1	0.34	0.87	98.10%	93.40%	1.4 cm^3
	P1 Plan 2	0.81	0.80	99.50%	98.10%	2.2 cm^3
Patient 2	P2 Plan 1	0.54	0.83	99.50%	95.40%	1.1 cm^3
	P2 Plan 2	0.81	0.78	99.80%	98.10%	1.6 cm^3
Patient 3	P3 Plan 1	0.42	0.77	99.10%	94.20%	1.1 cm^3
	P3 Plan 2	0.43	0.81	99.20%	94.30%	3.0 cm^3
Patient 4	P4 Plan 1	0.57	0.83	98.60%	95.70%	1.2 cm^3
	P4 Plan 2	0.50	0.82	98.20%	95.10%	1.8 cm^3
Patient 5	P5 Plan 1	0.68	0.85	99.50%	96.80%	1.3 cm^3
	P5 Plan 2	0.68	0.84	99.50%	96.80%	2.7 cm^3
Patient 6	P6 Plan 1	0.56	0.84	98.30%	95.60%	1.2 cm^3
	P6 Plan 2	0.44	0.82	97.80%	94.50%	1.7 cm^3
Patient 7	P7 Plan 1	0.27	0.81	97.30%	92.70%	1.3 cm^3
	P7 Plan 2	0.22	0.80	97.50%	92.30%	1.5 cm^3
Patient 8	P8 Plan 1	0.0002	0.75	95.00%	90.00%	1.5 cm^3
	P8 Plan 2	0.44	0.749	97.70%	94.40%	4.1 cm^3
Patient 9	P9 Plan 1	0.55	0.68	99.80%	95.50%	1.3 cm^3
	P9 Plan 2	0.59	0.78	98.50%	95.90%	6.4 cm^3

4.4. Clinical Acceptance Results

Following our methodology, we discuss the clinical acceptance results upon performing the computational experiments and selecting BRIGHT plans for clinical evaluation. Before the clinical acceptance results, we illustrate preliminary experiments to shape our view in terms of clinical plan differences compared to BRIGHT plans.

4.4.1. Comparing BRIGHT plans to clinical plans

For visualization purposes and as guides for our method improvement, we have compared a selection of BRIGHT plans to the clinical plans in figures 4.9, 4.10, 4.11. This preliminary evaluation allowed us to better view the differences between clinical and BRIGHT plans.

(a) BRIGHT interface showing Plan 19.

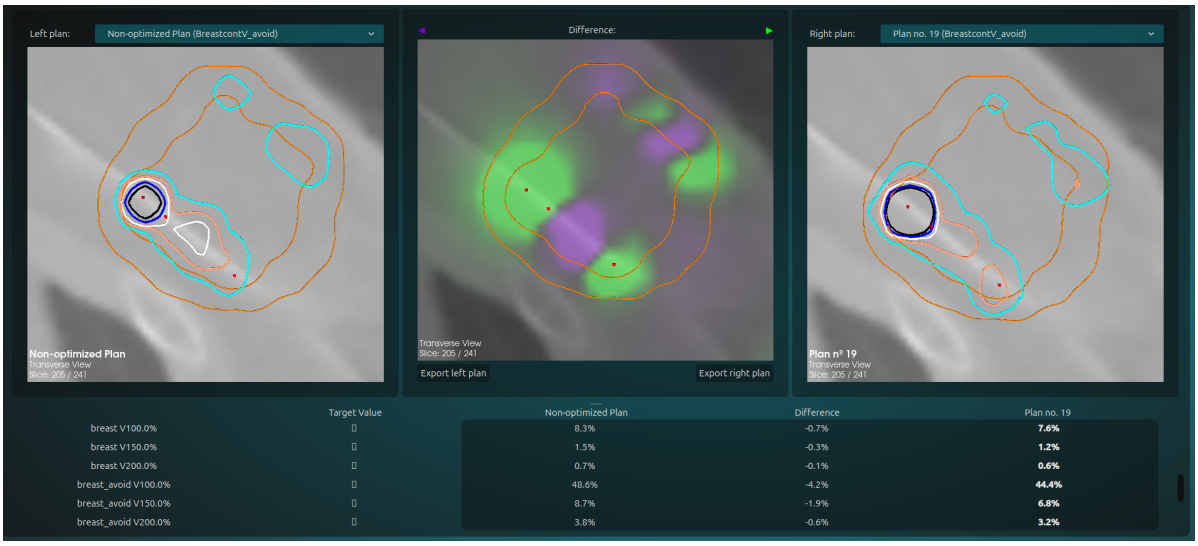

(b) Comparison of a slice from Plan 19 and the clinical plan (leftmost image). The middle image shows the difference between the clinical plan and Plan 19, and the rightmost image represents Plan 19.

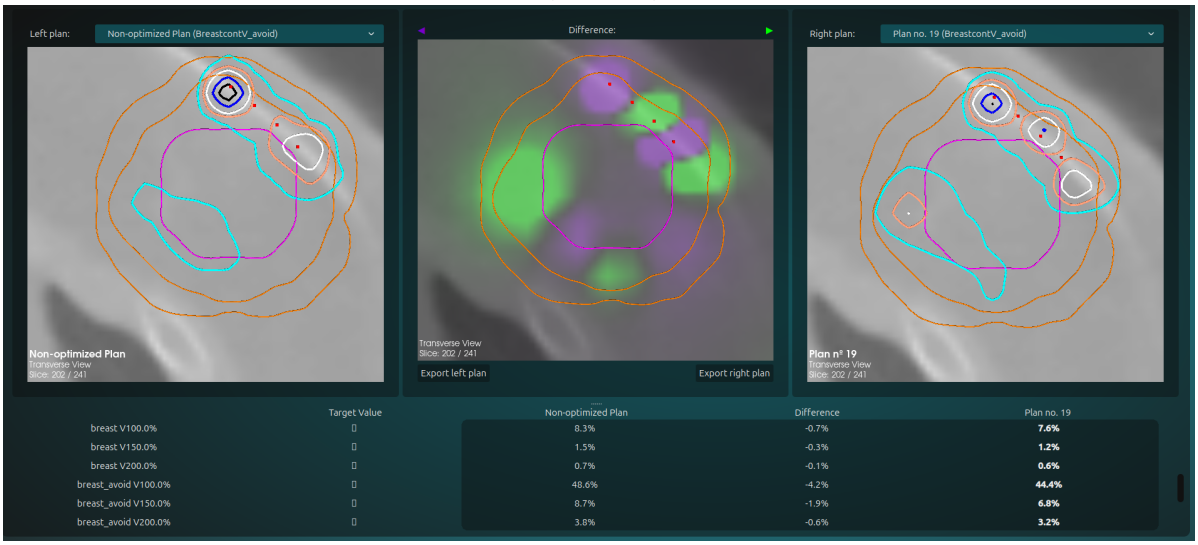

(c) Another slice comparison between Plan 19 and the clinical plan.

Figure 4.9: Visualization of Plan 19 compared to the clinical plan using the BRIGHT interface. The top image shows Plan 19, while the two bottom images display slice comparisons and the differences between Plan 19 and the clinical plan, with purple symbolizing that the dose is higher in the left plan and green that it is higher in the right plan.

(a) BRIGHT interface showing Plan 96.

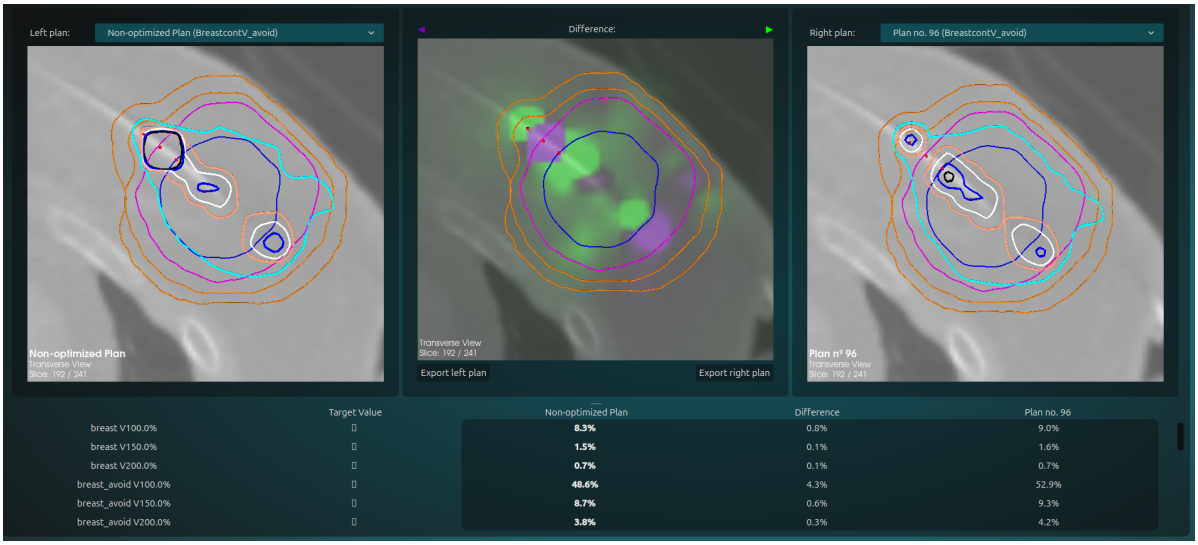

(b) Comparison of a slice from Plan 96 and the clinical plan (leftmost image). The middle image shows the difference between the clinical plan and Plan 96, and the rightmost image represents Plan 96.

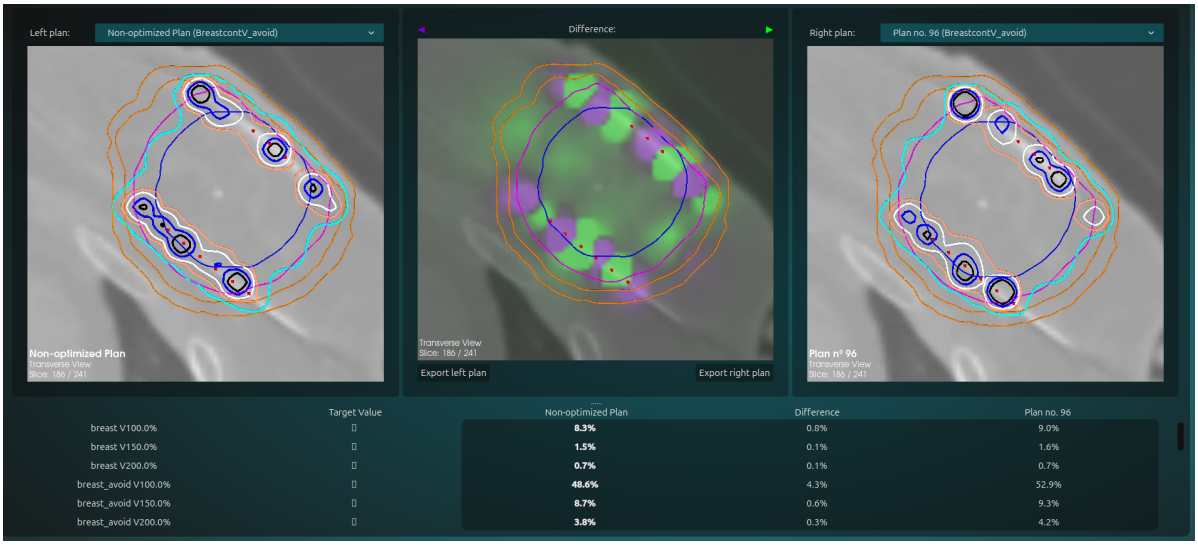

(c) Another slice comparison between Plan 96 and the clinical plan.

Figure 4.10: Visualization of Plan 96 compared to the clinical plan using the BRIGHT interface. The top image shows Plan 96, while the two bottom images display slice comparisons and the differences between Plan 96 and the clinical plan, with purple symbolizing that the dose is higher in the left plan and green that it is higher in the right plan.

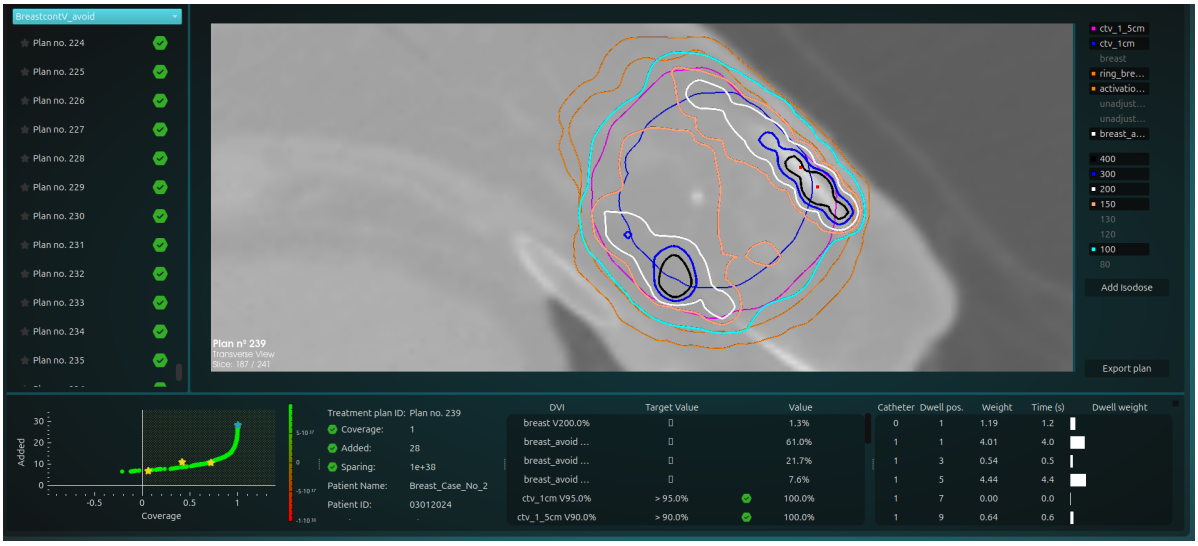

(a) BRIGHT interface showing Plan 239.

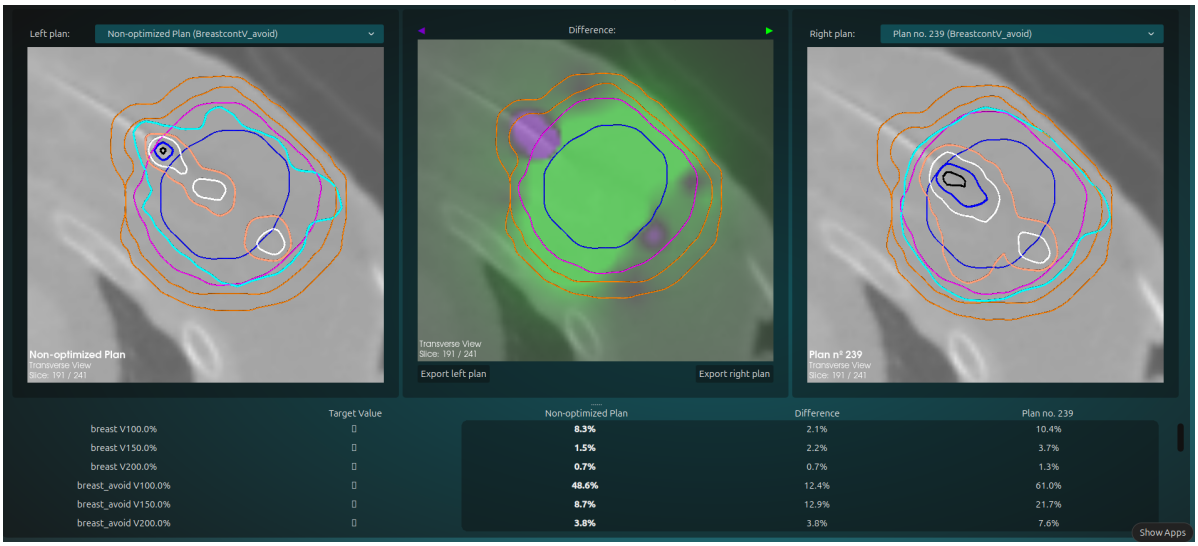

(b) Comparison of a slice from Plan 239 and the clinical plan (leftmost image). The middle image shows the difference between the clinical plan and Plan 239, and the rightmost image represents Plan 239.

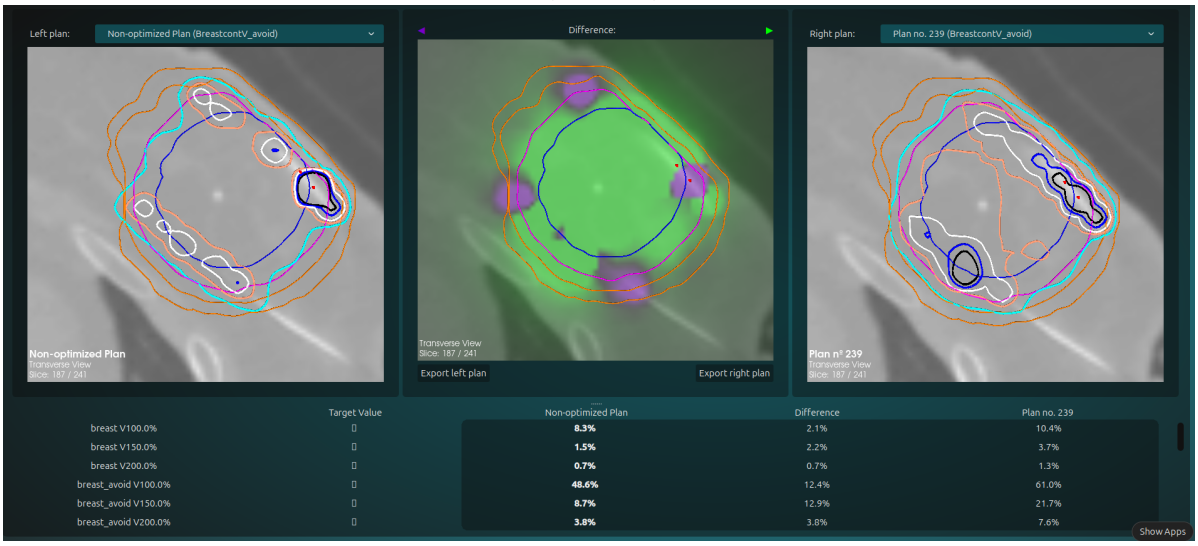

(c) Another slice comparison between Plan 239 and the clinical plan.

Figure 4.11: Visualization of Plan 239 compared to the clinical plan using the BRIGHT interface. The top image shows Plan 239, while the two bottom images display slice comparisons and the differences between Plan 239 and the clinical plan, with purple symbolizing that the dose is higher in the left plan and green that it is higher in the right plan.

4.4.2. Clinical Evaluation

Following the final experiments and plan selection, our expert reviewed the plans individually based on DVH parameters, dose homogeneity, and isodose lines visualized within the BRIGHT interface. Table 4.2 summarizes the clinical acceptance results.

The evaluations revealed consistently positive feedback, with most plans rated clinically acceptable. Experts highlighted the satisfactory DVH parameter values and substantial homogeneity across both optimization techniques, citing values such as 0.83 or 0.85 as excellent. However, some plans exhibited challenges in specific scenarios; for example, Plan 1 for Patient 9 was deemed unacceptable due to suboptimal homogeneity. When looking at this plan, we observe that V100% in the ring is 1.3 cm^3. However, inspecting the clinical plan, we note that V100% within the ring has a value of 7.94 cm^3. This significant difference implies that the hard constraint introduced in the 2-objective DHI optimization is too restrictive for Patient 9, who is deemed as having a suboptimal implant. This means potentially good solutions might be thrown away during hard constraint checking. Meanwhile, employing the additional LSI objective for the 3-objective contV approach guides the optimization toward a softer bound of 6.4 cm^3.

Feedback frequently included recommendations for fine-tuning dwell positions and addressing specific high-dose areas to improve conformity. Overall, the evaluations indicate that BRIGHT plans hold significant promise for clinical use, with areas for refinement identified to enhance future implementations.

Table 4.2: Clinical Evaluation of Plan 1 (2-objective DHI) and Plan 2 (3-objective contV).

Patient	Plan	DVH Satisfaction	Homogeneity Satisfaction	Acceptable	Additional Feedback
Patient 1	1	Strongly agree	Strongly agree (0.87)	Yes	Great coverage and excellent DHI.
	2	Strongly agree	Strongly agree (0.8)	Yes	Adjust dwell positions to reduce skin dose (slice 47-10).
Patient 2	1	Strongly agree	Strongly agree (0.83)	Yes	Excellent coverage, very good DHI.
	2	Strongly agree	Strongly agree (0.78)	Yes	Excellent coverage, very good DHI. Prefers Plan 1 for higher DHI.
Patient 3	1	Strongly agree	Strongly agree (0.77)	Yes	Adjust coverage (excellent) to improve DHI.
	2	Strongly agree	Strongly agree (0.81)	Yes	Improved DHI but higher irradiated tissue outside CTV_1.5cm.
Patient 4	1	Strongly agree	Strongly agree (0.83)	Yes	Excellent coverage.
	2	Strongly agree	Strongly agree (0.82)	Yes	Practically indistinguishable from Plan 1.
Patient 5	1	Strongly agree	Strongly agree (0.85)	Yes	Can activate posterior catheter (slice 42-56).
	2	Strongly agree	Strongly agree (0.84)	Yes	Prefers Plan 1 as it is more conformal.
Patient 6	1	Strongly agree	Strongly agree (0.84)	Yes	Excellent coverage, very good DHI.
	2	Strongly agree	Strongly agree (0.82)	Yes	Prefers Plan 1.
Patient 7	1	Strongly agree	Strongly agree (0.81)	Yes	Excellent coverage, very good DHI.
	2	Strongly agree	Strongly agree (0.80)	Yes	Prefers Plan 1.
Patient 8	1	Strongly agree	Strongly agree (0.75)	Yes	Reduce skin dose exceeding 100%.
	2	Strongly agree	Strongly agree (0.749)	Yes	Improved coverage over Plan 1.
Patient 9	1	Strongly agree	Disagree (0.68)	No	Excessive coverage; improve DHI.
	2	Strongly agree	Disagree (0.78)	Yes	Reduce coverage to improve DHI and conformality.

5

Discussion

This study investigates and confirms that BRIGHT provides an effective method for generating high-quality treatment plans by simultaneously optimizing multiple competing objectives, such as minimizing hotspots and maximizing target coverage.

The results of the time and DCP experiments demonstrate that careful choices for optimization time and number of dose calculation points can result in improvements over the clinical plans for our available patients. As observed in the figures comparing various time budgets, longer optimization times (e.g., 600 seconds) led to improved solutions. This trend was evident across all patients, with the optimization algorithm producing more refined Pareto approximation fronts as time increased, enabling more effective balancing of the conflicting objectives. Moreover, even a time budget of 180 seconds can prove sufficient, for example, when using the contV + LCI optimization setup for patients 1, 2, and 4.

Furthermore, increasing the number of dose calculation points from 10,000 to 100,000 had a noticeable effect on the homogeneity (and sparing constraint for 3-objective) of the results for the majority of patients, especially for the 3-objective cases. More DCPs allow for a finer resolution in dose distribution calculations, which, in turn, leads to better accuracy in controlling homogeneity. However, this increase in precision comes at the cost of higher computational complexity, making it more challenging to achieve a high-quality solution within shorter time budgets. In terms of the minimum resources required to generate results that are at least on par with clinical benchmarks, DCP, and time experiments display similar trends. While the highest number of DCPs generally showed the best results for most patients, we observed that 100,000 and occasionally 50,000 DCPs are often sufficient to provide results that match or exceed the reference clinical plan across most setups.

Despite the observed improvements with more extended time budgets and higher DCPs, the diminishing returns of further increasing either parameter were also apparent. While the gains in solution quality were substantial when moving from 30 to 600 seconds or 10,000 to 100,000 DCPs, additional increases (e.g., 900 seconds or 300,000 DCPs) showed relatively marginal improvements compared to the added computational burden. This highlights the need to balance computation time, accuracy, and clinical applicability.

A main takeaway from the results is the effectiveness of using the 2-objective optimization approach focused on contiguous volumes and the Least Coverage Index. The comparison between the 2- and 3-objective techniques (adding the Least Sparing Index) highlighted key differences. While the 3-objective technique provided more varied solutions, it also introduced more complexity to the decision-making process. The 2-objective approach explored the search space thoroughly and generated Pareto approximation fronts that were easier to interpret and use for clinical decisions. This simplicity is critical, as it directly influences how effectively clinicians can compare and select treatment plans.

The expert feedback further highlights the practical importance of BRIGHT in the clinic. Of the 18 selected treatment plans, the clinician described the coverage as either *great* or *excellent* in 16 instances and deemed 17 plans to be clinically applicable. This comes in addition to the objective metrics, in which BRIGHT solutions outperform clinical reference plans consistently for numerous configurations.

It is noteworthy that these results originate from a run with a time budget of just 10 minutes on contemporary hardware. Providing appropriate plan selection mechanisms underlines not only the viability of applying BRIGHT for breast brachytherapy but also the immediate value it can provide to clinics and patients today.

Regarding limitations, a significant observation from the clinical acceptance evaluation is the risk that experts, despite the blinding, may recognize the clinical plan when comparing it with BRIGHT-generated solutions. This familiarity could introduce bias in the subjective assessments of plan quality. The sample size of 9 patients is another limitation, as it restricts the generalizability of the findings. While the patients represented a range of anatomical complexities, a broader sample would provide a more robust evaluation of BRIGHT's capabilities.

Two other limitations stem from the scope of this study. While we addressed several formulations of contV- and DHI-based optimization problems, the detailed analysis of the individual plans' features falls outside this scope. The limited dataset and clinical exposure make it difficult to compare the fine details of when each method may excel or under which circumstances they fall short. Such analyses would require a broader dataset and, as such, are left to future work. Similarly, integrating BRIGHT with external and complementary optimization methods is also left to future work. Such integrations could materialize as optimization problems with more degrees of freedom by also considering, for instance, catheter placement optimization simultaneously.

The thesis demonstrates the advantages of using BRIGHT for breast brachytherapy treatment planning, and the use of evolutionary algorithms in this context shows significant promise. Future work could focus on validating these findings with a larger patient cohort and investigating the impact of expert bias in clinical evaluations. Moreover, exploring hybrid models that integrate human expertise with automated planning may further enhance the clinical usability of BRIGHT-generated plans.

6

Conclusion

This study demonstrates the use of BRIGHT to optimize breast cancer brachytherapy treatment plans, focusing on objectives that balance dose homogeneity, contiguous volume, and clinical coverage and sparing.

We compare different objective formulations—primarily between two-objective and three-objective, as well as the homogeneity index and hotspot volumes optimization techniques—by varying optimization time and the number of dose calculation points.

After performing computational experiments, we determined that the 2-objective contV + LCI approach with a time budget of 600 seconds and 250,000 DCPs provided the best balance between solution quality and computational complexity. We evaluate the obtained BRIGHT plans with the help of a clinical expert and receive insightful feedback for further improvement and use of BRIGHT.

The results obtained in this study and previous BRIGHT applications set a strong foundation for adopting BRIGHT in clinical practice for further use cases and promoting automated optimization methods such as evolutionary algorithms.

7

Acknowledgements

I want to thank our clinical expert from Virginia Commonwealth University, Dorin Todor, for his support and guidance throughout this project. His input was invaluable in helping me connect my research to the real-world field of brachytherapy and gain a deeper appreciation for its challenges and prospects. I appreciate the time you put into sharing your feedback, which has made this a more meaningful learning experience for me.

I also want to thank my professors, Peter Bosman and Tanja Alderliesten, and supervisors Anton Bouter and Leah Dickhoff, for their guidance and encouragement along the way. Your insights and support have been essential in my journey, and I appreciate the time I spent at CWI getting a peek into the research group experience.

References

[1] Anton Bouter et al. “GPU-accelerated bi-objective treatment planning for prostate high-dose-rate brachytherapy”. In: *Medical physics* 46.9 (2019), pp. 3776–3787.

[2] Kiven Erique Lukong. “Understanding breast cancer–The long and winding road”. In: *BBA clinical* 7 (2017), pp. 64–77.

[3] Kathryn P Trayes and Sarah EH Cokenakes. “Breast cancer treatment”. In: *American family physician* 104.2 (2021), pp. 171–178.

[4] Elaine M Zeman, Eric C Schreiber, and Joel E Tepper. “Basics of radiation therapy”. In: *Abeloff’s clinical oncology*. Elsevier, 2020, pp. 431–460.

[5] Vratislav Strnad et al. “Phase III multicenter trial: interstitial brachytherapy alone versus external beam radiation therapy after breast conserving surgery for low-risk invasive carcinoma and low-risk ductal carcinoma in-situ of the female breast–one-year toxicities”. In: *Brachytherapy* 12 (2013), S15.

[6] Salvatore Cozzi et al. “Advantages of intraoperative implant for interstitial brachytherapy for accelerated partial breast irradiation either frail patients with early-stage disease or in locally recurrent breast cancer”. In: *Journal of contemporary brachytherapy* 10.2 (2018), pp. 97–104.

[7] Salvatore Cozzi et al. “The role of interstitial brachytherapy for breast cancer treatment: an overview of indications, applications, and technical notes”. In: *Cancers* 14.10 (2022), p. 2564.

[8] Peter Nygren. “What is cancer chemotherapy?” In: *Acta Oncologica* 40.2-3 (2001), pp. 166–174.

[9] CD Lee. “Recent developments and best practice in brachytherapy treatment planning”. In: *The British journal of radiology* 87.1041 (2014), p. 20140146.

[10] Jaroslaw T Hepel et al. “American Brachytherapy Society consensus report for accelerated partial breast irradiation using interstitial multicatheter brachytherapy”. In: *Brachytherapy* 16.5 (2017), pp. 919–928.

[11] Chirag Shah et al. “Modern approaches for breast brachytherapy”. In: *Seminars in Radiation Oncology*. Vol. 30. 1. Elsevier. 2020, pp. 61–67.

[12] Kalyanmoy Deb. “Multi-objective optimisation using evolutionary algorithms: an introduction”. In: *Multi-objective evolutionary optimisation for product design and manufacturing*. Springer, 2011, pp. 3–34.

[13] A Karabis, S Giannouli, D Baltas, et al. “40 HIPO: a hybrid inverse treatment planning optimization algorithm in HDR brachytherapy”. In: *Radiotherapy and Oncology* 76.2 (2005), S29.

[14] Daniel Faverly, Roland Holland, and Lambert Burgers. “An original stereomicroscopic analysis of the mammary glandular tree”. In: *Virchows Archiv A* 421 (1992), pp. 115–119.

[15] Bruce G Haffty et al. “Local recurrence versus new primary: clinical analysis of 82 breast relapses and potential applications for genetic fingerprinting”. In: *International Journal of Radiation Oncology* Biology* Physics* 27.3 (1993), pp. 575–583.

[16] Hiroshi Imamura et al. “Relationship between the morphological and biological characteristics of intraductal components accompanying invasive ductal breast carcinoma and patient age”. In: *Breast cancer research and treatment* 62 (2000), pp. 177–184.

[17] Neal S Goldstein, Larry Kestin, and Frank Vicini. “Factors associated with ipsilateral breast failure and distant metastases in patients with invasive breast carcinoma treated with breast-conserving therapy: a clinicopathologic study of 607 neoplasms from 583 patients”. In: *American journal of clinical pathology* 120.4 (2003), pp. 500–527.

[18] Yimin Wang et al. “Prevention and treatment for radiation-induced skin injury during radiotherapy”. In: *Radiation Medicine and Protection* 1.02 (2020), pp. 60–68.

[19] Danique LJ Barten 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”. In: *Brachytherapy* 22.2 (2023), pp. 279–289.

[20] Etienne Lessard and Jean Pouliot. “Inverse planning anatomy-based dose optimization for HDR-brachytherapy of the prostate using fast simulated annealing algorithm and dedicated objective function”. In: *Medical physics* 28.5 (2001), pp. 773–779.

[21] Timmy Siauw et al. “IPIP: A new approach to inverse planning for HDR brachytherapy by directly optimizing dosimetric indices”. In: *Medical Physics* 38.7 (2011), pp. 4045–4051.

[22] Timo M Deist and Bram L Gorissen. “High-dose-rate prostate brachytherapy inverse planning on dose-volume criteria by simulated annealing”. In: *Physics in Medicine & Biology* 61.3 (2016), p. 1155.

[23] Ngoc Hoang Luong et al. “Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment”. In: *Swarm and Evolutionary Computation* 40 (2018), pp. 37–52.

[24] Anton Bouter et al. “The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm”. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. 2017, pp. 537–544.

[25] Danique LJ Barten et al. “Artificial intelligence based planning of HDR prostate brachytherapy: first clinical experience”. In: (2021).

[26] Leah RM Dickhoff et al. “Keeping your best options open with AI-based treatment planning in prostate and cervix brachytherapy”. In: *Brachytherapy* (2024).

[27] Peter AN Bosman and Dirk Thierens. “More concise and robust linkage learning by filtering and combining linkage hierarchies”. In: *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. 2013, pp. 359–366.

[28] Peter AN Bosman and Dirk Thierens. “Linkage neighbors, optimal mixing and forced improvements in genetic algorithms”. In: *Proceedings of the 14th annual conference on Genetic and evolutionary computation*. 2012, pp. 585–592.

[29] Kalyanmoy Deb et al. “A fast and elitist multiobjective genetic algorithm: NSGA-II”. In: *IEEE transactions on evolutionary computation* 6.2 (2002), pp. 182–197.

[30] 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*. 2014, pp. 357–364.

[31] 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*. 2010, pp. 351–358.

[32] Roja Zakariaee et al. “Association of bladder dose with late urinary side effects in cervical cancer high-dose-rate brachytherapy”. In: *Brachytherapy* 16.6 (2017), pp. 1175–1183.

[33] Texas University MD Anderson Cancer Center. *Side effects of radiation therapy for breast cancer*. 2023. URL: `https://www.mdanderson.org/cancerwise/side-effects-of-radiation-therapy-for-breast-cancer.h00-159615489.html` (visited on 06/23/2024).

[34] Jeffrey M Straub et al. “Radiation-induced fibrosis: mechanisms and implications for therapy”. In: *Journal of cancer research and clinical oncology* 141 (2015), pp. 1985–1994.

[35] Joost Commandeur. “Improving the homogeneity of brachytherapy treatment plans generated by BRIGHT using a hotspot registration method based on connected component analysis”. In: (2021).

[36] Björn Morén, Torbjörn Larsson, and Åsa Carlsson Tedgren. “Preventing hot spots in high dose-rate brachytherapy”. In: *Operations Research Proceedings 2017: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR), Freie Universiät Berlin, Germany, September 6-8, 2017*. Springer. 2018, pp. 369–375.

[37] Marjolein C van der Meer et al. “Bi-objective optimization of catheter positions for high-dose-rate prostate brachytherapy”. In: *Medical physics* 47.12 (2020), pp. 6077–6086.

[38] S Giannouli et al. “Autoactivation of source dwell positions for HDR brachytherapy treatment planning”. In: *Medical Physics* 27.11 (2000), pp. 2517–2520.

[39] Anton Bouter et al. “Large-scale parallelization of partial evaluations in evolutionary algorithms for real-world problems”. In: *Proceedings of the genetic and evolutionary computation conference*. 2018, pp. 1199–1206.

[40] Leah RM Dickhoff et al. “Adaptive objective configuration in bi-objective evolutionary optimization for cervical cancer brachytherapy treatment planning”. In: *Proceedings of the Genetic and Evolutionary Computation Conference*. 2022, pp. 1173–1181.

A

Plan selection results

The following figures describe the highlighted plans during the selection process described in Section 3.4. The plans used in the clinical evaluation cover two main techniques: 2-objective DHI and 3-objective HV plans.

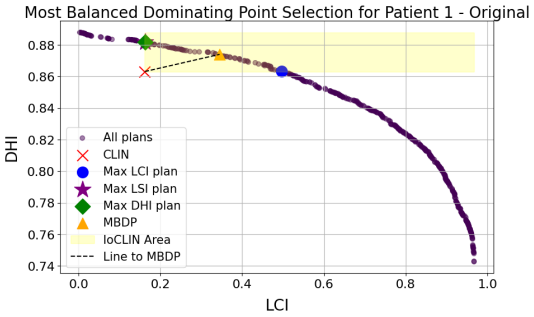

(a) Patient 1: Original plot (DHI + LCI).

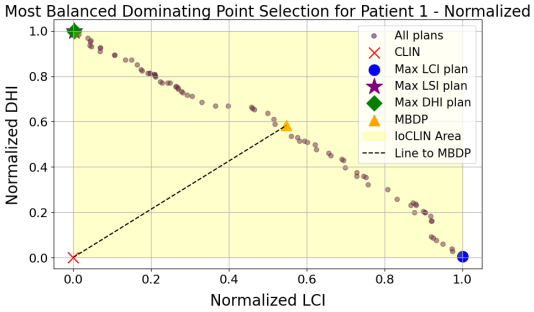

(b) Patient 1: Normalized plot (DHI + LCI).

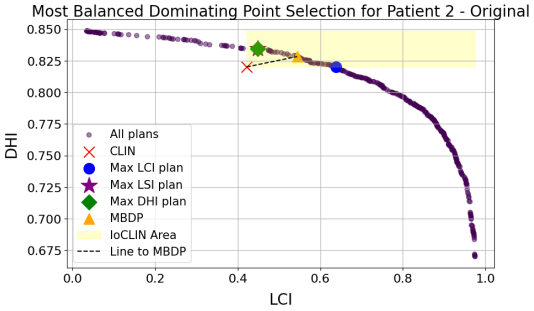

(c) Patient 2: Original plot (DHI + LCI).

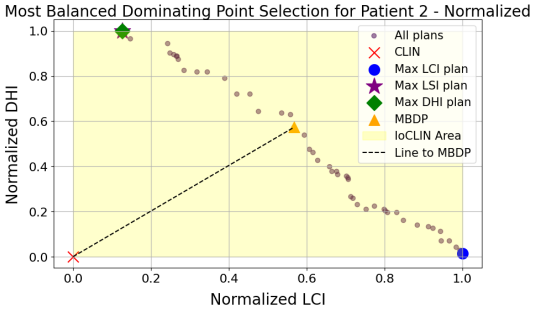

(d) Patient 2: Normalized plot (DHI + LCI).

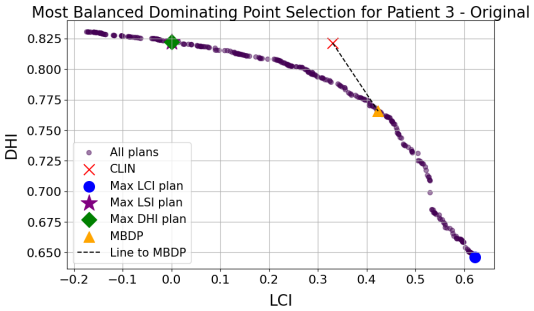

(e) Patient 3: Original plot (DHI + LCI).

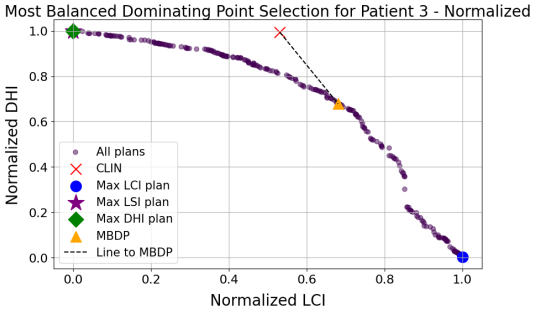

(f) Patient 3: Normalized plot (DHI + LCI).

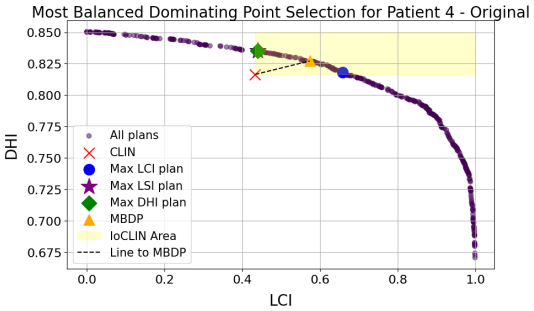

(g) Patient 4: Original plot (DHI + LCI).

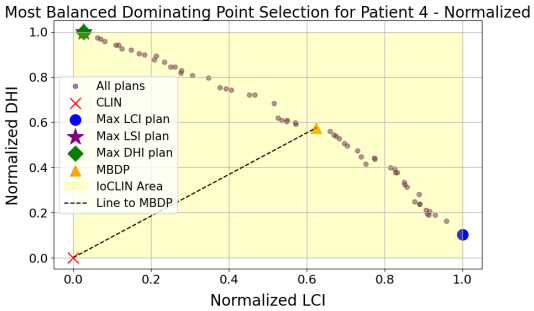

(h) Patient 4: Normalized plot (DHI + LCI).

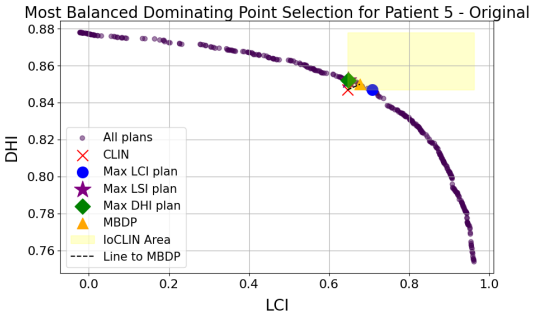

(i) Patient 5: Original plot (DHI + LCI).

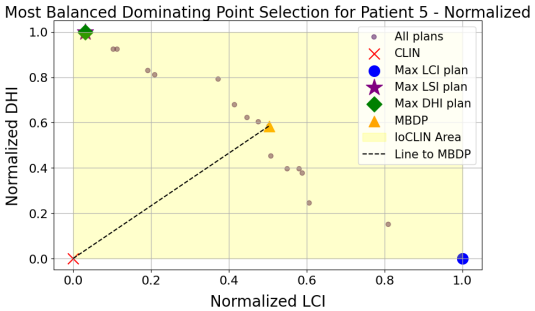

(j) Patient 5: Normalized plot (DHI + LCI).

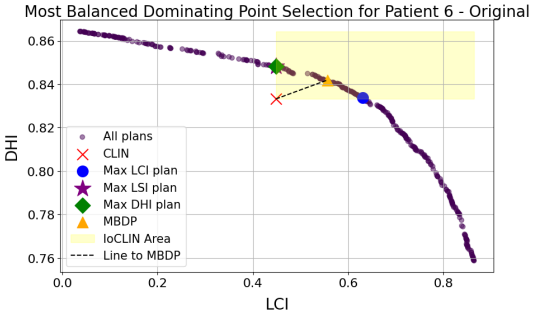

(k) Patient 6: Original plot (DHI + LCI).

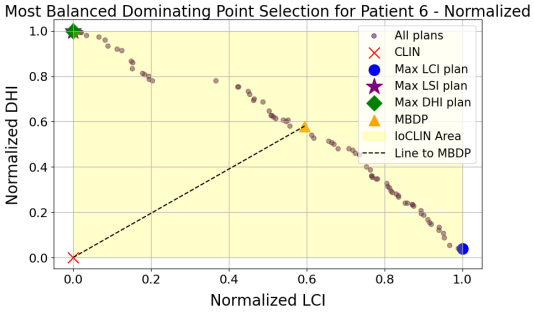

(l) Patient 6: Normalized plot (DHI + LCI).

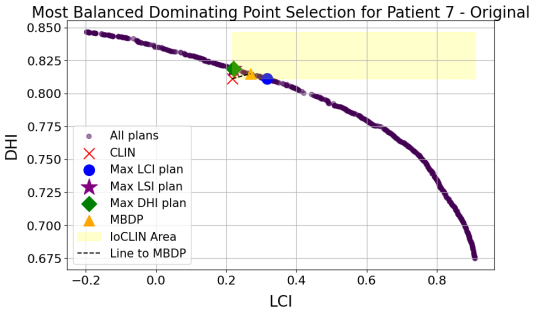

(m) Patient 7: Original plot (DHI + LCI).

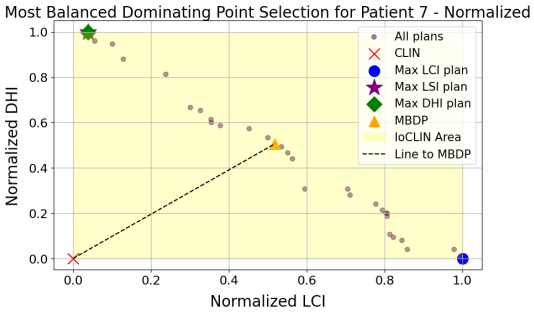

(n) Patient 7: Normalized plot (DHI + LCI).

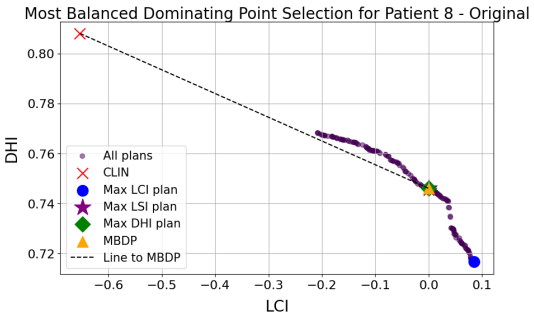

(o) Patient 8: Original plot (DHI + LCI).

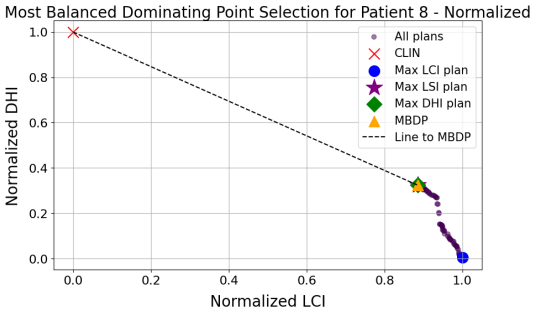

(p) Patient 8: Normalized plot (DHI + LCI).

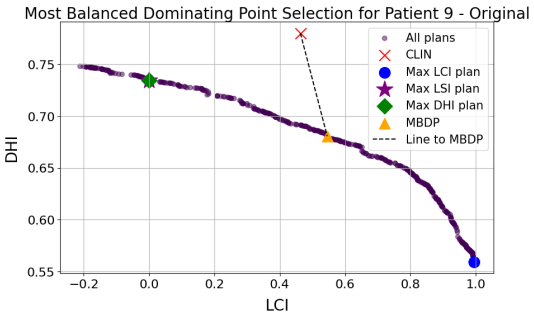

(q) Patient 9: Original plot (DHI + LCI).

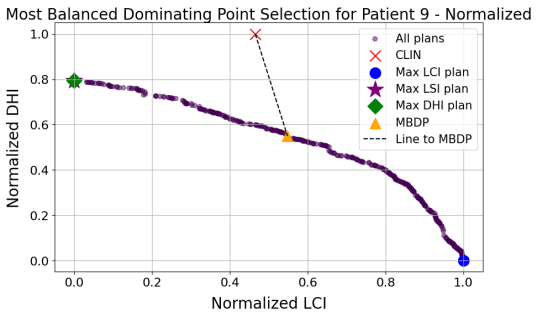

(r) Patient 9: Normalized plot (DHI + LCI).

Figure A.1: Comparison of original and normalized plots for patients 6 to 9 in the 2-objective (2O) DHI optimization setup.

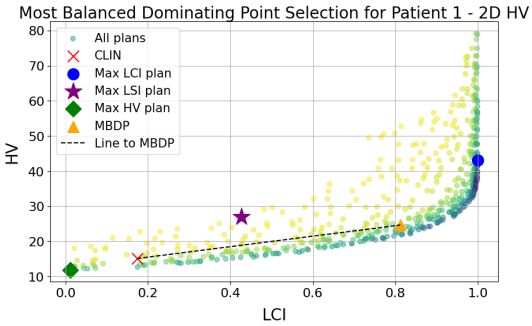

(a) 2D HV plot for Patient 1.

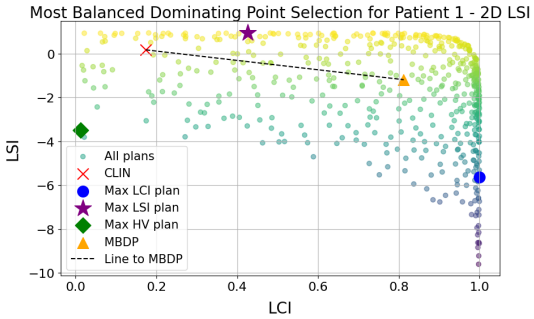

(b) 2D LSI plot for Patient 1.

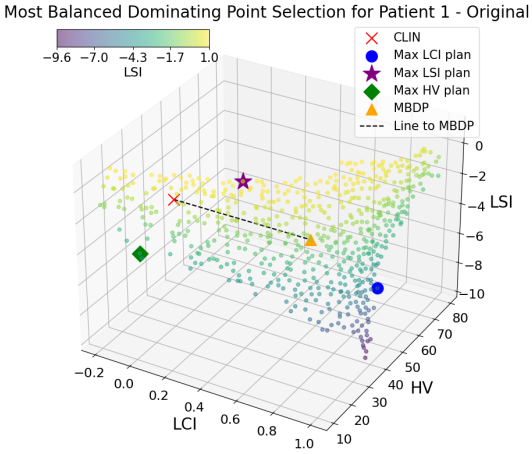

(c) Original 3D plot for Patient 1.

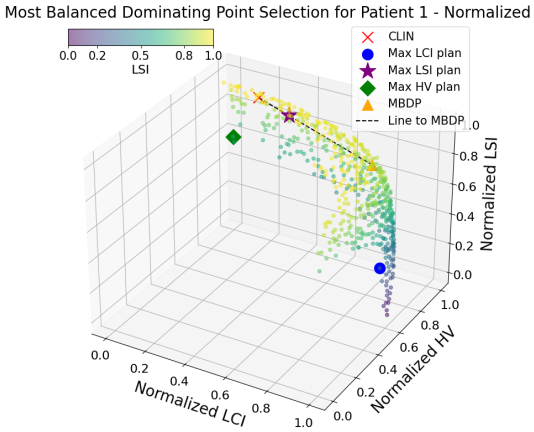

(d) Normalized 3D plot for Patient 1.

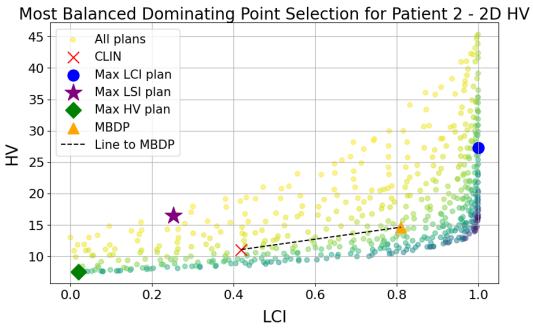

(e) 2D HV plot for Patient 2.

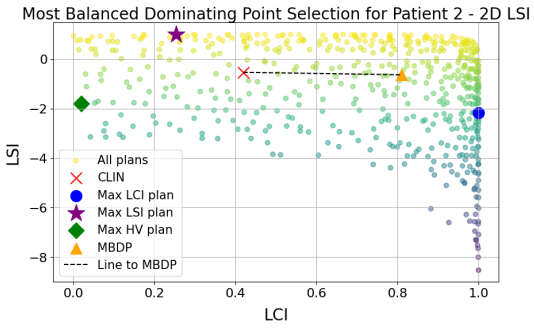

(f) 2D LSI plot for Patient 2.

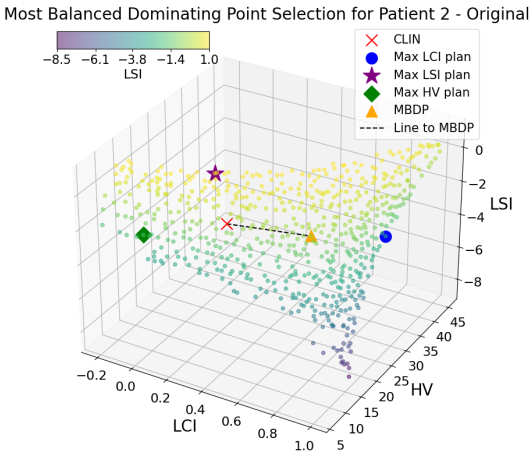

(g) Original 3D plot for Patient 2.

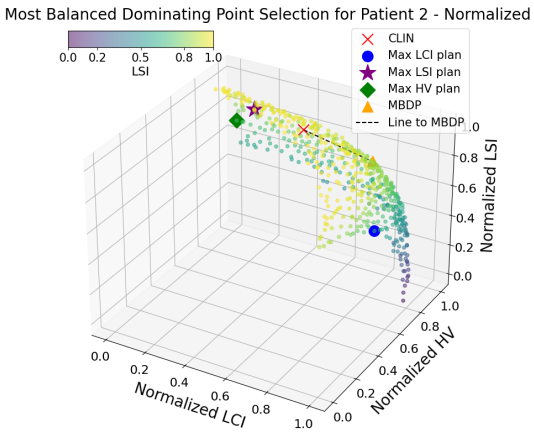

(h) Normalized 3D plot for Patient 2.

Figure A.2: Plan selection results for Patients 1,2, showing the comparison of 2D and 3D plots highlighting the most balanced points.

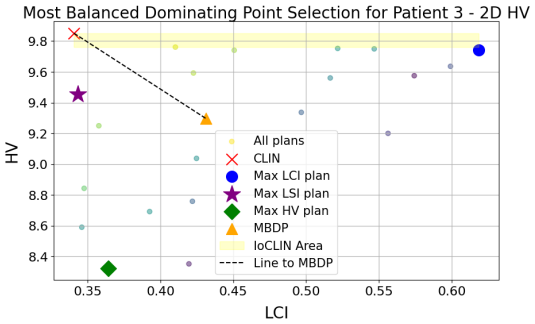

(a) 2D HV plot for Patient 3.

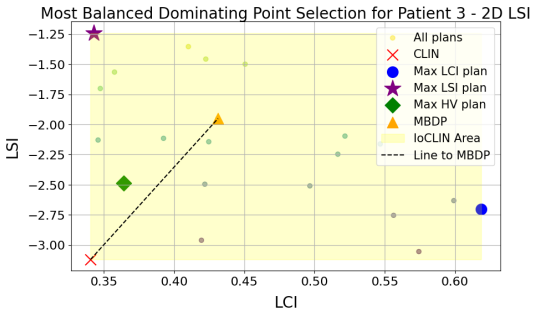

(b) 2D LSI plot for Patient 3.

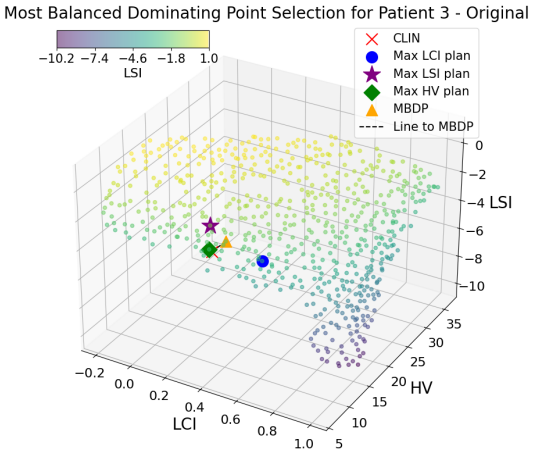

(c) Original 3D plot for Patient 3.

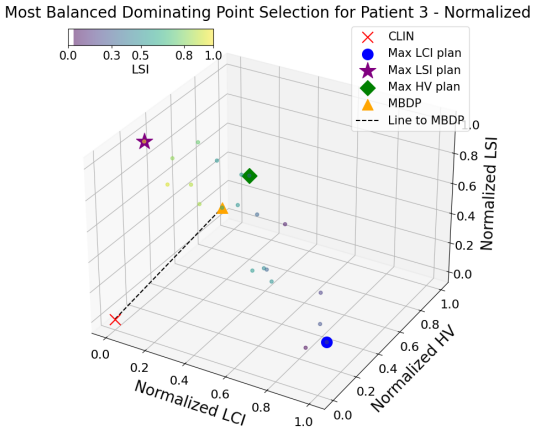

(d) Normalized 3D plot for Patient 3.

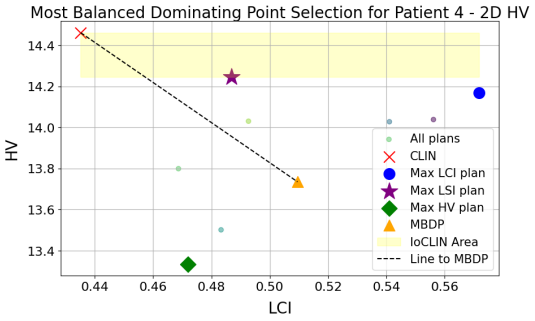

(e) 2D HV plot for Patient 4.

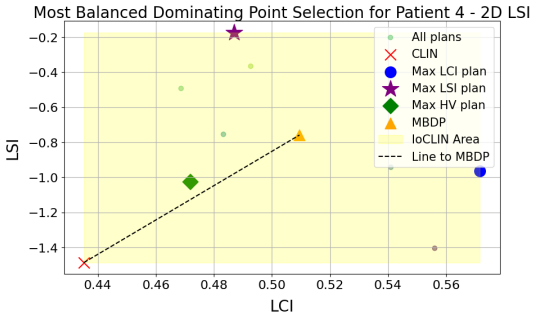

(f) 2D LSI plot for Patient 4.

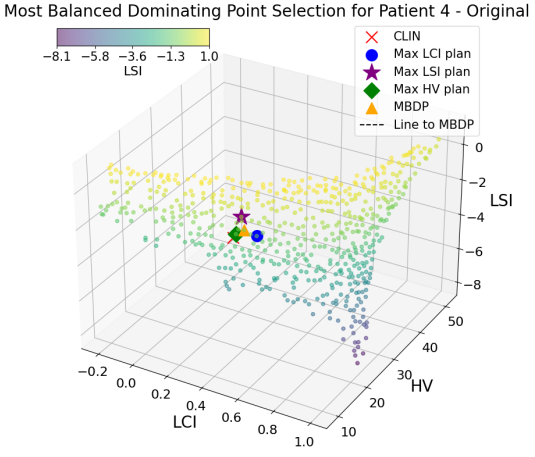

(g) Original 3D plot for Patient 4.

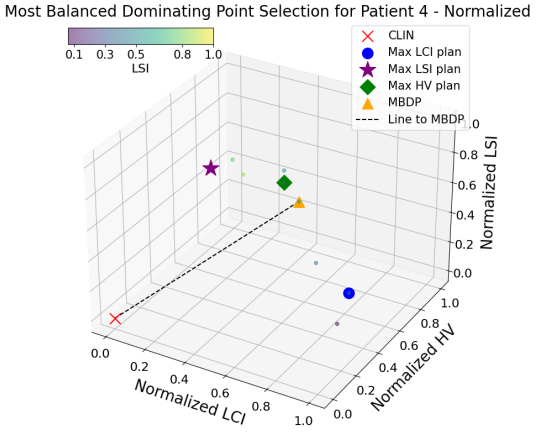

(h) Normalized 3D plot for Patient 4.

Figure A.3: Plan selection results for Patients 3,4, showing the comparison of 2D and 3D plots highlighting the most balanced points.

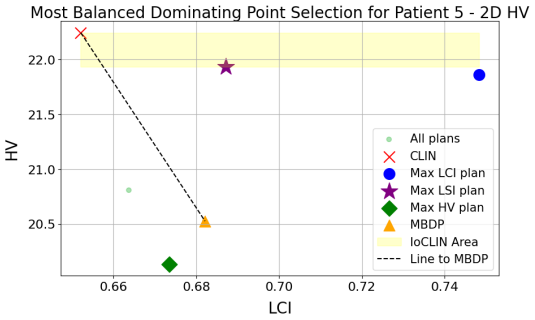

(a) 2D HV plot for Patient 5.

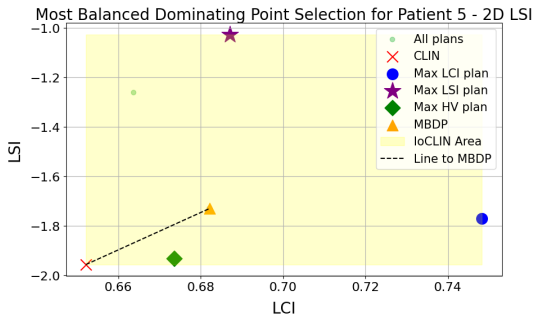

(b) 2D LSI plot for Patient 5.

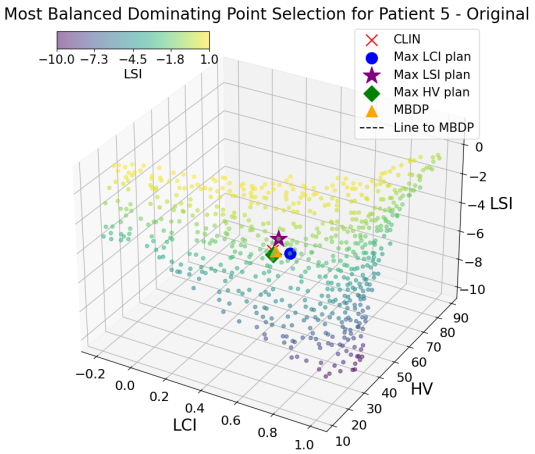

(c) Original 3D plot for Patient 5.

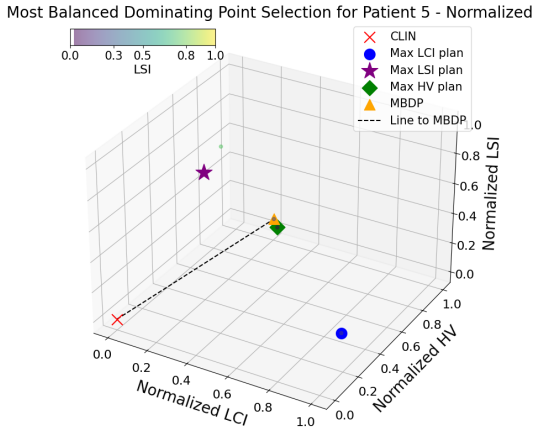

(d) Normalized 3D plot for Patient 5.

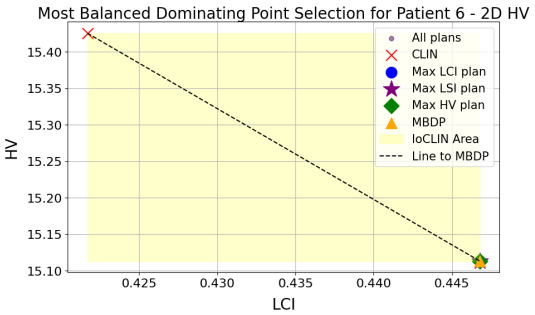

(e) 2D HV plot for Patient 6.

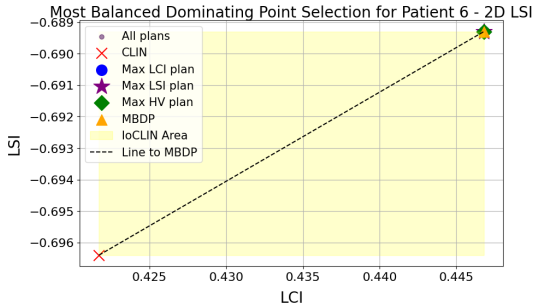

(f) 2D LSI plot for Patient 6.

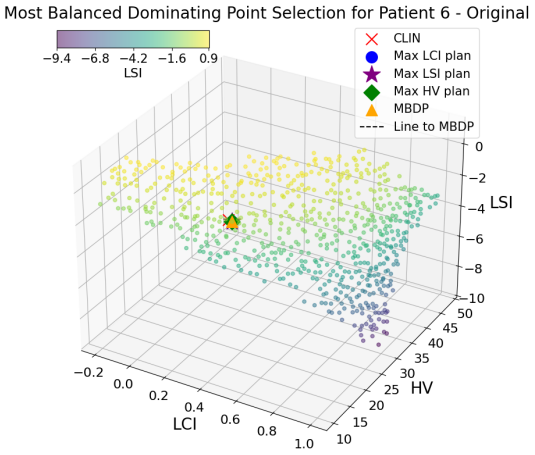

(g) Original 3D plot for Patient 6.

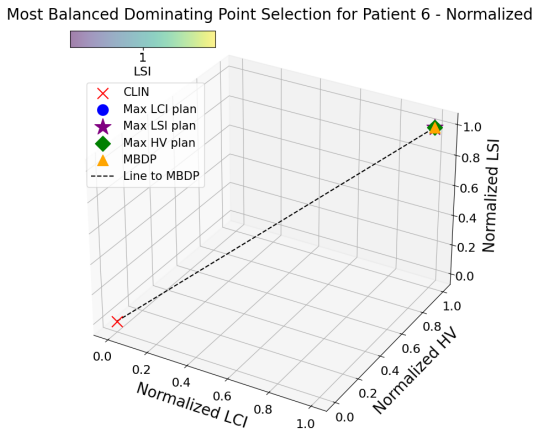

(h) Normalized 3D plot for Patient 6.

Figure A.4: Plan selection results for Patients 5,6, showing the comparison of 2D and 3D plots highlighting the most balanced points.

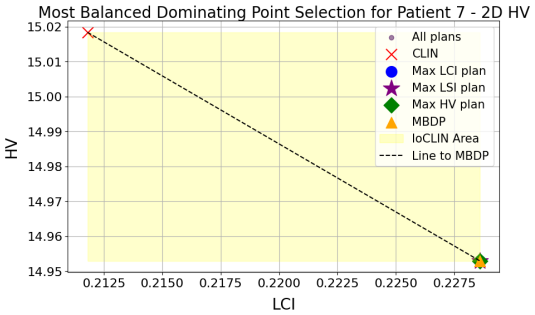

(a) 2D HV plot for Patient 7.

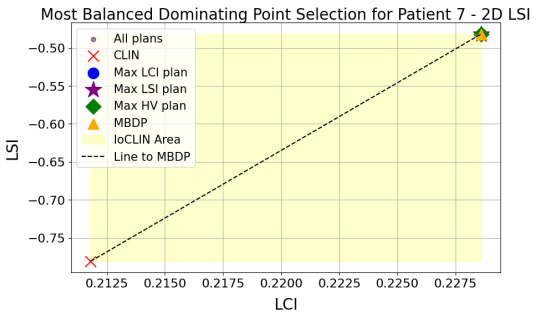

(b) 2D LSI plot for Patient 7.

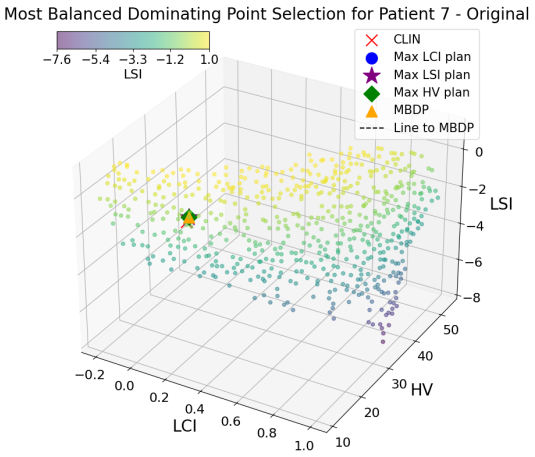

(c) Original 3D plot for Patient 7.

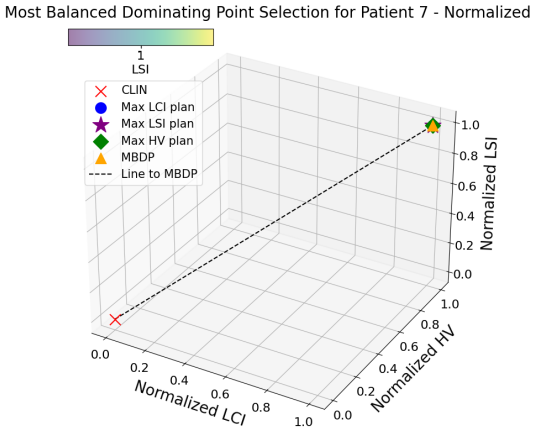

(d) Normalized 3D plot for Patient 7.

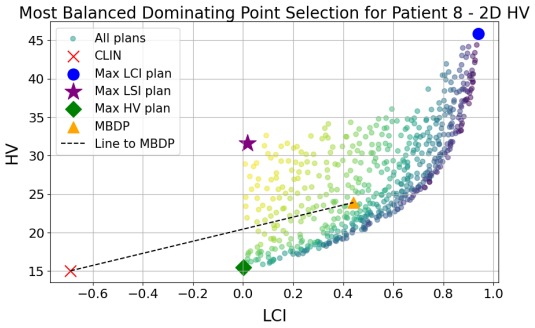

(e) 2D HV plot for Patient 8.

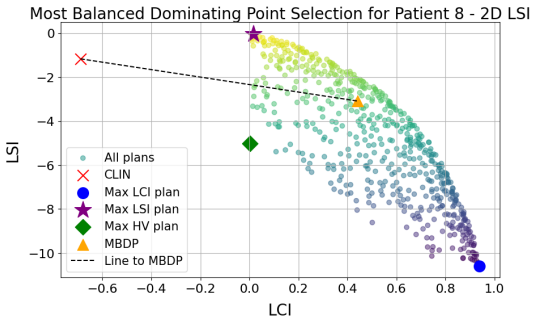

(f) 2D LSI plot for Patient 8.

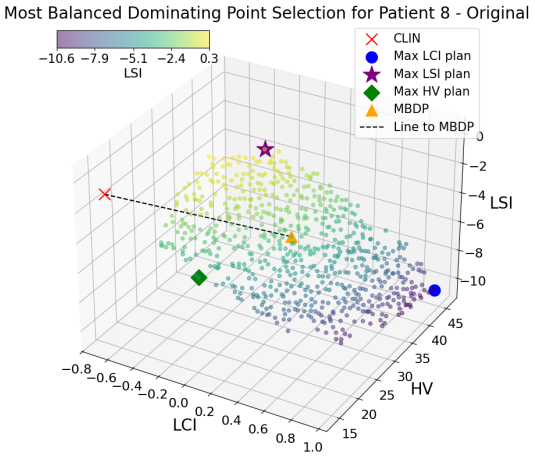

(g) Original 3D plot for Patient 8.

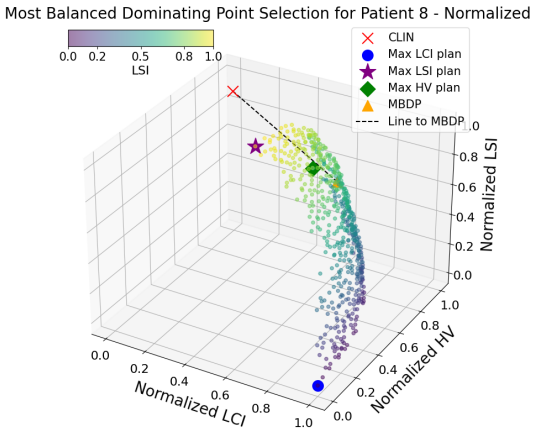

(h) Normalized 3D plot for Patient 8.

Figure A.5: Plan selection results for Patients 7,8, showing the comparison of 2D and 3D plots highlighting the most balanced points.

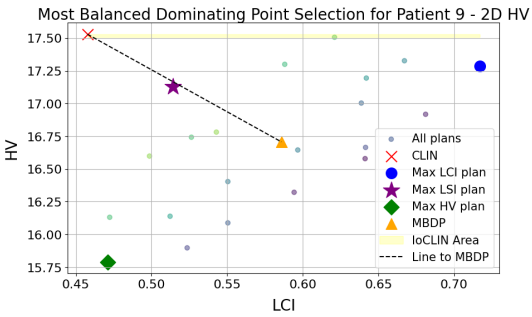

(a) 2D HV plot for Patient 9.

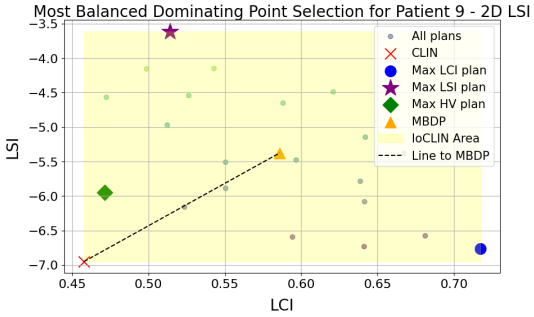

(b) 2D LSI plot for Patient 9.

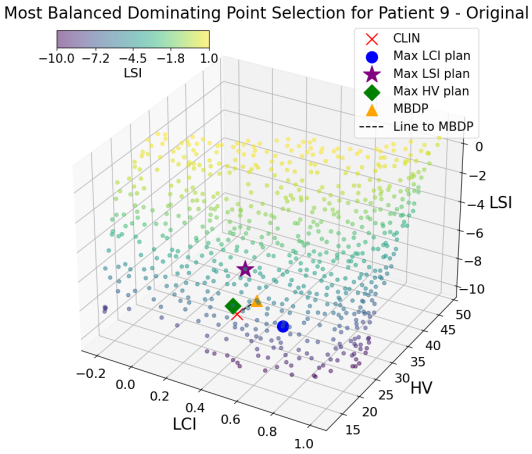

(c) Original 3D plot for Patient 9.

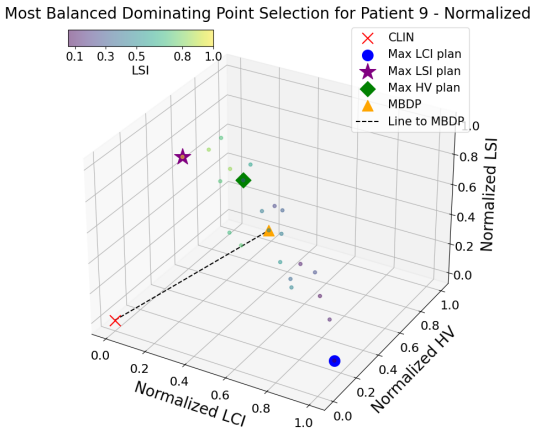

(d) Normalized 3D plot for Patient 9.

Figure A.6: Plan selection results for Patients 9, showing the comparison of 2D and 3D plots highlighting the most balanced points.

B

Contiguous volumes parameters experiments

B.1. 2-Objective Optimization: HV + LCI

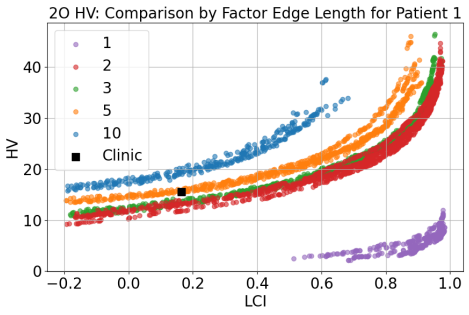

(a) Comparison by factor edge length for **Patient 1** (HV + LCI).

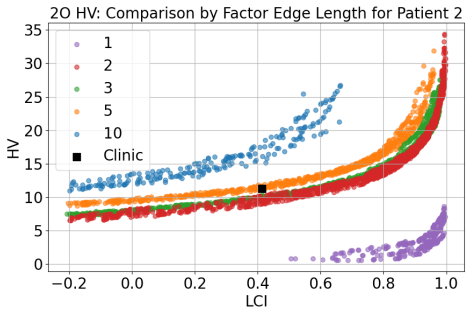

(b) Comparison by factor edge length for **Patient 2** (HV + LCI).

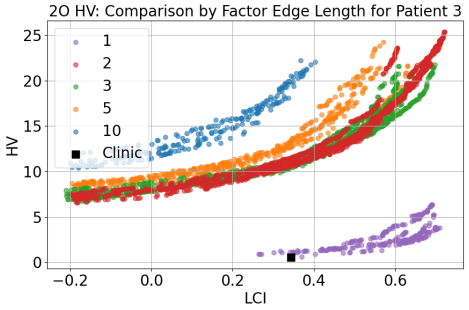

(c) Comparison by factor edge length for **Patient 3** (HV + LCI).

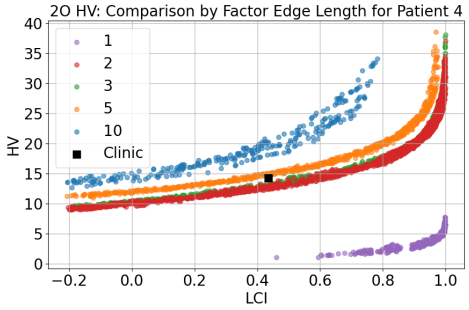

(d) Comparison by factor edge length for **Patient 4** (HV + LCI).

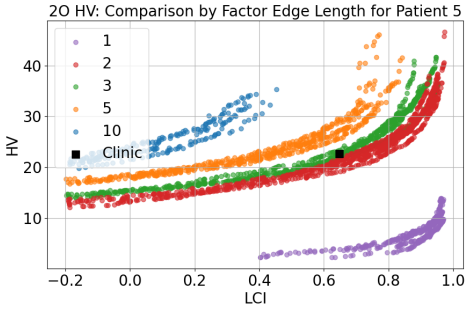

(e) Comparison by factor edge length for **Patient 5** (HV + LCI).

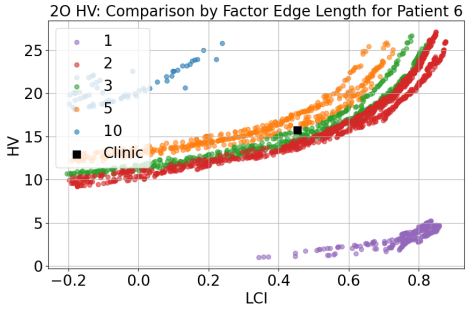

(f) Comparison by factor edge length for **Patient 6** (HV + LCI).

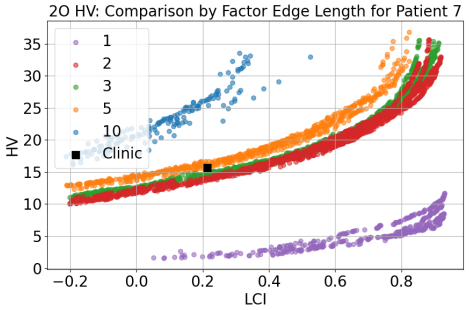

(g) Comparison by factor edge length for **Patient 7** (HV + LCI).

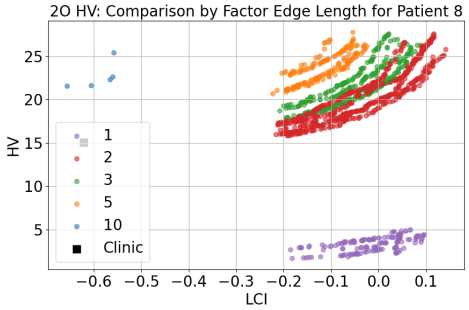

(h) Comparison by factor edge length for **Patient 8** (HV + LCI).

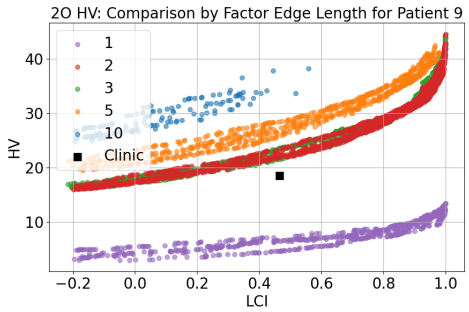

(i) Comparison by factor edge length for **Patient 9** (HV + LCI).

Figure B.1: Comparison of factor edge length for different patients based on the 2-objective optimization (HV + LCI) with varying factor edge lengths. Each setting has been run ten times to observe variance across the front. Each plot includes the clinical plan as a reference point.

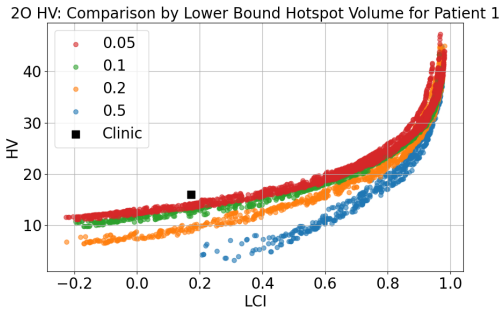

(a) comparison by lb hotspot volume for **Patient 1** (HV + LCI).

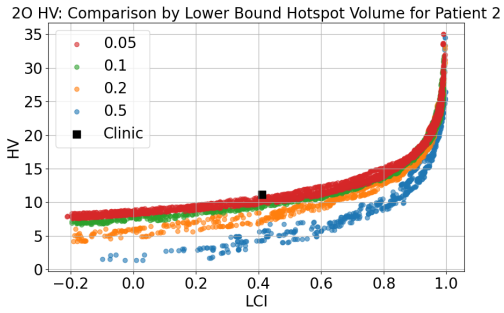

(b) comparison by lb hotspot volume for **Patient 2** (HV + LCI).

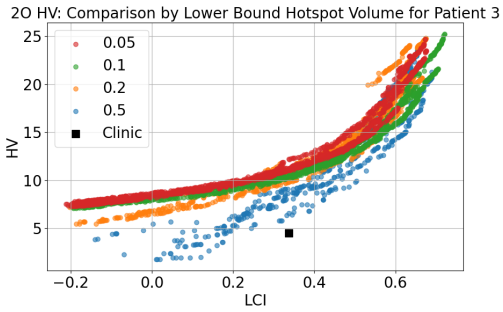

(c) comparison by lb hotspot volume for **Patient 3** (HV + LCI).

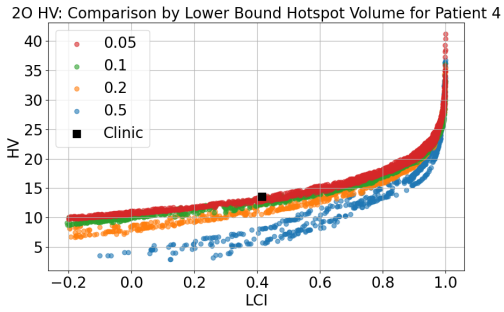

(d) comparison by lb hotspot volume for **Patient 4** (HV + LCI).

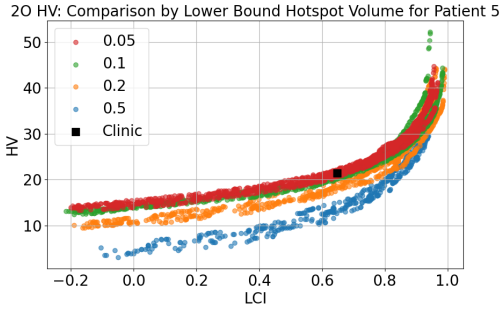

(e) comparison by lb hotspot volume for **Patient 5** (HV + LCI).

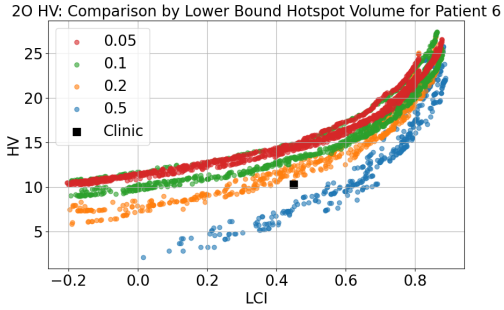

(f) comparison by lb hotspot volume for **Patient 6** (HV + LCI).

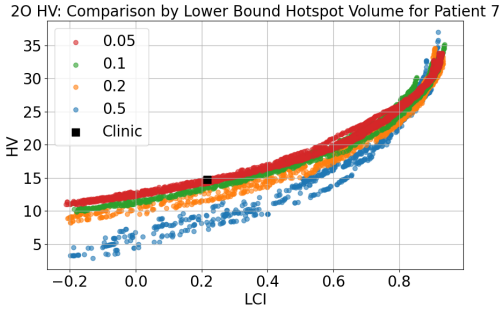

(g) comparison by lb hotspot volume for **Patient 7** (HV + LCI).

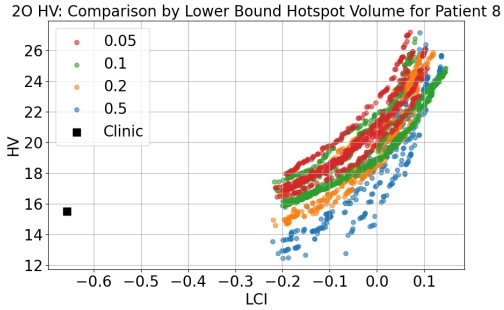

(h) comparison by lb hotspot volume for **Patient 8** (HV + LCI).

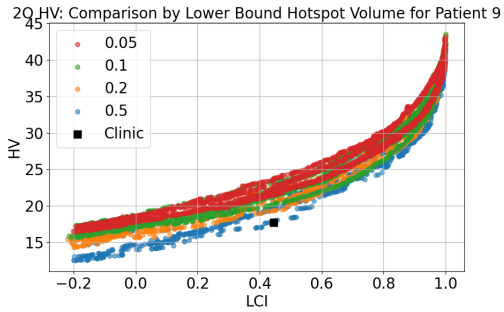

(i) comparison by lb hotspot volume for **Patient 9** (HV + LCI).

Figure B.2: Comparison of lower bound hotspot volume for different patients based on the 2-objective optimization (HV + LCI) with varying lower bound hotspot volumes. Each setting has been run ten times to observe variance across the front. Each plot includes the clinical plan as a reference point.

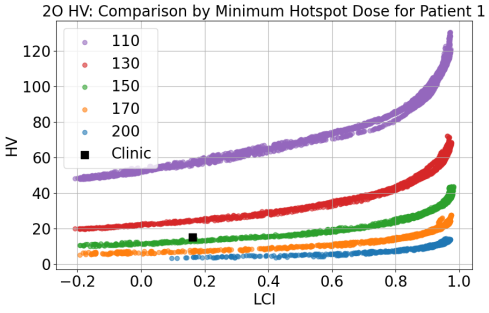

(a) Comparison by min hotspot dose for **Patient 1** (HV + LCI).

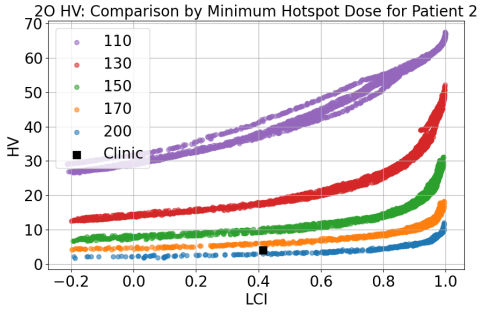

(b) Comparison by min hotspot dose for **Patient 2** (HV + LCI).

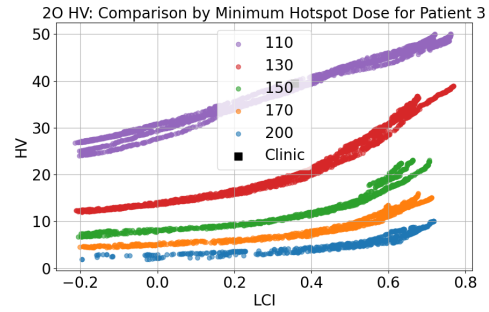

(c) Comparison by min hotspot dose for **Patient 3** (HV + LCI).

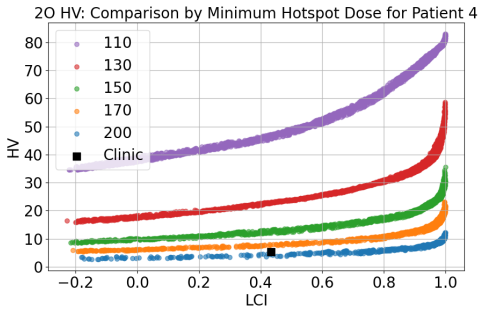

(d) Comparison by min hotspot dose for **Patient 4** (HV + LCI).

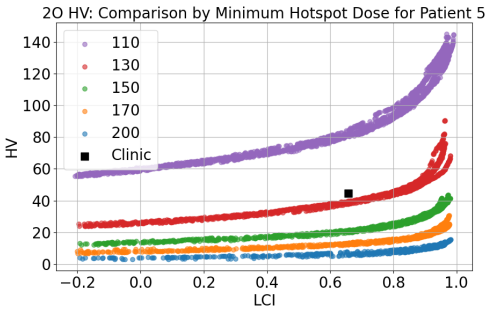

(e) Comparison by min hotspot dose for **Patient 5** (HV + LCI).

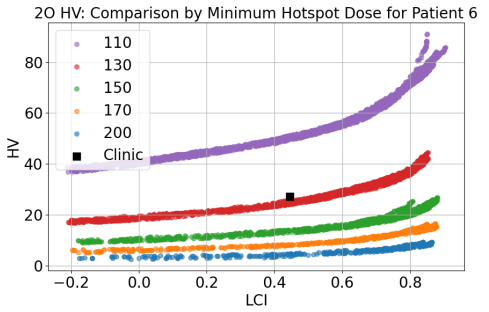

(f) Comparison by min hotspot dose for **Patient 6** (HV + LCI).

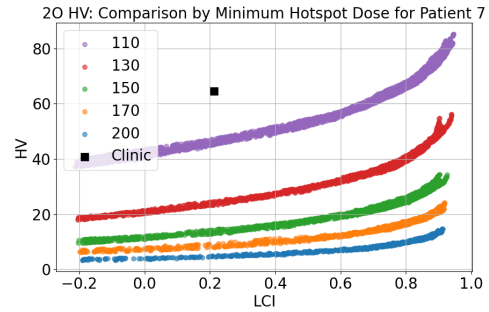

(g) Comparison by min hotspot dose for **Patient 7** (HV + LCI).

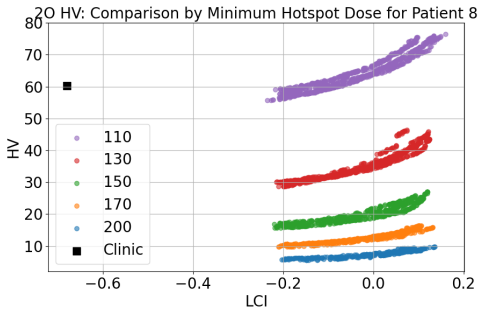

(h) Comparison by min hotspot dose for **Patient 8** (HV + LCI).

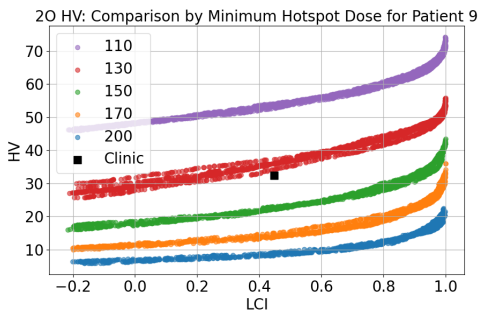

(i) Comparison by min hotspot dose for **Patient 9** (HV + LCI).

Figure B.3: Comparison of minimum hotspot dose for different patients based on the 2-objective optimization (HV + LCI) with varying minimum hotspot doses. Each setting has been run ten times to observe variance across the front. Each plot includes the clinical plan as a reference point.

B.2. 3-Objective Optimization: HV + LCI + LSI

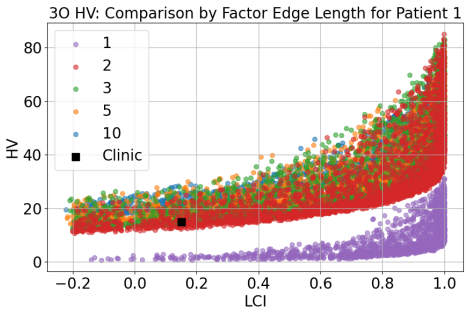

(a) Comparison by factor edge length for **Patient 1** (HV + LCI).

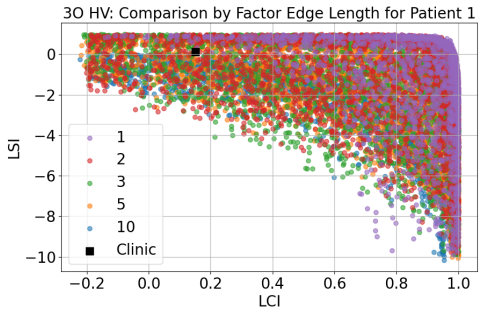

(b) Comparison by factor edge length for **Patient 1** (LSI + LCI).

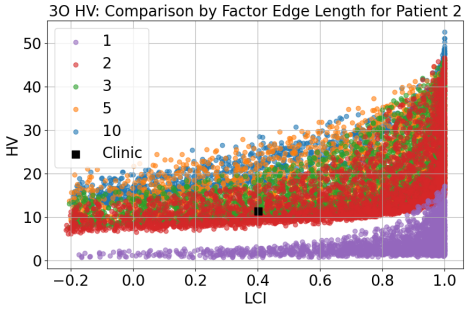

(c) Comparison by factor edge length for **Patient 2** (HV + LCI).

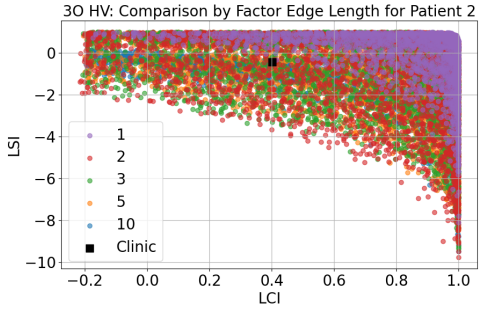

(d) Comparison by factor edge length for **Patient 2** (LSI + LCI).

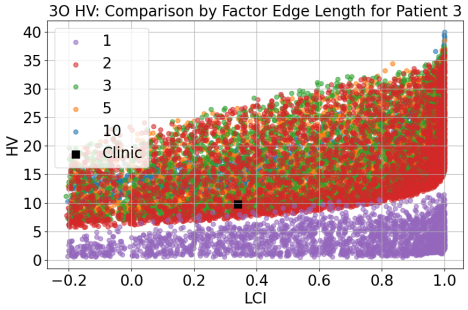

(e) Comparison by factor edge length for **Patient 3** (HV + LCI).

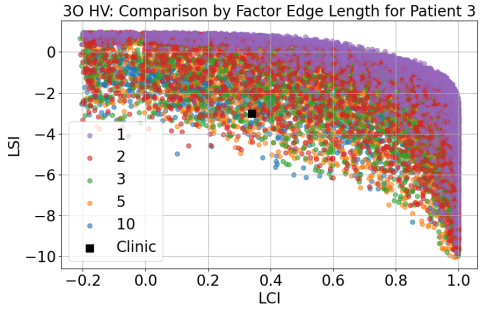

(f) Comparison by factor edge length for **Patient 3** (LSI + LCI).

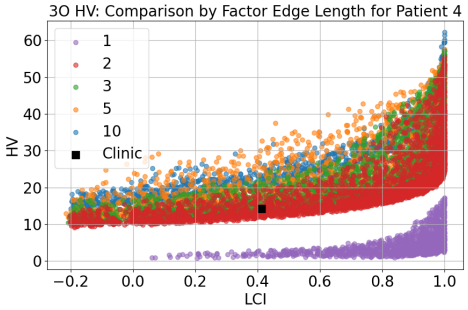

(g) Comparison by factor edge length for **Patient 4** (HV + LCI).

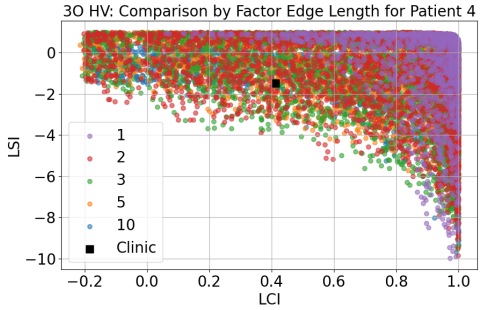

(h) Comparison by factor edge length for **Patient 4** (LSI + LCI).

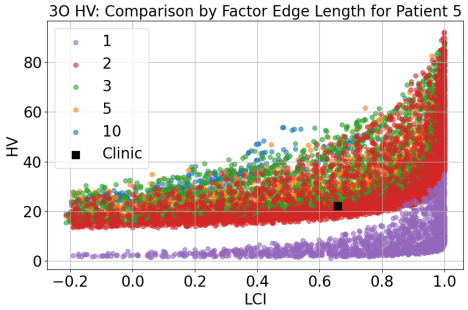

(i) Comparison by factor edge length for **Patient 5** (HV + LCI).

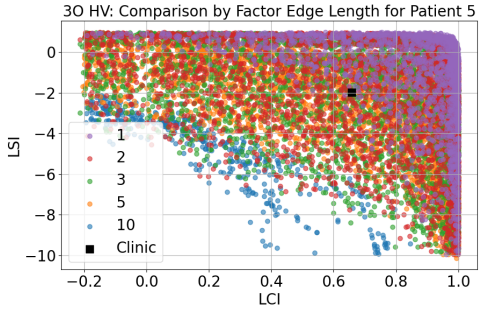

(j) Comparison by factor edge length for **Patient 5** (LSI + LCI).

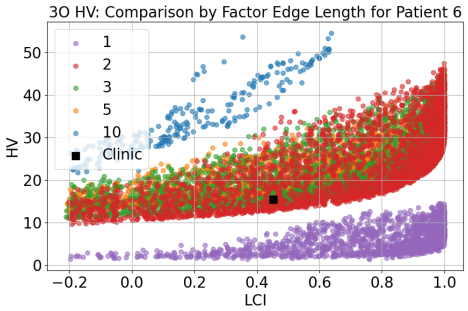

(k) Comparison by factor edge length for **Patient 6** (HV + LCI).

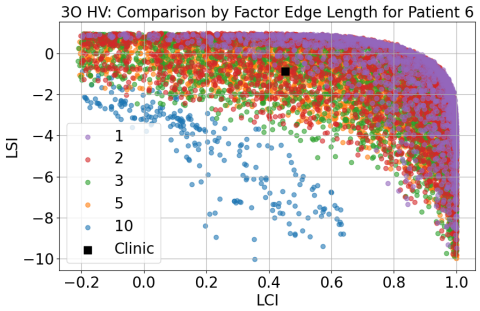

(l) Comparison by factor edge length for **Patient 6** (LSI + LCI).

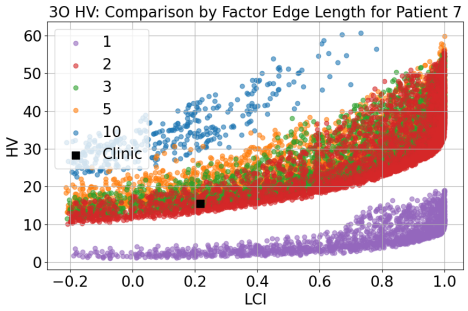

(m) Comparison by factor edge length for **Patient 7** (HV + LCI).

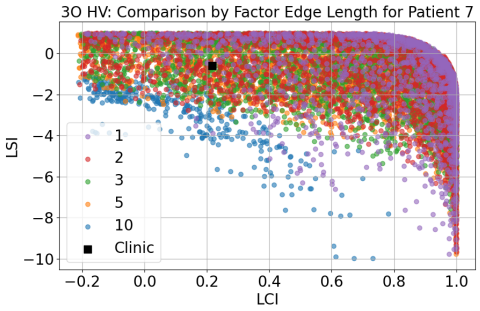

(n) Comparison by factor edge length for **Patient 7** (LSI + LCI).

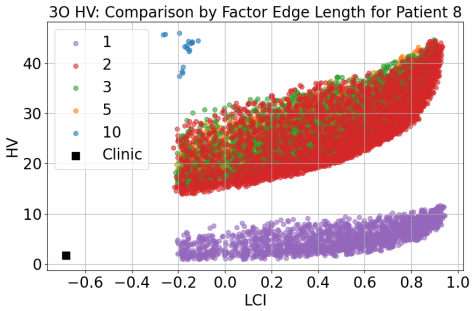

(o) Comparison by factor edge length for **Patient 8** (HV + LCI).

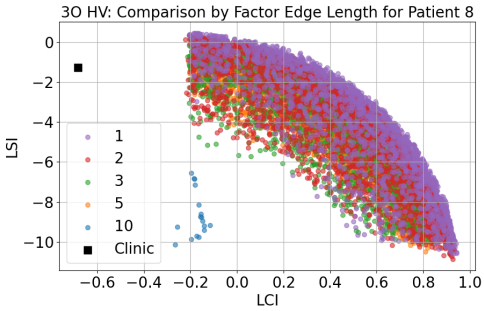

(p) Comparison by factor edge length for **Patient 8** (LSI + LCI).

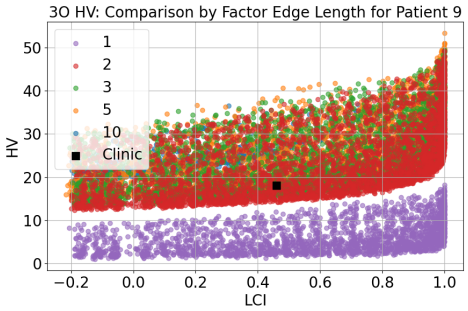

(q) Comparison by factor edge length for **Patient 9** (HV + LCI).

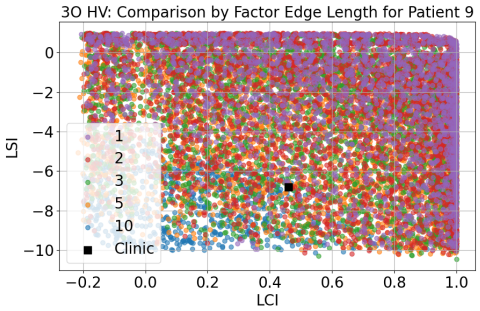

(r) Comparison by factor edge length for **Patient 9** (LSI + LCI).

Figure B.4: Approximation fronts for plans generated by BRIGHT for breast. We vary the factor edge length to understand its impact on HV.

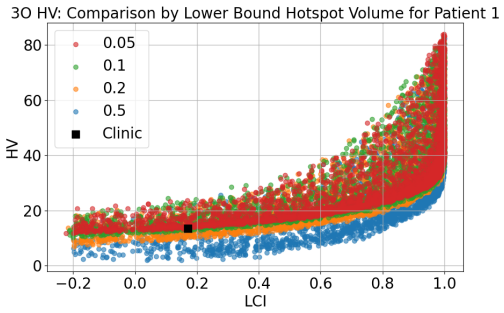

(a) comparison by lb hotspot volume for **Patient 1** (HV + LCI).

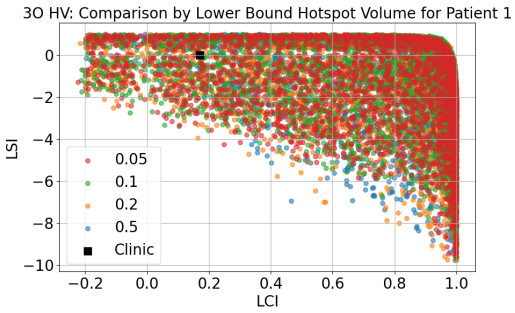

(b) comparison by lb hotspot volume for **Patient 1** (LSI + LCI).

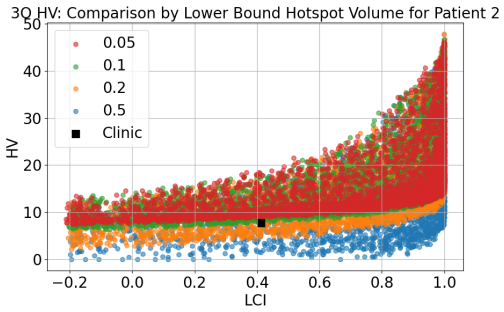

(c) comparison by lb hotspot volume for **Patient 2** (HV + LCI).

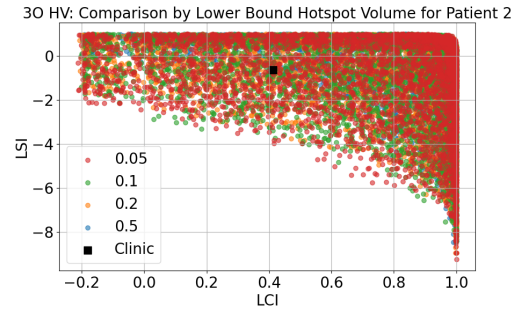

(d) comparison by lb hotspot volume for **Patient 2** (LSI + LCI).

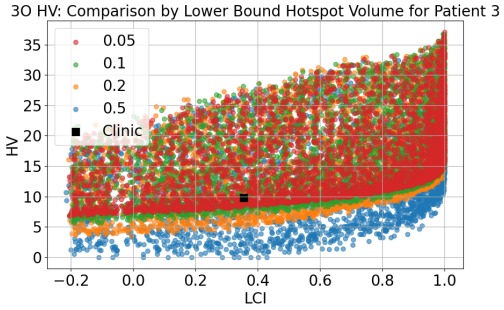

(e) comparison by lb hotspot volume for **Patient 3** (HV + LCI).

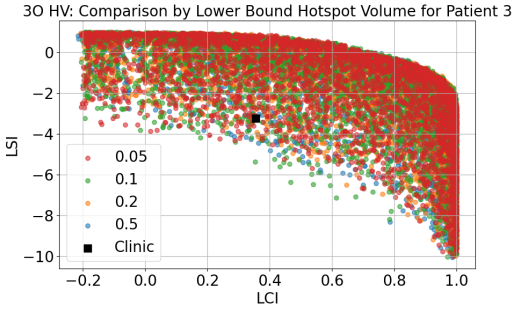

(f) comparison by lb hotspot volume for **Patient 3** (LSI + LCI).

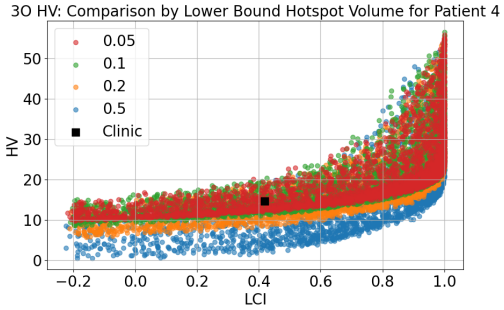

(g) comparison by lb hotspot volume for **Patient 4** (HV + LCI).

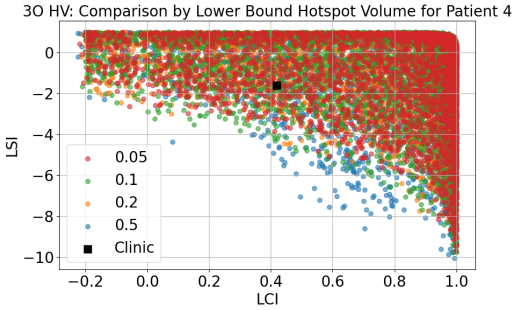

(h) comparison by lb hotspot volume for **Patient 4** (LSI + LCI).

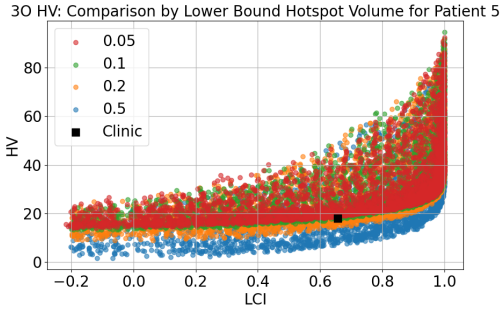

(i) comparison by lb hotspot volume for **Patient 5** (HV + LCI).

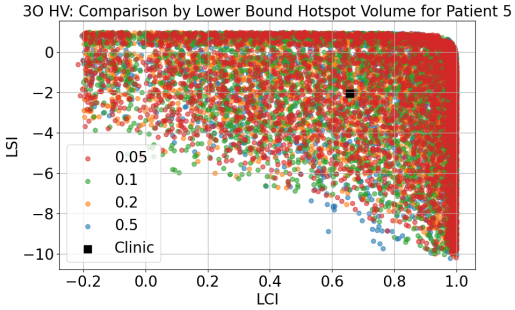

(j) comparison by lb hotspot volume for **Patient 5** (LSI + LCI).

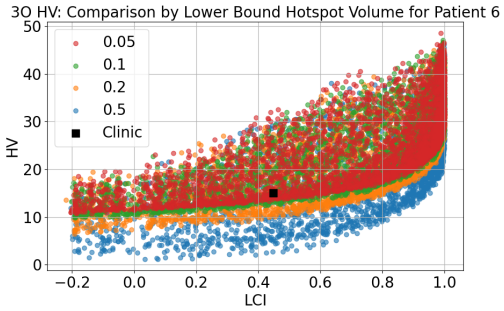

(k) comparison by lb hotspot volume for **Patient 6** (HV + LCI).

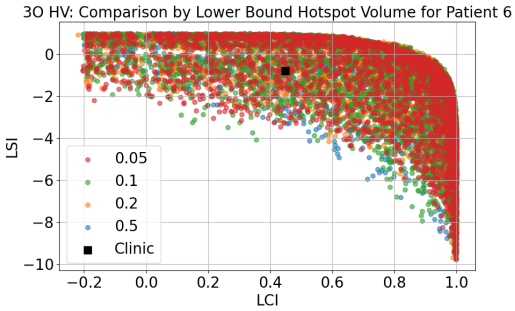

(l) comparison by lb hotspot volume for **Patient 6** (LSI + LCI).

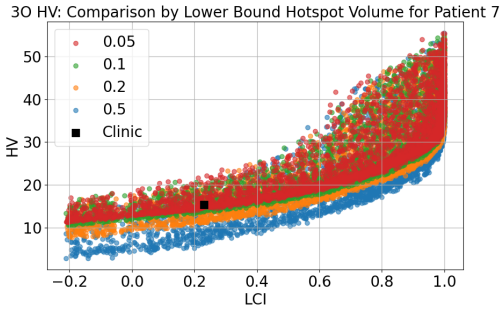

(m) comparison by lb hotspot volume for **Patient 7** (HV + LCI).

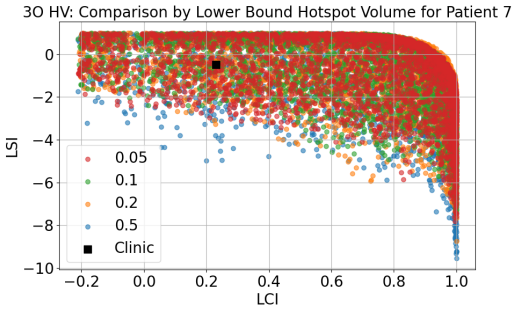

(n) comparison by lb hotspot volume for **Patient 7** (LSI + LCI).

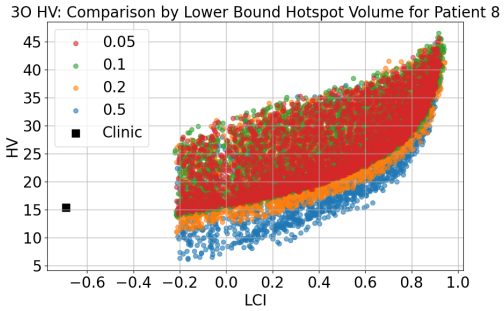

(o) comparison by lb hotspot volume for **Patient 8** (HV + LCI).

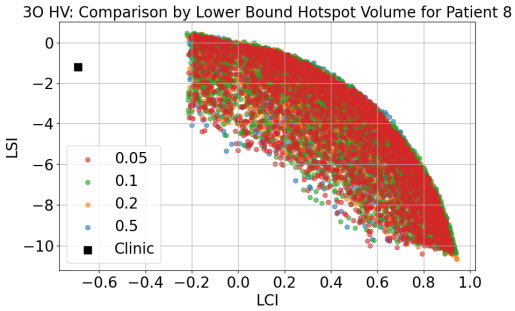

(p) comparison by lb hotspot volume for **Patient 8** (LSI + LCI).

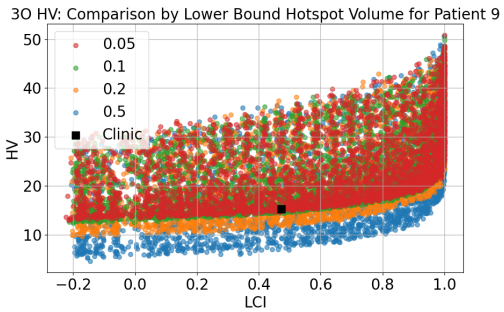

(q) comparison by lb hotspot volume for **Patient 9** (HV + LCI).

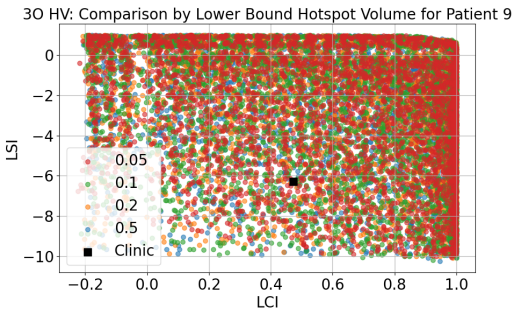

(r) comparison by lb hotspot volume for **Patient 9** (LSI + LCI).

Figure B.5: Approximation fronts for plans generated by BRIGHT for breast. We vary the lower bound hotspot volume to understand its impact on HV.

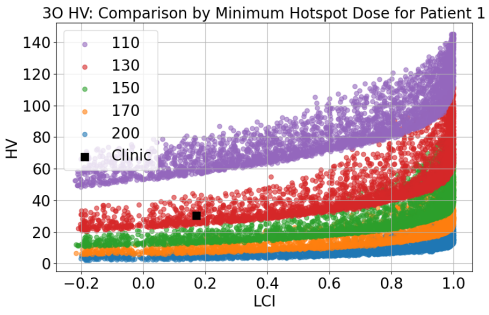

(a) Comparison by min hotspot dose for **Patient 1** (HV + LCI).

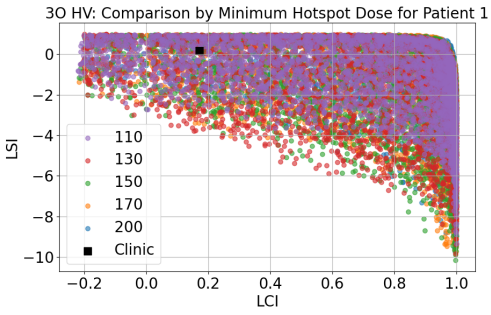

(b) Comparison by min hotspot dose for **Patient 1** (LSI + LCI).

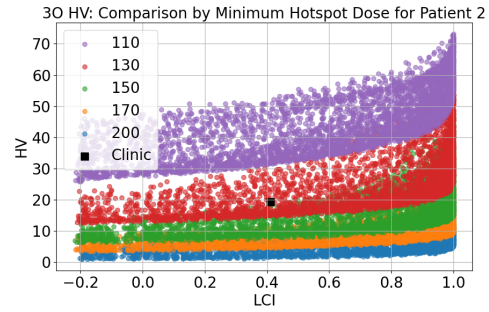

(c) Comparison by min hotspot dose for **Patient 2** (HV + LCI).

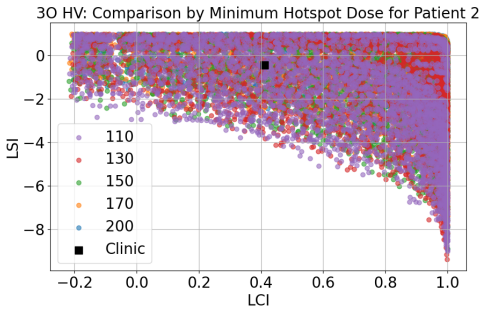

(d) Comparison by min hotspot dose for **Patient 2** (LSI + LCI).

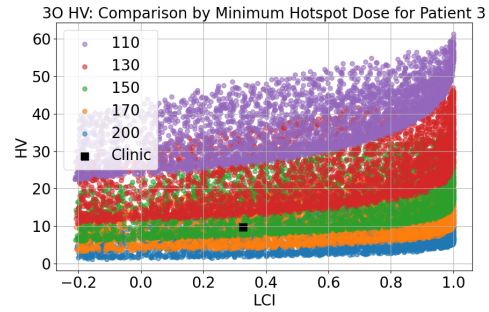

(e) Comparison by min hotspot dose for **Patient 3** (HV + LCI).

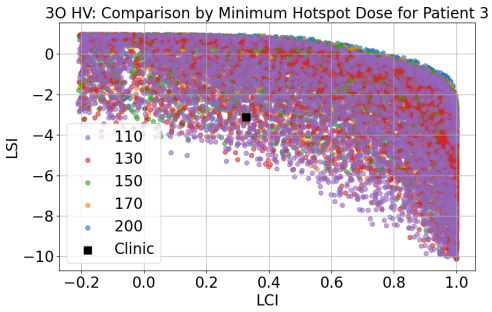

(f) Comparison by min hotspot dose for **Patient 3** (LSI + LCI).

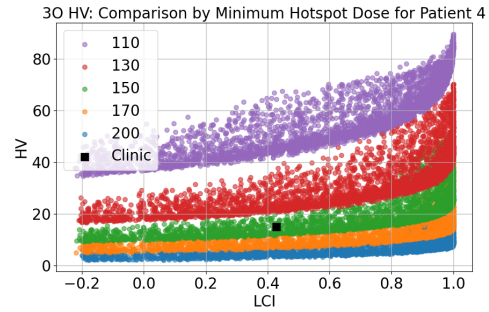

(g) Comparison by min hotspot dose for **Patient 4** (HV + LCI).

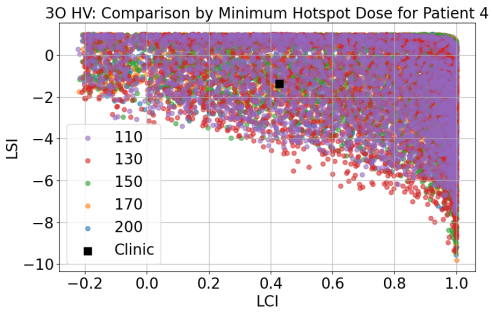

(h) Comparison by min hotspot dose for **Patient 4** (LSI + LCI).

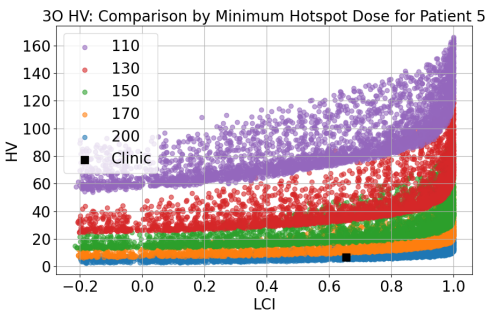

(i) Comparison by min hotspot dose for **Patient 5** (HV + LCI).

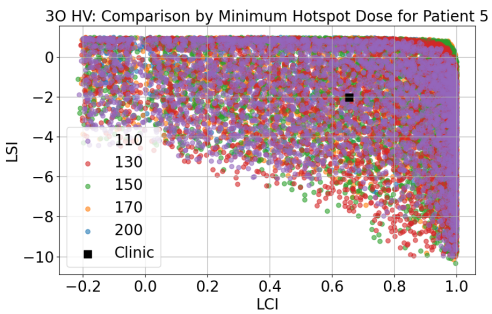

(j) Comparison by min hotspot dose for **Patient 5** (LSI + LCI).

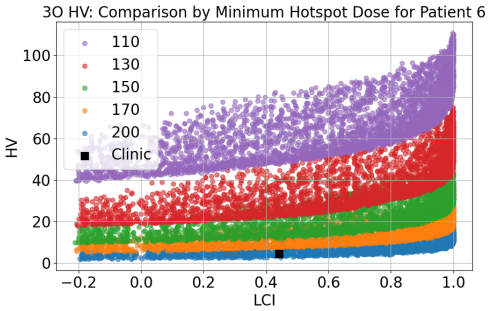

(k) Comparison by min hotspot dose for **Patient 6** (HV + LCI).

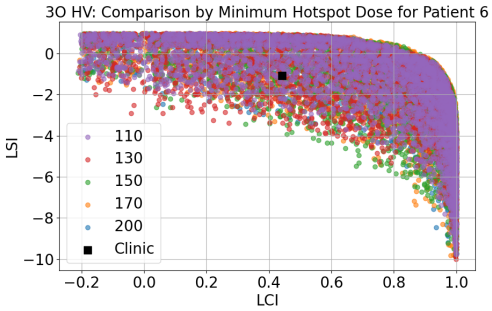

(l) Comparison by min hotspot dose for **Patient 6** (LSI + LCI).

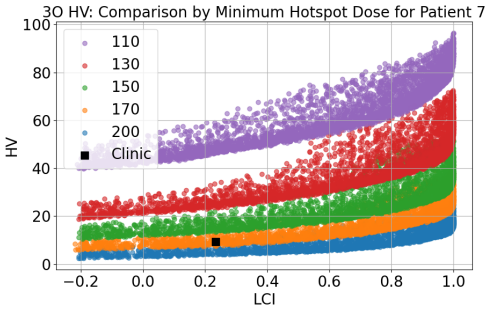

(m) Comparison by min hotspot dose for **Patient 7** (HV + LCI).

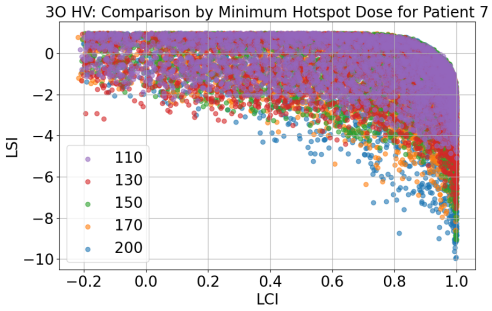

(n) Comparison by min hotspot dose for **Patient 7** (LSI + LCI).

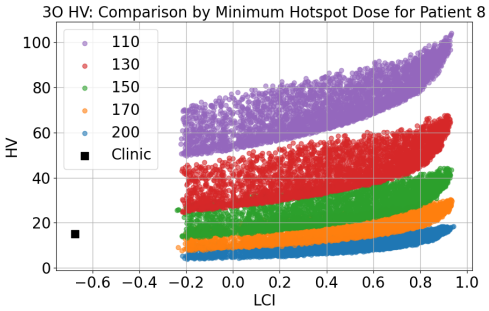

(o) Comparison by min hotspot dose for **Patient 8** (HV + LCI).

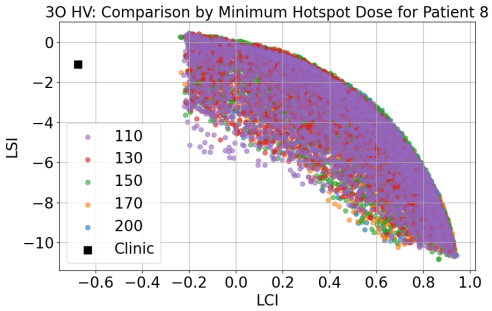

(p) Comparison by min hotspot dose for **Patient 8** (LSI + LCI).

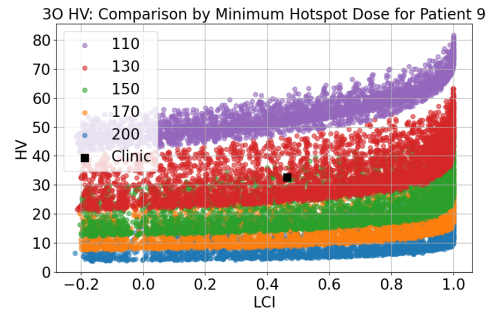

(q) Comparison by min hotspot dose for **Patient 9** (HV + LCI).

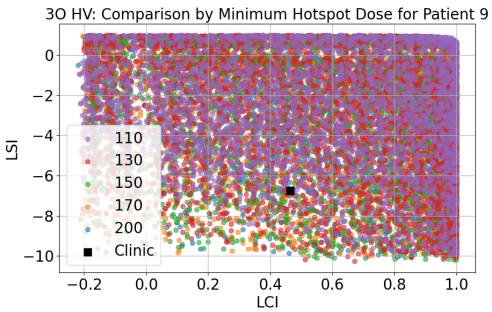

(r) Comparison by min hotspot dose for **Patient 9** (LSI + LCI).

Figure B.6: Approximation fronts for plans generated by BRIGHT for breast. We vary the minimum hotspot dose to understand its impact on HV.