Data-Driven Brachytherapy Protocol Generation for BRIGHT using RV-GOMEA Based Optimization

Author:
Miquel Ricart I Oltra
Student number 5240344

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

Thesis Committee:

Chair:	Prof. dr. P.A.N. Bosman, Faculty EEMCS, TU Delft
University supervisors:	Prof. dr. T. Alderliesten, LUMC
	MSc. R. J. Scholman, CWI
External supervisor:	Prof. dr. A.D. Todor, VCU
Committee Member:	Prof. dr. R. Marroquim, Faculty EEMCS, TU Delft

TU Delft Delft University of Technology

CWI Centrum Wiskunde & Informatica

Preface

This thesis marks the conclusion of my Master's degree in Embedded Systems at Delft University of Technology. Reflecting on my academic journey, it has been a period of significant personal and professional evolution. I have embraced the mistakes made along the way as essential learning opportunities, which have not only deepened my technical knowledge in software engineering but also contributed to my growth as an individual.

This research allowed me to dive deep into the field of evolutionary algorithms, applying them to a real-world challenge. This research was conducted within the Evolutionary Intelligence group at the Centrum Wiskunde & Informatica (CWI). The journey toward completing this work has been as challenging as it was rewarding, and it would not have been possible without the support of several key people.

First and foremost, I would like to thank MSc. Renzo Scholman for his insightful feedback and constant guidance. I am also deeply grateful to Prof. dr. Peter Bosman and Prof. dr. Tanja Alderliesten for their invaluable supervision and expertise throughout the past year.

Finally, I want to express my gratitude to my friends and family for their endless patience and encouragement during these months of research. A very special thanks goes to my partner, whose support and motivation kept me focused and driven to reach the finish line.

Miquel Ricart I Oltra
Delft, February 2026

Abstract

High-Dose-Rate brachytherapy is a critical component in the treatment of locally advanced cervical cancer. While automated treatment planning systems, such as BRIGHT, have demonstrated the ability to generate high-quality plans, their clinical adoption is hindered by the complexity of their configuration. Deploying such a system in a new hospital requires the manual definition of a clinical protocol that accurately reflects the local institution's specific standard of care. This "cold start" problem is time-consuming for both doctors and researchers.

This thesis proposes a novel framework for the Automated Discovery of Clinical Protocols. By formulating the protocol configuration as a bi-level optimization problem, we employ the Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm to autonomously extract implicit expert knowledge from a repository of historical clinical plans. The system evolves a set of protocol parameters that, when fed into BRIGHT, reproduce radiation dose distributions as preferred by human experts.

We validate this approach using anonymized patient data from Virginia Commonwealth University. Through a series of experiments with incrementally increasing complexity, ranging from optimizing simple dose thresholds to evolving the definitions of dosimetric metrics, we demonstrate that the proposed framework can successfully identify protocols that generate treatment plans that are quantitatively similar to the clinical ground truth. This research serves as a proof-of-concept, offering a pathway to rapidly deploy automated planning systems while ensuring alignment with local clinical expertise.

Contents

List of Figures

List of Abbreviations

Abbreviation	Definition
AFS	Average Front Similarity
AMS	Anticipated Mean Shift
AVS	Adaptive Variance Scaling
BRIGHT	BRachytherapy via artificially Intelligent GOMEA-Heuristic based Treatment planning
CPS	Closest Plan Similarity
CTV_{HR}	High-Risk Clinical Target Volume
CTV_{IR}	Intermediate-Risk Clinical Target Volume
DCP	Dose-Calculation Point
DVH	Dose-Volume Histogram
DVI	Dose-Volume Index
EA	Evolutionary Algorithm
EBRT	External Beam Radiation Therapy
EQD2	Equivalent Dose in 2 Gy Fraction
FOS	Family of Subsets
GBO	Gray-Box Optimization
GOM	Gene-pool Optimal Mixing
GOMEA	Gene-pool Optimal Mixing Evolutionary Algorithm
GP	Genetic Programming
GTV_{res}	Residual Gross Tumor Volume
HDR	High-Dose-Rate
KBP	Knowledge-Based Planning
MSE	Mean Squared Error
NT	Normal Tissue
OAR	Organ at Risk
PTV	Planning Target Volume
TPS	Treatment Planning System/s
ROI	Region Of Interest

Chapter 1

Introduction

Cancer remains one of the most fatal diseases of recent times, causing many deaths every year. Among diseases affecting women, cervical cancer is particularly prominent. According to the International Agency for Research on Cancer (IARC), it ranks as the fourth most common cancer in women globally, with an estimated 660,000 new cases and 350,000 deaths in 2022 [1]. Moreover, it is estimated that a million children became maternal orphans in 2020 because their mother died from cancer that year (15% of prevalent maternal orphans)[2], with close to one half of these orphans the result of maternal deaths from either female breast or cervical cancer.

This research focuses on the treatment of this second type, cervical cancer. The standard of care for locally advanced cervical cancer involves a combination of External Beam Radiation Therapy (EBRT), chemotherapy, and High-Dose-Rate (HDR) Brachytherapy [3, 4]. While EBRT targets the tumor and elective nodal volumes from outside the body, HDR brachytherapy allows for a highly localized "boost" of radiation. By placing a radioactive source (typically Iridium-192) directly into or adjacent to the tumor using specialized applicators, clinicians can deliver an aggressive dose to the targeted cancerous region while achieving a rapid dose fall-off that minimizes the radiation exposure to the surrounding tissue and healthy organs [5].

However, realizing the full potential of this high-dose precision requires exact planning. As clinical understanding deepens and imaging evolves from 2D X-rays to 3D MRI/CT-guided techniques [6, 7], the definition of an optimal treatment has become increasingly stringent. This higher level of insight demands methodologies that are not only more precise but also consistent in adhering to complex dose restrictions, a task that pushes the limits of manual optimization. Furthermore, the speed of the planning process is critical. Because internal organs are dynamic, significant movement can occur between the moment of image acquisition and the delivery of radiation. The longer the planning phase lasts, the greater the anatomical discrepancy caused by organ motion, which inevitably introduces error into the treatment [8].

To address these dual challenges of required precision and temporal urgency, the field of medical physics has turned to automated treatment planning. In recent years, a new tool, BRIGHT [9], has gained attention for its success in leveraging evolutionary algorithms to explore the multi-objective search space of treatment plans.

However, despite these technological advancements, the transition from manual to automated planning introduces a new set of challenges that slow its clinical adoption.

1.1 Problem Statement

While automated treatment planning systems (TPS) like BRIGHT have demonstrated the capability to generate mathematically superior plans, their clinical deployment faces a significant translation barrier: the system configuration. Integrating these solutions into a new clinical environment is not immediate, and it requires a lengthy onboarding phase to align the algorithm with the specific planning philosophy of the hospital.

This extensive setup is necessary because optimization algorithms do not possess clinical intuition, and they strictly minimize a mathematical objective function. Simply relying on international guidelines, such as EMBRACE-II [10], offers a starting point but are insufficient for direct algorithmic configuration for two reasons:

1. Clinical practice involves complex, often unwritten trade-offs between target coverage and organ sparing that vary significantly between institutions and individual experts [11, 12].

2. Translating a physician's subjective preferences into the numerical constraints required by a TPS is a complex and unintuitive task. Currently, this requires a group of doctors and researchers to perform a manual "trial-and-error" tuning process (adjusting protocol parameters, re-optimizing, and reviewing results), which can take weeks or months of discussion to set a satisfactory baseline for a new hospital [13].

To overcome these challenges, the field of radiotherapy has explored Knowledge-Based Planning (KBP) [14, 15, 16]. KBP typically uses statistical models or machine learning on prior plans to predict dose distributions (e.g., Dose-Volume Histograms) for new patients. These predictions guide the optimization process. While effective, standard KBP often focuses on predicting outcomes (the dose) rather than optimizing the instructions (protocol parameters) that drive planning.

This thesis proposes a novel evolutionary alternative. Instead of using predictive models or large training databases to estimate dose, we introduce an Automated Protocol Discovery framework. We use the Gene-pool Optimal Mixing Evolutionary Algorithm (GOMEA) to treat protocol configuration as an optimization problem. This approach 'reverse-engineers' explicit clinical objectives from historical expert data. It allows us to autonomously 'learn' the local standard of care to be used by BRIGHT without manual tuning. By automating this calibration, the framework has the potential to offload weeks of manual work of professionals to autonomous computation, thereby streamlining the onboarding of new institutions.

1.1.1 The Bi-Level Planning Problem

The generation of a brachytherapy treatment plan can be viewed as a complex optimization task governed by a set of instructions known as a Clinical Protocol.

The Clinical Protocol and DVIs:

The protocol is essentially a collection of constraints that the optimization algorithm must satisfy. In clinical practice, these are quantified using Dose-Volume Indices (DVIs). A DVI is a metric that relates a specific radiation dose to a volume of tissue or vice versa. For example, a common coverage aim is $D_{90\%} > 7.0\ Gy$, which requires that 90% of the tumor volume receives at least 7.0 Gray of radiation. Conversely, a sparing aim such as $V_{65\%} < 2.0\ cm^3$ restricts the volume of a healthy organ receiving high radiation.

The Optimization Engine, BRIGHT and GOMEA:

This research utilizes BRIGHT (BRachytherapy via artificially Intelligent GOMEA-Heuristic based Treatment planning), a treatment planning system that employs the GOMEA algorithm. GOMEA is a model-based evolutionary algorithm that excels at solving optimization problems where variables are highly interdependent.

The Bi-Level Optimization Challenge:

The core problem addressed in this research is that BRIGHT requires a strictly defined protocol (inputs) to generate a high-quality Pareto front of treatment plans (outputs). However, determining the optimal protocol parameters that result in a clinically acceptable plan is not straightforward.

Therefore, we formulate the selection of protocol parameters as a bi-level optimization problem:

- The Inner Loop (BRIGHT): Optimizes the dwell times of the radiation source to maximize plan quality based on a specific protocol.
- The Outer Loop (Proposed Framework): Optimizes the protocol parameters themselves (the DVIs and their thresholds) to minimize the difference between the generated plans and historical expert plans.

This hierarchical structure allows us to ”reverse-engineer” the implicit knowledge of clinical experts, effectively automating the configuration of the planning system.

1.2 Research Questions

The goal of this thesis is to investigate whether the 'art' of treatment planning can be formalized as a learnable parameter space that encodes clinical intent. Therefore, the Main Research Question is the following:

MRQ: *Can the implicit expert knowledge within accepted clinical plans be extracted via evolutionary optimization to reconstruct valid institution-specific protocols automatically?*

To answer this, we decompose the problem into four specific experiments, corresponding to the incremental complexity of the search space and the computational resources required.

First, let us investigate the problem to understand how the bi-level optimization behaves. We know that using a stochastic algorithm like an EA to evaluate the fitness of our candidate solutions will cause the search space of the problem to be noisy, where equal evaluations of the same protocol may yield different results. To determine if our framework is a logical approach to this problem, we must understand the topology of the search space. Thus, the first research question explores the search space and how modifying the evaluation of solutions might influence the speed of convergence:

RQ1: *How do fitness landscape smoothing techniques (such as stochastic averaging) and objective function shaping (such as gradient steepening) affect the search space and the convergence velocity and robustness of the evolutionary search?*

Second, we must establish the feasibility of the approach. Before attempting to learn complex metric definitions, we must verify if the evolutionary algorithm can optimize simple numerical thresholds (the aims) to reproduce expert results within a fixed protocol structure. This leads to the second research question, which investigates both the capability of the search algorithm to minimize geometric error and whether the similarity metric can lead to clinical acceptability:

RQ2: *Can an evolutionary search strategy identify protocol aims that minimize the difference between generated and clinically approved treatment plans, and does this convergence translate to clinically valid treatment plans?*

Third, we address the limitations of rigid protocol structures. The previous optimization only tunes the aim values, but expert protocols might benefit from varying what they optimize for (the metric of the DVI). The third research question investigates the impact of expanding the search space to include the evolution of the metric definitions themselves:

RQ3: *How does increasing the complexity of the optimization, by introducing the evolution of complete DVIs, affect the convergence behavior of the evolutionary algorithm and the quality of the resulting solutions?*

Finally, we address the issue of scalability. A fully automated discovery system must evolve a complete protocol, which means including at least one constraint per region of interest. RQ4 investigates the system's performance in this high-dimensional setting, specifically analyzing how additional resources can be leveraged to overcome the computation-versus-precision trade-off inherent of this stochastic optimization:

RQ4: *Does the framework scale effectively to high-dimensional, unconstrained protocol discovery, and to what extent does massive parallelization mitigate the stochastic noise when evaluating more complex genotypes?*

1.3 Outline of the Thesis

The remainder of this thesis is structured as follows. Chapter 2 presents the necessary medical and algorithmic context. It details the principles of HDR Brachytherapy, the RV-GOMEA framework, and BRIGHT. Chapter 3 defines the methodology for automated protocol discovery. It details the optimization problem and the functions used to quantify plan similarity. Chapter 4 describes the experimental design used to validate the framework. Chapter 5 presents the experimental results. Chapter 6 discusses the limitations and proposes directions for future research. Chapter 7 concludes the thesis and recapitulates the contributions.

Chapter 2

Background and Related Work

This chapter establishes the theoretical and clinical foundations necessary to understand the protocol discovery framework proposed in this thesis. We first provide an overview of cervical cancer treatment and the specifics of HDR brachytherapy, focusing on the definition of clinical protocols and plan evaluation. Subsequently, we introduce Evolutionary Algorithms, specifically detailing (MO)RV-GOMEA and its implementation within the BRIGHT treatment planning system.

2.1 HDR Brachytherapy Treatment

High-dose-rate brachytherapy refers to a type of brachytherapy in which a highly radioactive source is placed very close to the cancer cells for a short period of time with the objective of killing the cancerous cells. To introduce the radioactive source, a set of needles and catheters is inserted into the soft tissue of the patient with the use of an applicator. These catheters are hollow and facilitate the precise access of the radiation source to the tumour [17, 18]. The radioactive source is then moved through multiple dwell positions inside the catheters, where it is held stationary for a prescribed time, the dwell time, to apply the specified amounts of radiation to the target volumes. The treatment plan defined by the clinical experts is what determines the radiation time at each dwell position.

The goal is to deliver the prescribed dose of radiation to the target volumes, which include the entire uterine cervix and any residual tumor present at the time of treatment [19], while minimizing the dose to surrounding healthy tissues and Organs at Risk (OARs). This fundamental trade-off is the core challenge of treatment planning.

To create a patient-specific treatment plan, detailed anatomical information is acquired through medical imaging after the applicator has been inserted. While historical techniques relied on 2D X-ray imaging, modern brachytherapy is guided by 3D imaging, primarily from CT and MRI scans, which allows for a much more precise delineation of the targets and OARs. [6]

2.2 Treatment Plan and Evaluation

Physically, a brachytherapy treatment plan is defined by a set of dwell times. These dwell times are the fundamental instructions delivered to the treatment unit. Their magnitude and position determine the 3D spatial distribution of the radiation dose deposited in the patient [18].

However, determining these times is not the starting point, but the result of the planning process. In modern Inverse Planning, the workflow is reversed: the clinician defines the desired dose distribution first, and an optimization algorithm calculates the necessary dwell times to achieve it. Thus, while dwell times are the physical input for the treatment machine, they are the output of the optimization system.

To drive this optimization, the clinical intent must be formalized into a Clinical Protocol. The protocol consists of a set of objectives that distinguish a high-quality plan from an unacceptable one. Mathematically, it limits the feasible search space by balancing multiple objectives that we classify into two competing goals: maximizing the radiation dose to the tumor and minimizing the dose to the OARs, such as the bladder, rectum, and sigmoid.

Within the clinical protocol, these conflicting goals are quantified using DVIs. Each DVI serves as a structured rule for plan acceptance, defined by four components: The ROI that it is applied to, the metric, the operator, and the target goal or aim.

The metric determines the physical quantity being measured (either of dose or of volume). The operator establishes the clinical intent, distinguishing between coverage aims (which set a minimum lower bound to ensure the tumor receives enough radiation) and sparing aims (which set a maximum upper bound to protect healthy tissue). Finally, the aim defines the specific numerical threshold that the plan should satisfy.

Based on the metric used, these indices are classified into two types:

- Dosimetric indices (D_V): These limit the dose received by a specific percentage of a structure's volume. For example, a coverage aim for a volume O of $D^o_{90\%} > 95\%$ requires that 90% of the volume receives at least 95% of the prescribed dose.
- Volumetric indices (V_D): These limit the absolute volume of a structure that receives a certain dose. For example, a sparing aim for a volume O of $V^o_{60\%} < 2.0cm^3$ requires that the volume receiving 60% of the prescribed radiation dose should stay below 2 cubic centimeters.

The evolution of these protocols has been driven by decades of international research, particularly through the GEC-ESTRO GYN network [10]. This collaborative effort resulted in standardized, evidence-based guidelines for image-guided adaptive brachytherapy (IGABT), culminating in large-scale clinical studies, such as EMBRACE I and EMBRACE II. The EMBRACE II study, in particular, establishes a benchmark protocol with detailed dose prescription aims and limits for both tumor targets (GTV, CTV HR) and multiple OARs, including the bladder and rectum. The DVIs from the EMBRACE II protocol are summarized in Table 2.1.

While highly detailed protocols like EMBRACE II have significantly advanced the standard of care, manually creating a plan that satisfies their many competing objectives remains a complex process. This challenge has motivated the development of automated treatment planning systems that use artificial intelligence to navigate these trade-offs [20].

Table 2.1: Planning aims and constraints according to the EMBRACE-II protocol [10]. Note: Doses have been converted from EQD2 to fractional dose to maintain consistency with the rest of this thesis.

Target	CTV_{HR} $D_{90\%}$	CTV_{HR} $D_{98\%}$	GTV_{res} $D_{98\%}$	CTV_{IR} $D_{98\%}$	Point A
Planning Aims	$>$ 7.8 Gy $<$ 8.3 Gy	$>$ 5.8 Gy	$>$ 8.3 Gy	$>$ 3.5 Gy	$>$ 4.1 Gy
Limits	$>$ 7.1 Gy	–	$>$ 7.8 Gy	–	–
OAR	Bladder D_{2cm^3}	Rectum D_{2cm^3}	Recto-vaginal point D_{point}	Sigmoid D_{2cm^3}	Bowel D_{2cm^3}
Planning Aims	$<$ 5.5 Gy	$<$ 4.0 Gy	$<$ 4.0 Gy	$<$ 4.5 Gy	$<$ 4.5 Gy
Limits	$<$ 6.3 Gy	$<$ 6.3 Gy	$<$ 6.3 Gy	$<$ 6.3 Gy	$<$ 6.3 Gy

2.2.1 Clinical evaluation of the treatment plan

When evaluating these plans, the Dose-Volume Histogram (DVH) serves as the main tool for clinical experts to assess the quality of a treatment plan and ensure compliance with the protocol [21]. The DVH condenses the complex 3D dose distribution into a 2D cumulative radiation graph, where the x-axis represents the radiation dose (typically in *Gy* or percentage of the prescribed dose) and the y-axis represents the volume of the structure receiving at least that dose (typically in cm^3 or percentage of the volume).

By visually inspecting these curves, clinical experts can assess the trade-off between target coverage and organ sparing, allowing the professional to verify that objectives such as $D_{90\%}^{CTV_{HR}} > 95\%$ are satisfied. An ideal plan exhibits a "square" shoulder for target volumes (indicating high, uniform coverage), and a rapid fall-off near the origin for OARs to minimize toxicity. An example of DVH is illustrated in Figure 2.2.

2.3 Evolutionary Algorithms

While the DVH provides a good metric for evaluating the quality of a plan, it offers no mechanism for generating one. There is a disconnect between the clinical intent (visualized as a DVH) and the machine parameters (dwell times) required to achieve it. This is due to the physics of brachytherapy: radiation is cumulative, meaning a change in a single dwell time simultaneously affects multiple adjacent structures. This physical coupling transforms the planning process into a non-convex optimization problem with a vast search space. In this landscape, traditional gradient-based methods often get trapped in local optima.

To solve this inverse problem effectively, we turn to Evolutionary Algorithms (EAs). EAs have proven particularly robust in navigating such complex, non-linear search spaces [22]. As a class of metaheuristics grounded in evolutionary computation, EAs mimic biological processes, such as selection, crossover, and mutation, to iteratively evolve high-quality solutions [23]. The algorithm operates not on a single solution, but on a population of individuals, where each individual can, for example, represent a unique treatment plan defined by its specific genotype (the set of dwell times) [24].

In each generation, an evaluation function quantifies the quality of each solution through a metric called fitness. A selection mechanism then enforces the principle of "survival of the fittest," giving individuals with higher fitness a greater probability of passing their genes to the next generation [25, 26, 27, 28]. Through crossover (recombination), the genes of selected parents are combined to explore new regions of the search space [29]. Furthermore, to prevent the permanent loss of specific gene values and avoid premature convergence to local optima, a mutation operation is applied. This introduces a small probability of random variation, stimulating exploration beyond the initial genetic information present in the population or revisiting solutions that might have been discarded prematurely [30, 31].

Classical EAs are often viewed as black-box optimizers that make no assumptions about the underlying mathematical properties of the problem. However, optimization problems can be categorized based on the availability of domain knowledge: White-Box (full knowledge), Black-Box (no knowledge), and Grey-Box Optimization (GBO), where limited structure is known [32, 33]. Fortunately, our treatment planning problem includes information that allows us to define it as GBO problem, enabling more efficient evaluations. The specific mechanics of how this is achieved will be explained in the following section.

2.4 Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm

To improve upon standard EAs, this research utilizes the Real-Valued Gene-pool Optimal Mixing Evolutionary Algorithm (RV-GOMEA). Unlike "blind" operators that swap variables randomly, RV-GOMEA exploits the linkage structure (the dependencies between variables) to perform targeted variation [34, 35].

Central to RV-GOMEA is the Family of Subsets (FOS), denoted as $\mathfrak{F}$. The FOS is a collection of subsets of variable indices that are deemed dependent on each other. This structure is learned dynamically from the population (often as a Linkage Tree) or predefined [36, 37]. The FOS can be represented as a hierarchical clustering structure (detailed in Section 3.4.1, Figure 3.1), where leaves represent individual variables and higher nodes group interdependent variables together.

The high-level procedure of RV-GOMEA is outlined in Algorithm 1. In each generation, the algorithm performs the following steps: **Selection** (a subset of the best solutions (size $\lfloor \tau n \rfloor$) from the population P to form a selection pool S), **Model Learning** (for every subset of variables $\mathfrak{F}_j$ in the FOS, the algorithm estimates a Maximum-Likelihood multi-

variate normal distribution $\mathcal{N}(\hat{\boldsymbol{\mu}}_{\mathfrak{F}_j}, \hat{\boldsymbol{\Sigma}}_{\mathfrak{F}_j})$ based on the parameters in S), and **Variation** (new solutions are generated using the Gene-pool Optimal Mixing (GOM) operator).

The variation mechanism is the GOM operator (Algorithm 2). Unlike standard crossover, GOM iterates through each solution in the population (the "parent") and attempts to improve it by modifying small subsets of variables at a time, defined by the FOS.

To function effectively in continuous domains, RV-GOMEA integrates two key mechanisms from the AMaLGaM framework to guide the sampling process: Adaptive Variance Scaling and Anticipated Mean Shift [38].

- **Adaptive Variance Scaling (AVS):** To counteract the variance-diminishing effect of selection, the estimated covariance matrix $\hat{\boldsymbol{\Sigma}}_{\mathfrak{F}_j}$ is scaled by a distribution multiplier, denoted as $c^{\text{Multiplier}}_{\mathfrak{F}_j}$. This factor is adapted dynamically based on whether improvements are being found, expanding the search variance when strictly following the distribution leads to stagnation.

- **Anticipated Mean Shift (AMS):** To accelerate convergence along slopes in the search space, AMS shifts the mean of the sampling distribution in the direction the population is moving. This shift is denoted by the vector $\hat{\boldsymbol{\mu}}^{\text{Shift}}$. In the algorithm, δ^{AMS} represents the scaling factor applied to this mean shift.

Summarizing Algorithm 2, the variation for a single parent solution P_i is as follows: First, before modifying the variables in the current FOS subset $\mathfrak{F}_j$, the current state of P_i is saved into a backup vector b (Line 5). This allows the algorithm to revert changes if they are detrimental. Then, new values $\boldsymbol{y}$ are sampled from the normal distribution $\mathcal{N}$, using the covariance matrix scaled by the AVS multiplier (Line 7). After that, the best portion of the population are shifted by $\delta^{\text{AMS}}\hat{\boldsymbol{\mu}}^{\text{Shift}}$ (Line 8-9). Finally, the new values are inserted into P_i, and the solution is evaluated. If the objective value improves, the change is kept.

Note that if the objective value worsens, the solution is usually reverted to the backup state b. However, to prevent getting stuck in local optima, there is a small probability ($p^{\text{accept}} = 0.05$) that a worsening move is accepted (Lines 12-16).

2.4.1 Multi-Objective Optimization

While RV-GOMEA handles the continuous nature of the variables, the problem of designing a treatment plan is inherently multi-objective, requiring a balance between maximizing tumor coverage and minimizing OARs' sparing. To address this, BRIGHT uses the Multi-Objective version of RV-GOMEA, MO-RV-GOMEA [39].

Rather than aggregating all clinical goals into a single fitness value, MO-RV-GOMEA identifies a set of solutions that represent optimal trade-offs, collectively referred to as the Pareto front. A solution is considered to dominate another if it is superior in at least one objective and not inferior in any other. The Pareto front comprises all non-dominated solutions. By applying these dominance concepts, MO-RV-GOMEA maintains a diverse population that comprehensively covers the spectrum of possible trade-off compromises. As a

result, instead of a single output, researchers are provided with a Pareto approximation front with a range of valid options [40].

Algorithm 1 The basic structure of RV-GOMEA.

function RV-GOMEA
  $P \leftarrow$ Initial population, evaluated
  pop-NIS $\leftarrow t \leftarrow 0$
  for $P_i \in P$ **do** $\text{NIS}(P_i) \leftarrow 0$
  while not terminated **do**
    $S \leftarrow$ The best $\lfloor \tau n \rfloor$ solutions in P `/* τ = 0.35 [35]*/`
    $\mathfrak{F} \leftarrow$ Learn FOS from S `/* Unless FOS is fixed.*/`
    for $\mathfrak{F}_j \in \mathfrak{F}$ **do**
      Estimate $\mathcal{N}(\hat{\boldsymbol{\mu}}_{\mathfrak{F}_j}(t), \hat{\boldsymbol{\Sigma}}_{\mathfrak{F}_j}(t))$ with maximum-likelihood based on S
      $\hat{\boldsymbol{\mu}}^{\text{Shift}}_{\mathfrak{F}_j}(t) \leftarrow \hat{\boldsymbol{\mu}}_{\mathfrak{F}_j}(t) - \hat{\boldsymbol{\mu}}_{\mathfrak{F}_j}(t-1)$ `/* Skipped in the initial generation*/`
    for $i \in \{0, \ldots, n^{\text{elitist}} - 1\}$ **do** $P_i \leftarrow i^{\text{th}}$ best solution in P
    generateNewSolutions 2
    $t \leftarrow t + 1$

Algorithm 2 Generating new solutions, GOM in RV-GOMEA.

function GENERATENEWSOLUTIONS
  for $i \in \{n^{\text{elitist}}, \ldots, n-1\}$ **do** $\text{improved}(P_i) \leftarrow \text{False}$
  for $\mathfrak{F}_j \in \mathfrak{F}$ **do** /* In a random order.*/
    for $i \in \{n^{\text{elitist}}, \ldots, n-1\}$ **do**
      for $u \in \mathfrak{F}_j$ **do** $b[u] \leftarrow P_i[u]$ /* Save current values to 'b' to allow reversion.*/
      $f_b \leftarrow f_{P_i}$ /* Save current fitness*/
      $\boldsymbol{y} \leftarrow \mathcal{N}\left(\hat{\boldsymbol{\mu}}_{\mathfrak{F}_j}, c^{\text{Multiplier}}_{\mathfrak{F}_j}\hat{\boldsymbol{\Sigma}}_{\mathfrak{F}_j}\right)$ /* Use AVS multiplier c*/
      if $i < n^{\text{elitist}} + \lfloor \frac{1}{2}\tau n \rfloor$ **then**
        $\boldsymbol{y} \leftarrow \boldsymbol{y} + c^{\text{Multiplier}}_{\mathfrak{F}_j}\delta^{\text{AMS}}\hat{\boldsymbol{\mu}}^{\text{Shift}}_{\mathfrak{F}_j}$ /* Apply AMS to accelerate search*/
      for $u \in \mathfrak{F}_j$ **do** $P_i[u] \leftarrow \boldsymbol{y}[u]$ /* Insert new values into solution*/
      $f_{P_i} \leftarrow f(P_i)$
      if $f_{P_i} < f_b$ **then**
        $\text{improved}(P_i) \leftarrow \text{True}$
      else if $U(0,1) > p^{\text{accept}}$ **then**
        for $u \in \mathfrak{F}_j$ **do** $P_i[u] \leftarrow b[u]$ /* If no improvement, restore from backup 'b'*/
        $f_{P_i} \leftarrow f_b$
    AdaptMultiplier$\left(c^{\text{Multiplier}}_{\mathfrak{F}_j}\right)$ /* Adapts c based on the improvements*/
  for $i \in \{n^{\text{elitist}}, \ldots, n^{\text{elitist}} + \lfloor \frac{1}{2}\tau n \rfloor - 1\}$ **do**
    $b \leftarrow P_i; f_b \leftarrow f_{P_i}$
    $P_i \leftarrow P_i + \delta^{\text{AMS}}\hat{\boldsymbol{\mu}}^{\text{Shift}}$
    $f_{P_i} \leftarrow f(P_i)$
    if $f_{P_i} < f_b$ **then**
      $\text{improved}(P_i) \leftarrow \text{True}$
    else if $U(0,1) > p^{\text{accept}}$ **then**
      $P_i \leftarrow b$
      $f_{P_i} \leftarrow f_b$
  for $i \in \{n^{\text{elitist}}, \ldots, n-1\}$ **do**
    if $\text{improved}(P_i)$ **then** $\text{NIS}(P_i) = 0$
    else $\text{NIS}(P_i) \leftarrow \text{NIS}(P_i) + 1$
    if $\text{NIS}(P_i) > \text{NIS}^{\text{MAX}}$ **then**
      ForcedImprovements(P_i) /* [36] with pseudocode in [35]*/

2.5 BRIGHT

The theoretical framework culminates in the development of BRIGHT (BRachytherapy via artificially Intelligent GOMEA-Heuristic based Treatment planning). BRIGHT represents the clinical implementation of the algorithmic principles described in Section 2.4. It is an inverse treatment planning system that utilizes MO-RV-GOMEA applied to the multi-objective problem definition of HDR brachytherapy planning. Unlike manual planning, where a clinician must iteratively adjust parameters by trial and error, BRIGHT directly optimizes the dwell times to automatically generate high-quality dose distributions.

The core innovation of BRIGHT lies in its formulation of the planning problem. While clinical protocols contain dozens of specific constraints, BRIGHT simplifies this into a bi-objective optimization task [9]. It balances the two conflicting goals of brachytherapy: maximizing radiation coverage of tumor targets versus minimizing exposure to healthy tissue. To quantify these opposing forces, BRIGHT aggregates the various clinical aims into two unified objectives using a "weighted worst-case" approach: the wegihted least coverage index (LCI_w) and the Least Sparing Index (LSI_w). For a plan p are computed as the sum over either every coverage or sparing aim a as:

$$\begin{aligned} LCI_w(p) &= \sum_{a \in \text{coverage aims}} w_a(\delta(\text{DVI}^a)) \\ LCI_w(p) &= \sum_{a \in \text{coverage aims}} w_a(\delta(\text{DVI}^a)) \end{aligned} \tag{2.1}$$

To determine how close a DVI is to the desired clinical aims, BRIGHT first calculates the normalized deviation for each DVI. These values are normalized between -1 and 1 which is essential because clinical protocols usually define constraints in conflicting units (e.g., Gy for dose vs. % for volume) [9].

For coverage aims (maximizing dose/volume) δ_c^{norm} is computed as:

$$\delta_c^{norm}(V_d^o) = (V_d^o - V_d^{o,aim}) \;/\; (100.0 - V_d^{o,aim}) \tag{2.2}$$

$$\delta_c^{norm}(D_v^o) = (D_v^o - D_v^{o,aim}) \;/\; (0.85 \times D_v^{o,aim}) \tag{2.3}$$

For sparing aims (minimizing dose/volume) δ_s^{norm} is calculated by:

$$\delta_s^{norm}(V_d^o) = (V_d^{o,aim} - V_d^o) \;/\; V_d^{o,aim} \tag{2.4}$$

$$\delta_s^{norm}(D_v^o) = (D_v^{o,aim} - D_v^o) \;/\; D_v^{o,aim} \tag{2.5}$$

Then, the weighted worst-case aggregation does not simply take the single worst value. Instead, it employs a dynamic weighting strategy to optimize the "weakest links" while still improving other values.

The normalized delta DVIs are sorted in the order of most violated to least violated and assigned a weight $w_o(\delta)$ to each DVI [9].

Finally, for the grouped objective O with weights w_o^i for a treatment plan P is calculated as:

$$O(P) = \sum_{\delta_o^{norm} \in O} w_o(\delta_o^{norm}) \tag{2.6}$$

In this formulation, a positive δ^{norm} value indicates that the specific clinical aim has been met, while a negative value indicates a violation. Consequently, BRIGHT seeks to maximize both LCI_w and LSI_w. Of particular clinical interest is the "Golden Corner," the region of the Pareto approximation front where both $LCI > 0$ and $LSI > 0$. As illustrated in Figure 2.1, the 'Golden Corner' is the upper-right quadrant where both LCI and LSI are positive. Plans located in this region satisfy every aim defined in the clinical protocol, representing the ideal solution space. Then, the clinician can select the most appropriate plan for the patient from that region or the surrounding space, assuming the compromise if they consider it pertinent.

Currently, BRIGHT has successfully bridged the gap from research to clinic and is being used at Amsterdam UMC for the treatment of prostate cancer using HDR brachytherapy [41].

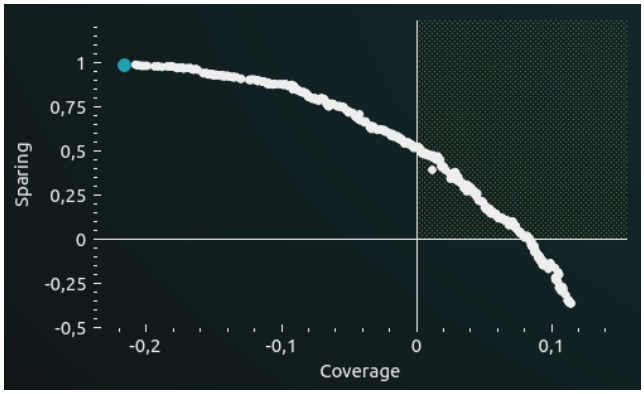

Figure 2.1: Pareto approximation front of generated plans on the coverage-sparing space with the "golden corner" marked in yellow, as shown in the BRIGHT application.

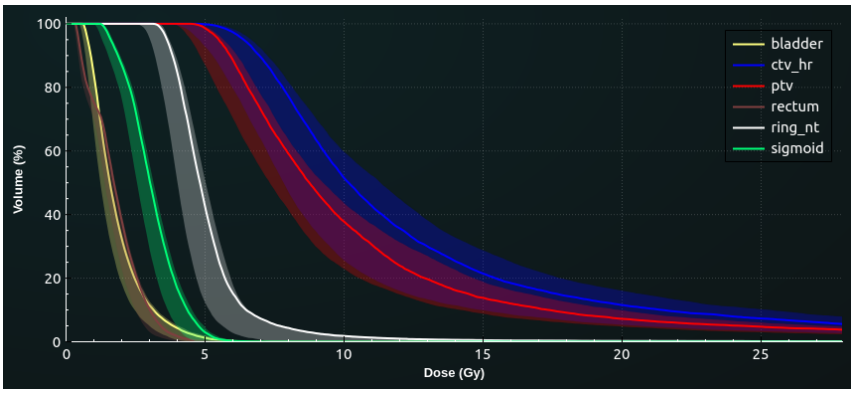

Figure 2.2: Dose-Volume Histogram plotting the radiation at the different ROIs, as shown in the BRIGHT application.

Chapter 3

Development of a Clinical Protocol (Methodology)

This chapter details the methodology for the automated discovery of clinical protocols. We first formalize the bi-level optimization problem, in which an outer evolutionary loop tunes protocol parameters to minimize the difference in radiation between generated and clinically approved treatment plans, and the inner loop generates treatment plans and evaluates the results.

Subsequently, we describe the specific algorithmic implementation and the experimental framework used to validate our hypotheses.

3.1 Problem Formulation

The fundamental premise of this research is that the "implicit knowledge" of the standard of care of a medical institution is encoded within its historical data of treated patients. Therefore, instead of manually tuning parameters to achieve a desired output, we invert the workflow: we use the desired outputs (historical plans) to solve for the input parameters (the protocol) that produced them.

We can formulate this as a minimization problem.

Let $\mathcal{D} = \{p^1_{cli}, \ldots, p^N_{cli}\}$ be a dataset of N historical cases, where p^i represents the patient anatomy and p^i_j the treatment plan for that anatomy, cli refers to the clinically approved treatment plan, the "ground truth". For generated plans, j is the plan number and i is the patient ID.

We define the treatment planning system BRIGHT as a stochastic function B that maps a clinical protocol $\mathbf{x}$ and patient anatomy p^i to a generated Pareto approximation front $\mathcal{F}^i_{gen}$:

$$\mathcal{F}^i_{gen} = B(p^i, \mathbf{x})$$

The protocol $\mathbf{x}$ is formalized as a vector of decision variables $\theta \in \mathbb{R}^d$, representing the numerical aims and weights of the Dose-Volume Indices (as defined in Section 2.2).

The optimization goal is to find the optimal protocol configuration $\mathbf{x}^*$ that minimizes the difference between the generated fronts $\mathcal{F}_{gen}$ and the clinical reference plans across the population.

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^{n} \mathcal{L}(B(p^i, \mathbf{x}), p^i_{cli})$$

Where $\mathcal{L}$ is the loss function quantifying the difference between the generated front and the expert plan. The specific formulations of this loss function ($\mathcal{L}$) correspond to the Closest Plan Similarity and Average Front Similarity metrics, which are detailed in the following section.

3.1.1 Loss Function

To quantify the similarity between the treatment plans in the generated front $\mathcal{F}_{gen}$ and the clinical plan p^i_{cli}, we randomly sample a set of radiation positions S in the space of each of the ROIs V and calculate the distance between any two individual plans using the Mean Squared Error (MSE) of the dose D at each point in S as in equation 3.1.

$$dist(p^i_j, p^i_{cli}) = \frac{1}{|V|} \frac{1}{|S|} \sum_{v \in V} \sum_{s \in S} (D^i_j(v,s) - D^i_{cli}(v,s))^2 \tag{3.1}$$

In this research, we investigated two distinct hypotheses for what best constitutes a ”similar” front, leading to two different loss functions.

Loss Function 1: Closest Plan Similarity (CPS)

This formulation hypothesizes that a good protocol must be capable of generating at least one plan that closely matches the expert's choice. The error for a patient p^i is therefore the distance of the closest plan on the generated front to the clinical plan.

$$\mathcal{L}^{(CPS)}_i(B(p^i, \mathbf{x}), x^i_{cli}) = \min_{x_j \in \mathcal{F}_{i,gen}} (dist(p^i_j, p^i_{cli})) \tag{3.2}$$

The overall objective is to minimize the average of this loss over all patients. Therefore, we define the fitness objective as:

$$\underset{\mathbf{x}}{\text{minimize}}\ F^{(CPS)}(\mathbf{x}) = \frac{1}{|N|} \sum_{i \in N} \mathcal{L}^{(CPS)}_i(B(p^i, \mathbf{x}), x^i_{cli}) \tag{3.3}$$

Loss Function 2: Average Front Similarity (AFS)

This alternative formulation hypothesizes that a good protocol should produce a front where the entire set of trade-off solutions is, on average, located near the expert's choice. This measures the overall quality and location of the front. The error for patient i is the average distance of all plans on the generated front to the clinical plan.

$$\mathcal{L}_i^{(AFS)}(B(p^i,\mathbf{x}),x_{cli}^i) = \frac{1}{|\mathcal{F}_{i,gen}|}\sum_{x_{i,j}\in\mathcal{F}_{i,gen}} dist(p_j^i,p_{cli}^i) \tag{3.4}$$

The overall objective is to minimize the average of this loss over all patients. Therefore, we define the fitness objective as:

$$\underset{\mathbf{x}}{\text{minimize}}\ F^{(AFS)}(\mathbf{x}) = \frac{1}{|N|}\sum_{i\in N}\mathcal{L}_i^{(AFS)}(B(p^i,\mathbf{x}),x_{cli}^i) \tag{3.5}$$

In the initial evaluation experiments of this thesis, we will determine which fitness calculation method is more suitable for this use case.

3.2 BRIGHT Configuration

In the proposed bi-level optimization framework, BRIGHT functions as the evaluation engine for the outer loop. To generate a Pareto approximation front of treatment plans, BRIGHT requires two distinct categories of inputs: Static hyperparameters that control the resolution and computational cost of the underlying optimization process. Specifically, the number of Dose Calculation Points (DCPs), the allowed runtime per run, and the clinical protocol $\mathbf{x}$, which defines the DVIs that guide the search.

In this research, the hyperparameters are fixed to ensure consistent fitness evaluations, while the protocol $\mathbf{x}$ serves as the variable input encoded within the genotype of the outer evolutionary algorithm. The evaluated output of a BRIGHT run against the clinical ground truth ($F^{(CPS)}(\mathbf{x})$ or $F^{(AFS)}(\mathbf{x})$) constitutes the fitness value used to drive the evolution of the population.

The following subsections detail the calibration mini experiment used to determine the optimal number of DCPs (Section 3.2.1) and the fixed protocol structures used to constrain the search space (Section 3.2.2).

3.2.1 Dose Calculation Points and Runtime

To determine the optimal number of DCPs for the experiments, four scenarios with varying DCP densities were evaluated. These scenarios were tested against a group including an atypical patient extra who required a higher coverage, forcing the clinical treatment plans to be set near the LCI limit for the baseline protocol. This setup was designed to be sensitive to the DCP sampling, creating a risk of clinical plans being marked as infeasible, being too far from the golden corner (less than -0.25 lci or lsi), due to the random selection of points.

The objective was to assess the consistency of the results across the different DCP sample sizes while keeping all other variables constant. The experiment was run twice and consisted of 10 runs of 10 evaluations for 7 different patients (700 BRIGHT runs with different seeds). The evaluated scenarios were as follows:

1. 30,000 DCPs (5,000 per ROI)
2. 60,000 DCPs (10,000 per ROI)
3. 120,000 DCPs (20,000 per ROI)
4. 300,000 DCPs (50,000 per ROI)

The selection of the BRIGHT running time was 15s or 30s for the different scenarios, based on [42]. The results of this evaluation are summarized in Table 3.1.

Table 3.1: Number of evaluations where the clinical plan is considered too far from the golden corner (lower is better).

	First Run (7 Patients)		**Second Run (Again 7 Patients)**	
DCPs	**Problematic Runs**	**Problematic Evaluations**	**Problematic Runs**	**Problematic Evaluations**
30,000 (15 s)	10/10	35/700	10/10	33/700
60,000 (30 s)	7/10	16/700	10/10	26/700
120,000 (30 s)	8/10	13/700	6/10	8/700
300,000 (30 s)	**3/10**	**3/700**	**2/10**	**3/700**

3.2.2 Protocol parameters

Several protocol parameters were standardized to ensure the focus of this study remained on the established research questions. This simplification was necessary to manage experimental complexity and isolate the effects of the variables under investigation. The parameters held constant were: Nominal prescribed dose, DVI limits, number of DVIs per ROI, and set of optimization objectives.

- **Nominal Prescribed Dose:** The prescribed dose was fixed at 7.0 Gy per fraction. Varying this value would only create a series of proportionally scaled plans (e.g., simultaneously decreasing the nominal dose and increasing DVI coverage targets proportionally), which would not add any additional value to the research and just confuse clinicians who are used to this prescribed dose level for HDR brachytherapy.
- **DVI Limits:** To reduce the size of the genotypes, the limits for each DVI were fixed proportionally. This is a hard constraint set on BRIGHT, slightly lower for coverage aims and slightly higher for sparing aims, to allow for some more flexibility in the exploration. The limit for coverage DVIs was set to 10% lower than the target aim, and the limit for sparing DVIs was set to 150% higher than the OAR aim, as deviations beyond these magnitudes typically result in clinically unacceptable plans. This approach standardizes the acceptable deviation from the DVI targets.

- **Number of DVIs per ROI:** The protocol was limited to a single DVI per ROI. While utilising multiple objectives for a single volume has proven valuable in the manual design of treatment plans, incorporating it here would have significantly increased the complexity of these initial experiments. Consequently, exploring multi-DVI protocols is proposed as a potential avenue for future work.
- **Optimization Objectives:** Finally, the number of optimisation objectives was limited to the LCI_w and the LSI_w. While adding a third objective can be beneficial [43], it was deemed to introduce unnecessary complexity for the scope of this research.

3.3 Manual Protocol

To derive a protocol that accurately reflects specific institutional expertise, we established a collaboration with the medical physics team at Virginia Commonwealth University, who actively treat patients with cervical cancer. This collaboration aimed to capture the specific, often implicit, preferences used in their daily clinical practice.

The formulation of this protocol was the result of a multi-stage process involving mutual familiarization and iterative feedback. Initially, the BRIGHT planning system was demonstrated to the clinical experts to establish a shared understanding of the optimization engine's capabilities and limitations. Following this, we engaged in a series of knowledge-exchange sessions to bridge the gap between the researchers' algorithmic perspective and the doctors' clinical intent.

We tested the system using a set of historical patient cases. The medical team reviewed the generated plans, identifying discrepancies between the automated output and their preferred dose distributions. Through this iterative tuning process, we synthesized the specific set of objectives and constraints that replicated their planning style. This derived configuration is designated as the Baseline Protocol and is explicitly defined in Table 3.2.

Table 3.2: The Baseline Protocol derived from discussions with medical professionals.

Coverage DVIs	$D_{90}^{CTV_{HR}}$	D_{90}^{PTV}		
Planning Aims	$> 94\ \%$	$> 80\ \%$		
Limits	$> 85\ \%$	$> 72\ \%$		
Sparing DVIs	$D_{65}^{Bladder}$	D_{65}^{Rectum}	$D_{65}^{Sigmoid}$	$D_{90}^{Ring_{NT}}$
Planning Aims	$< 2\ cm^3$	$< 2\ cm^3$	$< 2\ cm^3$	$< 10\ cm^3$
Limits	$< 3\ cm^3$	$< 3\ cm^3$	$< 3\ cm^3$	$< 10\ cm^3$

3.3.1 Volume Ring Normal Tissue

A critical component of the Baseline Protocol identified during these discussions is the inclusion of an auxiliary sparing volume, referred to as the "Ring of Normal Tissue" ($Ring_{NT}$).

In the absence of constraints on the surrounding tissue, the optimizer may allow high-dose radiation to "spill" into healthy areas to easily achieve target coverage. To prevent this and enforce a margin, the $Ring_{NT}$ is introduced. The ring is defined as a shell surrounding the combined target volumes (CTV_{HR} and PTV) with a thickness of 2 mm. Additionally, to ensure that valid OAR sparing takes precedence, any volume of this shell that overlaps with an OAR (Bladder, Rectum, Sigmoid) is subtracted as those are more restrictive.

3.4 Implementation of the Protocol Discovery Algorithm

The solution is implemented as a direct instance of RV-GOMEA. The algorithmic procedure is detailed in Algorithm 3.

The process is initialized by generating a population P of N random candidate protocols (Line 1). The evolutionary cycle is driven by the GOM operator (Line 6). The Linkage Model is detailed in Section 3.4.1. This ensures that interdependent protocol parameters (metric and corresponding aim) are considered together.

The connection to the clinical domain is handled entirely within the fitness function f (Lines 2 and 7). When the algorithm requires a fitness value for a candidate $\mathbf{x}$, it executes the EvaluateFitness routine. This routine performs the k stochastic BRIGHT runs on the training dataset D and computes the aggregate loss $F(\mathbf{x})$ as defined in Equations 3.3 and 3.5. If the resulting loss is lower than that of the parent, the offspring is accepted into the population immediately (Line 9), ensuring a strict improvement in solution quality.

The general GOMEA algorithm and GOM operator work as defined in [42] and can be seen in Section 2.4.

3.4.1 Linkage Model Configuration

Since the efficiency of the GOM operator relies heavily on the quality of the linkage model. While GOMEA typically uses a linkage learning mechanism to infer dependencies dynamically, this research utilizes a fixed, pre-defined linkage tree constructed based on problem-specific domain knowledge.

The structure of this fixed model is illustrated in Figure 3.1. A critical design choice in this configuration is the exclusion of the univariate FOS (the set containing all individual genes). This decision is motivated by the inherent "atomic" nature of the protocol parameters in the full optimization. For instance, a single clinical objective, such as $CTV_{HR}\ D_{x_1} > x_7\%$, is defined by a pair of two highly correlated variables: the dose metric x_1 and the aim value x_7. Modifying x_1 without a corresponding adjustment to x_7 is likely to result in an ineffective constraint. Due to search space pruning strategies explained in the following section, wrongly combined genes would end up re-initialized.

By defining a fixed linkage tree that groups these interdependent parameters at the lowest level (i.e., placing x_1 and x_7 in the same leaf subset), we force GOMEA to treat them as an indivisible building block. This prevents the GOM operator from disrupting valid metric-constraint pairings, thereby reducing the generation of non-viable offspring (e.g., for the $D^{CTV_{HR}}$ we might want to explore [90, 90] and [50, 150], but not [90, 150] or [50, 90] which might fall into the discarded regions detailed in Section 3.5). Higher levels of the tree aggregate these blocks into logical clinical categories (e.g., grouping all Coverage aims versus Sparing aims), allowing the algorithm to swap complete objective blocks of the protocol during crossover.

Algorithm 3 Outer loop

Input:

I: Protocol encoded individual
N: Population size
T: Termination condition
FOS: Linkage Model
Bright configuration:
D: Training data (anonymized example patients)
k: runs per patient averaged
r: runtime
$dcps$: Dose calculation points.

Output:

The best individual found in the population

$P \leftarrow$ InitializePopulation(N) ▷ Create N random individuals
EvaluateFitness(P, f) ▷ Calculate fitness for each individual
$t \leftarrow 0$
while not T **do**
    for $i = 1$ to N **do**
        $o_i \leftarrow$ GeneOptimalMixing(P_i, P, FOS) ▷ Create one offspring for P_i
        EvaluateFitness($\{o_i\}, f$)
        if $f(o_i) \geq f(P_i)$ **then**
            $P_i \leftarrow o_i$ ▷ Replace parent if offspring is better or equal
        #Store the best individual
    $t \leftarrow t + 1$
return best individual from P

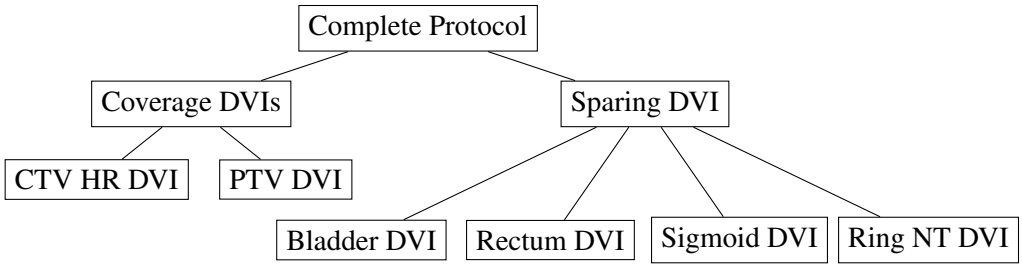

Figure 3.1: Diagram of the general linkage tree structure for the complete DVI experiments.

3.5 Computational challenges and Optimizations

The research methodology described in section 3.1 presents a significant computational challenge: the prohibitive cost of the algorithm described in section 3.4.

The computational bottleneck comes from the fitness evaluations within RV-GOMEA. The fitness evaluation of a single candidate protocol $\mathbf{x}$ (an individual in the EA) requires executing BRIGHT k times for every patient for every individual of the population for every partial evaluation of the GOM operator.

Therefore, the total number of BRIGHT simulations required is:

$$\begin{aligned}\text{Total Runs} = (\text{Generations}) \times (\text{Population Size } N + \text{Partial Evaluations}) \\ \times (\text{Training Patients } P) \times (\text{Stochastic Runs } k)\end{aligned} \tag{3.6}$$

where,

$$\max \text{Partial Evaluations} = N \times (2l - 1) \tag{3.7}$$

This huge number of evaluations makes the number of fitness evaluations computationally expensive. While BRIGHT itself is highly optimized for GPU parallelization, due to hardware constraints, there is a limit on the number of BRIGHT runs that we can run in parallel. This limitation, combined with the sheer volume of runs required for the evolutionary search, makes the ”naive” execution of this methodology computationally infeasible for a proof-of-concept study.

Optimization Strategies
To ensure the feasibility of the experiments, the strategy was not to optimize BRIGHT further, but to minimize the total number of BRIGHT simulations performed by the outer evolutionary algorithm. Four key optimizations were implemented:

1. **Data-Driven Search Space Bounding:** To reduce the search space, the range of possible values for the protocol parameters was heuristically constrained using data from the N training set plans. This process was tailored specifically for target volumes versus OARs:

   - **For Target Volumes:** The DVHs from all training patients were aggregated for the specific volume O to determine the lower bound of clinically-achieved values (e.g., the lowest $D_{90\%}$). The search range for the corresponding objective was then set to the lower bound of the aggregated DVH with a 25% margin*: $[DVH_{o,lb} \times 0.75, DVH_{o,lb} \times 1.25]$.
   - **For Organs at Risk:** The DVHs were similarly aggregated to determine the upper bound of clinically-achieved values (e.g., the highest $V_{65\%}$). The search range for the corresponding objective was set to the upper bound of the aggregated DVH with a 50% margin*: $[0, DVH_{o,ub} \times 1.50]$.

   * These margins were set heuristically based on the results for Section 5.2.1 Figure 5.1.

   This optimization is based on the assumption that the optimal protocol $\mathbf{x}^*$ will define objectives reasonably close to those already achieved by expert clinicians, allowing the EA to focus its search on a smaller region of the parameter space.

2. **Rejection by Clinical Plan Feasibility:** This optimization uses the central hypothesis: the expert's clinical plans (p_{cli}) are high-quality, feasible solutions. A candidate protocol $\mathbf{x}$ defines the objectives used to calculate BRIGHT's LCI and LSI scores (as defined in Section 2.5). A validation step was introduced to check if p_{cli} would fail the candidate protocol $\mathbf{x}$. For each patient, the $LCI(p_{cli})$ and $LSI(p_{cli})$ scores were calculated using the DVIs from $\mathbf{x}$. If $LCI(p_{cli}) < -0.25$ or $LSI(x_{cli}) < -0.25$, it signifies that the expert's own plan would be too far from the golden corner and considered unacceptable by the candidate protocol $\mathbf{x}$. Such protocols were deemed "overly restrictive" and were immediately assigned a worst-possible (invalid) fitness score, saving all $N \times k$ simulation runs.

3. **Limited Stochastic Averaging:** A pragmatic trade-off was made between the accuracy of the fitness estimate and the available computational time. The number of repeated runs for the fitness calculation was limited to $k = 2$ per patient. While a higher k would provide an easier-to-navigate search space, $k = 2$ was deemed a sufficient compromise. It helps to smooth minor variations in front generation without rendering the experiments infeasible, which is acceptable for this proof-of-concept investigation.

4. **Termination Criteria:** To balance search efficiency with computational feasibility, a dual termination strategy was employed.

   a) **Adaptive Early Stopping:** The primary criterion uses a convergence-based heuristic. The evolutionary run is terminated if no improvement is observed in the population's best-found solution for n consecutive generations. Based on preliminary convergence observations, this threshold was set to $n = 2$, effectively preventing the algorithm from expending resources on a stagnated search.

   b) **Maximum Computational Budget:** As a safeguard, a hard upper limit on the total time the experiment can run. This ensures that in cases of slow convergence, the experiment stays feasible.

   The algorithm terminates when the first of these two conditions is met.

Chapter 4

Experiments

This chapter details the experimental framework used to evaluate the automated protocol discovery method proposed in Chapter 3. We begin by describing the clinical dataset, which serves as the ground truth for our research. Then, we outline the progression of four experiments, incrementally increasing the complexity of the genotype representation and formulating hypotheses on how it will affect the results.

4.1 Design of Experiments

4.1.1 Dataset

The dataset used in this research consists of anonymized clinical data from 7 patients who underwent brachytherapy for cervical cancer at Virginia Commonwealth University (VCU). The dataset was partitioned into:

- A training set of 6 patients, used to guide the evolutionary optimization algorithm and draw initial conclusions from the progressively more complex experiments.
- A validation patient, randomly selected before the study, held out and used exclusively for the final evaluation of the optimized protocol, for an unbiased assessment of its generalization capabilities.

For each patient case, the dataset includes:

- The 3D patient geometry derived from CT scans, which show the applicator catheters in place.
- The pre-defined dwell positions for the radiation source within the catheters.
- A set of delineated 3D structures, including: the Clinical Target Volume - High Risk (CTV_{HR}) and Planning Target Volume (PTV), and the relevant OARs, the bladder, rectum, and sigmoid.

- The clinically-approved treatment plan, specified as a vector of dwell times. As previously mentioned, this plan represents the "ground truth" for this study, as it was created by a medical physicist and deemed safe and effective for treatment.

Given the sample size, this research should be considered a proof-of-concept. The primary objective is to establish the feasibility and potential of using an evolutionary algorithm to automate the development of a clinical protocol for BRIGHT, rather than to produce a protocol that can create improved treatment plans.

4.1.2 Protocol Representation and Genotype Encoding

As defined in Section 3, a clinical protocol $\mathbf{x}$ consists of a collection of inequalities, where each inequality represents a DVI imposed on a specific ROI as defined in Section 2.2.

Our optimization approach explores the protocol space by evolving different parts of the DVI. We distinguish between two types of parameters:

1. Right-Hand Side (RHS): The Aim and Limit of the DVI (e.g., optimizing the $2cm^3$ threshold).

2. Left-Hand Side (LHS): The type and quantity of the Metric (e.g., optimizing the 65% isodose level).

The genotype of an individual in the evolutionary algorithm is represented as a real-valued vector $\vec{g} \in \mathbb{R}^d$. Each gene g_i in this vector corresponds to a specific parameter of the protocol. The dimensionality d and the specific mapping of genes to protocol parameters vary across the experiments to test the impact of representational freedom (Hypothesis 2). The specific encoding schemes for the structured (RHS-only), partial (Single or Double DVI), and unconstrained (Full Protocol) experiments are detailed in Sections 4.3, 4.4, 4.5, respectively.

4.2 Search Space Exploration Experiment

Before proceeding to protocol optimization, it is crucial to understand the search space to better analyze the results from the experiments described in the following Sections.

4.2.1 Objective

The objective of this experiment is to perform a comprehensive landscape analysis by isolating a single DVI for the CTV_{HR} volume. By systematically mapping the fitness values across a grid of possible metric-aim combinations, we will be able to visualize the "noise" introduced by BRIGHT and determine the computational resources required to mitigate it. Furthermore, we investigate the impact of manipulating the fitness function via normalization to potentially accelerate the search process.

This experiment is guided by **RQ1**, which we have split into two specific research questions:

Research Question 1.1 (Stochastic Mitigation)

Does increasing the computational cost per evaluation, the number of averaging runs k per patient, yield a smoother fitness landscape?

Research Question 1.2 (Convergence Velocity)

Can the convergence rate of the automated discovery process be accelerated by introducing a normalization factor that artificially steepens the fitness gradient in regions distant from the clinical ground truth?

4.2.2 Hypotheses

The main factor in this experiment is the stochastic nature of the fitness function. We postulate the hypothesis (**H1**) that applying specific landscape smoothing and shaping techniques will provide a sufficiently stable signal for the evolutionary algorithm to converge, despite the noise of the inner loop. To expand on this, we propose the following specific hypotheses for each research question:

- **Hypothesis 1.1 - Variance Reduction via Averaging:** We hypothesize that the fitness landscape constructed using a low number of averaging runs ($k = 2$) will exhibit significantly more noise, characterized by spurious local optima that may mislead the search. In contrast, increasing the averaging to $k = 10$ or higher is expected to smooth the noise, revealing the region where the optimal should be and facilitating a more reliable navigation for the optimizer.
- **Hypothesis 1.2 - Gradient Steepening via Normalization:** We hypothesize that the introduction of a normalization factor will improve convergence speed. By penalizing protocols that generate plans far from the region of interest, the normalization factor

is expected to worsen the fitness scores of "uninteresting" areas of the search space. This should theoretically create a steeper descent gradient, effectively "pushing" the population members toward the clinically viable region more rapidly than a non-normalized objective function.

4.2.3 Experimental Configuration

To rigorously characterize the topology and noise of the search space, we designed a grid-search experiment isolating the effects of a single DVI for CTV_{HR}. In this configuration, the DVIs for the PTV and all OARs were held constant at their baseline protocol values ($\mathbf{x}_{base}$) 3.3 to prevent effects from objective dependencies.

The search space was defined as a two-dimensional plane formed by varying the metric and the aim of the CTV_{HR} DVI. The sampling grid was constructed with a uniform step size of 2% along both axes:

- **Metric Parameter (Volume):** Sampled within the range $[7\%, 99\%]$* of the ROI.
- **Aim Value (Dose):** Sampled within the range $[1\%, 249\%]$* of the nominal prescribed dose (7.0 Gy).

* The radiation limit was set at 250% as an extreme value that we don't expect as a result of the optimization. Then the volume starts at 7%, which is the value of the aggregated DVH at 250% of dose.

To accurately quantify the variance introduced by the underlying randomized optimization in BRIGHT, we performed $k = 50$ independent planning runs for each discrete coordinate pair sampled within this grid. This way we can compare three averaging scenarios: $k = 2$, $k = 10$, $k = 50$.

Then, to investigate the impact of gradient steepening (Hypothesis 1.2), we introduced a normalization factor to the objective function. This factor penalizes DVIs that include low values (i.e., near-zero dose or volume), theoretically accelerating convergence by steering the search away from uninteresting regions. The normalized fitness function for coverage DVIs is defined as:

$$F_{norm}^{(AFS)} = F^{(AFS)} + \frac{1}{\frac{x_{vol}}{100} * \frac{x_{dose}}{100}} \tag{4.1}$$

where the fitness corresponds to the calculated AFS metric, and x_{vol} and x_{dose} represent the percentage values of the metric and aim parameters, respectively.

4.3 Initial Evaluation Experiment

4.3.1 Objective

The objective of this experiment is to establish the feasibility of the RV-GOMEA-based optimization framework by applying it to a constrained, well-defined problem: the optimization of protocol aims (RHS) within a fixed structure. As previously mentioned, this initial test serves as a validation of both the algorithm's ability to navigate the search space and the utility of the proposed Average Front Similarity ($F^{(AFS)}$) fitness function defined in Equation 3.4. The optimized protocols for this experiment will be referred as $\mathbf{x}_{E1}$

This experiment is guided by **RQ2**, which we have split into two more specific research questions:

Research Question 2.1 (Algorithmic Convergence)

Can the RV-GOMEA-based algorithm minimize the $F^{(AFS)}$ error by optimizing the six aims of a fixed-structure protocol, achieving a fitness score superior to that of the pre-defined clinical Baseline Protocol?

Research Question 2.2 (Metric Validation)

Does minimizing the $F^{(AFS)}$ objective function correlate with clinical relevance? Specifically, does the optimized protocol $\mathbf{x}^*_{E1}$ generate treatment plans that are quantitatively and qualitatively similar to the expert-approved clinical plans?

4.3.2 Hypotheses

The premise of this experiment is that an evolutionary algorithm can extract implicit expert knowledge from a dataset of approved clinical plans to discover a single, generalized protocol $\mathbf{x}^*_{E1}$. We hypothesize (**H2**) that when applied to unseen patient cases, this discovered structured protocol will enable BRIGHT to automatically generate treatment plans that are quantifiably similar to those created by human experts.

Specifically, we propose the following specific hypotheses for each research question.

- **Hypothesis 2.1 - Convergence in a Structured Protocol:** We hypothesize that the EA will successfully converge to a region of the search space with low error, identifying a protocol vector $\mathbf{x}^*_{E1}$ that generates treatment plans with quantifiable higher similarity to the clinical ground truth than those produced by the baseline protocol. This expectation is because of the following two factors:

  1. The fixed protocol structure restricts the optimization to 6 dimensions, reducing the search complexity sufficiently for the algorithm to navigate within the limited computational budget.

2. Since the baseline protocol values lie within the search bounds, the algorithm should, at a minimum, rediscover these values, expecting at least a performance equal to that of the baseline protocol.

- **Hypothesis 2.2 - Clinical Alignment:** We expect that the optimized parameter values in $\mathbf{x}^*_{E1}$ will align closely with the baseline protocol. By enforcing clinically relevant metrics from the start (rather than allowing the algorithm to explore arbitrary proportional isodoses), we constrain the search to a clinically interpretable subspace, which we hypothesize will lead to parameters that mirror established expert consensus.

4.3.3 Experimental Configuration

Genotype Encoding

The fixed structure of the clinical protocol is the following: The metrics for the targets are fixed to $D_{90\%}$, and the metrics for the OARs are fixed to $V_{65\%}$ and $V_{90\%}$, based on the baseline protocol described in Section 3.3. Therefore, the optimization task of the aims for these metrics is as follows.

The genotype is a 6-dimensional real-valued vector, $\mathbf{x}_{E1} = (x_1, x_2, x_3, x_4, x_5, x_6)$, where each gene maps directly to a specific aim:

$$\begin{aligned} D_{90\%}^{CTV_{HR}} &> x_1\% \\ D_{90\%}^{PTV} &> x_2\% \\ V_{65\%}^{Bladder} &< x_3 \text{ cm}^3 \\ V_{65\%}^{Rectum} &< x_4 \text{ cm}^3 \\ V_{65\%}^{Sigmoid} &< x_5 \text{ cm}^3 \\ V_{90\%}^{Ring_{NT}} &< x_6 \text{ cm}^3 \end{aligned}$$

Population Size and Runtime

To determine the size of the population for these experimental runs, we used the univariate formula 4.2 [34].

$$\textit{Population Size} \approx 10\sqrt{\textit{Genotype Size}} \tag{4.2}$$

$$10\sqrt{6} \approx 25$$

The experimental runs were set to run for a maximum of 1 month ($\approx 720h$).

Algorithm Configuration

The outer loop uses RV-GOMEA to evolve the protocol population. The inner loop uses BRIGHT to evaluate the fitness of each candidate protocol. The configurations are detailed in Table 4.1.

Table 4.1: Parameter values for the Initial Evaluation Experiment. Optimization algorithm (Outer loop). Fitness function (Inner loop, BRIGHT).

Parameter	Value
Outer Loop: RV-GOMEA (Protocol Search)	
Training data	6 patients
Population Size	25 4.2
Linkage Model	Univariate
Termination criteria	Runtime $\leq$ 720 h
Inner Loop: BRIGHT (Fitness Evaluations)	
Total runs	12 (2×6 patients)

Performance Metrics

The experiment results are evaluated during the optimization process, and on the outcome:

1. Optimization Process Metrics: The evolutionary run is guided by the AFS objective function 3.4. AFS was chosen for this initial experiment as a result of small observations where using CPS would consistently set the clinical plan on the edge of the Pareto approximation front, resulting in fewer plans close to the clinical plan in the lsi vs lci space. AFS was chosen because it evaluates the entire Pareto front, rewarding protocols that consistently produce good trade-off solutions and generalizing much better to unseen patients rather than a single plan that happens to be close to the ground truth.

The fitness/loss of the individuals generated by the EA will be the main metric for analyzing the process. This will be visualized by plotting the "Best-so-far" fitness, the fitness of the individual with the lowest $F^{(AFS)}$ discovered up to and including each generation. This fitness value doesn't necessarily point directly to the best individual of the population due to the inherent noise, yet it is a good indicator that the solution falls within the set of better solutions worth investigating.

2. Final Protocol Quality Metrics: The output of the experiment is the final, optimized 6-dimensional protocol vector, $\mathbf{x}^*_{E1}$ (the best-so-far individual from the final generation). The quality of $\mathbf{x}^*_{E1}$ is assessed in two ways:

- Qualitative Analysis of the Protocol: The optimized parameter vector $\mathbf{x}^*_{E1}$ will be compared against the baseline values. This analysis will assess whether the algorithm's solution is similar to the expert-derived one or if it discovered a very distinct set of aim values.
- Qualitative Analysis of the generated plans: The protocol's ability to generate high-quality treatment plans. Plans generated using $\mathbf{x}^*_{E1}$ will be compared to plans generated with the baseline protocol to asses if there is a noticeable improvement past the numerical fitness value.

4.4 Partial Protocol Optimizations

4.4.1 Objective

The objective of this experiment is to evaluate the performance and robustness of the evolutionary algorithm in a higher-dimensional, more flexible search space. Unlike in the initial investigation, in which we optimized protocol aim values within a pre-defined structure, the investigation of simultaneously optimizing the metric definition and the aim values should give us a better insight into the complete protocol optimization, as we are more likely to get stuck on metrics irrelevant to clinical practice. The optimized protocols for this experiment will be referred as $\mathbf{x}_{E2}$

By allowing the algorithm to evolve what to measure in addition to what limit to enforce, we aim to determine whether the EA can autonomously discover clinically relevant DVIs. This experiment is guided by **RQ3**:

Research Question 3

How does increasing the complexity of the genotype, allowing the evolution of metric definitions, affect the convergence behavior of the EA and the quality of the resulting solutions?

To systematically isolate the effects of this increased complexity and validate the findings from the search space research in Section 4.2 regarding fitness landscape topology, the investigation is divided into three incremental sub-experiments:

- CTV_{HR} **optimization:** Optimization of the metric definition (x_{vol}) and aim value (x_{dose}) for the CTV_{HR}. To validate the impact of search-space steepening (Hypothesis 1.2), this sub-experiment is performed twice: once with the standard $F^{(AFS)}$ fitness function and once with the normalized one $F^{(AFS)}_{norm}$.
- PTV **optimization:** Analogous optimization isolated to a single DVI for PTV, also comparing the normalized versus non-normalized fitness functions.
- CTV_{HR} **+** PTV **optimization:** Simultaneous optimization of the metric definitions and aim values for both the CTV_{HR} and PTV (4 dimensions total). This final stage utilizes only the fitness function configuration ($F^{(AFS)}$ or $F^{(AFS)}_{norm}$) that demonstrated superior convergence behavior in the previous experiments.

4.4.2 Hypotheses

We propose the hypothesis (**H3**) that the quality of the discovered protocol is positively correlated with the flexibility allowed in the genotype. Specifically, we postulate that an optimization of both metric definitions and aim values will yield a superior solution compared to optimizing aim values alone, despite the increased complexity of the search space.

To analyze the impact of this complete DVI optimization, we propose the following specific hypotheses.

- **Hypothesis 3.1 - Search Space Complexity:** As the dimensionality and interdependency of the search space increase from the isolated cases to the combined case, the optimization task will become more difficult. We expect this to manifest as slower convergence rates, requiring a greater number of generations to reach a stable fitness plateau compared to the lower-dimensional baselines.
- **Hypothesis 3.2 - Solution Variance:** The individuals generated in the combined experiment are expected to exhibit greater variance across independent runs compared to the initial ones. This is attributed to the expanded search space, which likely contains multiple distinct local optima (different metric-value combinations) that satisfy the clinical criteria equally well.
- **Hypothesis 3.3 - Clinical Plausibility:** Despite the increased complexity, we hypothesize that the EA will robustly identify clinically plausible parameters. Specifically, we expect the optimized metrics to converge near standard clinical indices (e.g., $D_{90\%}$), demonstrating that the algorithm can "rediscover" established clinical wisdom without any explicit hard-coding.

4.4.3 Experimental Configuration

This experiment builds upon the findings of the search space investigation. Based on the "funneling" effect observed in the search space analysis (Section 5.2.2), we explicitly test the hypothesis that penalizing trivial solutions accelerates convergence by using both $F^{(AFS)}$ and $F^{(AFS)}_{norm}$ fitness functions.

Genotype Encoding

The genotype encodes the parameters to be optimized, while all inactive constraints (including all OARs) are held constant at their baseline values.

1. CTV_{HR} optimization:
   - Genotype: $\mathbf{x}_{E2a} = (x_1, x_2)$
   - Mapping: $D^{CTV_{HR}}_{x_1\%} > x_2\%$
   - Comparison: Non-normalized fitness function ($F^{(AFS)}$) vs. normalized fitness function ($F^{(AFS)}_{norm}$)
2. PTV optimization:
   - Genotype: $\mathbf{x}_{E2b} = (x_1, x_2)$
   - Mapping: $D^{PTV}_{x_1\%} > x_2\%$
   - Comparison: Non-normalized fitness function ($F^{(AFS)}$) vs. normalized fitness function ($F^{(AFS)}_{norm}$)
3. $CTV_{HR} + PTV$ (Complete Targets) optimization:

- Genotype: $\mathbf{x}_{E2c} = (x_1, x_2, x_3, x_4)$
- Mapping: $D^{CTV_{HR}}_{x_1\%} > x_3\%$ and $D^{PTV}_{x_2\%} > x_4\%$
- Only non-normalized fitness function. Explained in Section 5.4.

Population Size and Runtime

To determine the size of the population for these experimental runs, we used the multivariate formula 4.3 [34].

$$\textit{Population Size} \approx 3 \times \textit{Genotype Size}^{1.5} + 17 \tag{4.3}$$

$$3 \times 2^{1.5} + 17 \approx 25$$

$$3 \times 4^{1.5} + 17 = 41$$

The experimental runs were set to run for a maximum of 1 week each ($\approx 160h$).

Algorithm Configuration

The evolutionary settings are summarized in Table 4.2. Note that, due to computational constraints, this experiment ran for only a few days to ensure it could be repeated several times.

Table 4.2: Parameter values for Experiment 3. Optimization algorithm (Outer loop). Fitness function (Inner loop, BRIGHT).

Parameter	Value
Outer Loop: RV-GOMEA (Protocol Search)	
Training data	6 patients
Population Size	25 / 41 4.3
Linkage Model	Described in Section 3.4.1
Termination criteria	Runtime $\leq$ 160 h
Repetitions	6
Inner Loop: BRIGHT Configuration (Fitness Evaluations)	
Total runs	12 (2×6 patients)

Performance Metrics

Consistent with previous experiments, performance is assessed at both the process and outcome levels. However, given the expanded search space and the additional comparison of this experiment (Normalized vs. Non-normalized), that additional comparison is considered:

1. Optimization Process Metrics: We track the generation-mean and best-so-far fitness trajectories to quantify the "speed" of evolution between using $F^{(AFS)}$ or $F_{norm}^{(AFS)}$ (in the single DVI experiments), serving to validate the search space steepening effect (Hypothesis 1.2).

For a fair comparison, the normalization penalty factor has been removed from the final fitness values shown in the plot.

The plots used to evaluate the results are:

- Generation-mean fitness: The average $F^{(AFS)}$ error of all individuals in the population at each generation.
- Best-so-far fitness: The lowest $F^{(AFS)}$ error found in the population at or until each generation.

2. Final Protocol Quality Metrics: The output of the experiment is evaluated not just on performance, but also on clinical plausibility and search space characteristics:

- Qualitative Analysis of the Protocol: The optimized parameter vector $\mathbf{x}_{E2}^*$ will be compared on the different scenarios to determine how it varies on the different runs, as well as compared to the baseline protocol.
- Qualitative Analysis of the generated plans: The protocol's ability to generate high-quality treatment plans. Plans generated using $\mathbf{x}_{E2c}^*$ will be compared to plans generated with the baseline protocol to assess if there is a noticeable improvement.

4.5 Complete Protocol Optimization

4.5.1 Objective

The objective of this final experiment is to assess the scalability of the optimization approach on the complete, unconstrained protocol discovery problem. This involves the simultaneous optimization of both the metric definitions and the aim values for all six DVIs (two target volumes and four OARs), resulting in a 12-dimensional search space.

Furthermore, this experiment investigates the relationship between computational power, landscape smoothness, and convergence velocity. By leveraging massive parallelization, we aim to determine if increasing the stochastic averaging factor (k) (thereby smoothing the fitness landscape as previously observed in 4.2) enables the EA to navigate this complex space more efficiently. The complete unconstrained set of protocol parameters optimized on this experiment will be referred to as ($\mathbf{x}_{E3}$).

This experiment is guided by **RQ4**, which we have split into two research questions:

Research Question 4.1 (Feasibility & Cohesion)

Can GOMEA successfully navigate the 12-dimensional search space to discover a clinical protocol that produces treatment plans similar enough to the plans made by experts?

Research Question 4.2 (Search Space Smoothing Results)

Does mitigating the stochastic noise of the fitness function through high-throughput parallelization (increasing averaging runs from $k = 2$ to $k = 16$) result in faster convergence and superior final solution quality?

4.5.2 Hypotheses

This final experiment addresses the scalability of the proposed framework. We postulate the hypothesis (**H4**) that the RV-GOMEA framework, specifically its GOM operator, scales positively to high-dimensional search spaces. However, we can mitigate the additional overhead by providing sufficient parallelization.

To evaluate this scalability and the impact of computational resources, we propose the following specific hypotheses.

- **Hypothesis 4.1 - Dimensionality Challenge:** We hypothesize that the significant increase in dimensionality (from 4 variables in the previous experiment to 12) will present a substantial challenge for the resource-constrained scenario ($k = 2$), potentially getting trapped in a local optimum.

- **Hypothesis 4.2 - Acceleration via Smoothing:** Based on the landscape analysis, we hypothesize that the "High-Compute" scenario ($k = 16$) will allow the EA to converge in fewer generations and avoid the local optima traps predicted in **H4.1**. Justifying the increased computational resource usage.

4.5.3 Experimental Configuration

Genotype Encoding

The genotype is a 12-dimensional real-valued vector $\mathbf{x}_{E3} = (x_1, \ldots, x_{12})$. Each pair of parameters (x_i, x_{i+6}) defines a complete DVI for one of the six ROIs. The mapping is defined as follows:

$$
\begin{aligned}
D_{x_1\%}^{CTV_{HR}} &> x_7\% \\
D_{x_2\%}^{PTV} &> x_8\% \\
V_{x_3\%}^{Bladder} &< x_9 \text{ cm}^3 \\
V_{x_4\%}^{Rectum} &< x_{10} \text{ cm}^3 \\
V_{x_5\%}^{Sigmoid} &< x_{11} \text{ cm}^3 \\
V_{x_6\%}^{Ring_{NT}} &< x_{12} \text{ cm}^3
\end{aligned}
$$

Computational Scenarios

To address Research Question 4.2, the experiment is executed under two distinct computational configurations:

1. **Scenario A (Standard Compute):** The baseline configuration using 3 GPUs.
    - Averaging: $k = 2$ runs per patient (12 total simulations per fitness evaluation).
    - Rationale: Represents a resource-constrained environment and more representative comparison to previous experiments.
2. **Scenario B (High-Performance Compute):** An accelerated configuration using 48 GPUs.
    - Averaging: $k = 16$* runs per patient (96 total simulations per fitness evaluation).
    - Rationale: Utilizing massive parallelization to reduce fitness noise. Expected to converge in less generations.

* The k value was selected to fully utilize the computational resources provided. We do not necessarily expect a huge difference between $k = 10$ and $k = 16$.

Population Size and Runtime

To determine the size of the population for these experimental runs, we used the multivariate formula 4.3:

$$3 \times 12^{1.5} + 17 \approx 142$$

The experimental runs were set to run for a maximum of 2 months ($\approx 1400h$).

Algorithm Configuration

The evolutionary settings are summarized in Table 4.3.

Table 4.3: Parameter values for the complete protocol experiment. Optimization algorithm (Outer loop). Fitness function (Inner loop, BRIGHT).

Parameter	Value
Outer Loop: RV-GOMEA (Protocol Search)	
Training data	6 patients
Population Size	142
Linkage Model	Described in Section 3.4.1
Termination criteria	Runtime $\leq$ 1400 h
Inner Loop: BRIGHT Configuration (Fitness Evaluations)	
Total runs	12 / 96 (2 or 16 $\times$ 6 patients)

Performance Metrics

To assess the scalability of the proposed framework and the clinical validity of the fully evolved protocols, performance is evaluated across three distinct aspects: optimization process, final protocol quality, and clinical acceptability.

1. **Optimization Process Metrics:** Consistent with the previous experiments, the primary quantitative indicators of success remain the evolutionary process metrics:
    - Generation-mean fitness: The average $F^{(AFS)}$ error of all individuals in the population at each generation.
    - Best-so-far fitness: The lowest $F^{(AFS)}$ error found in the population at or until each generation.
2. **Final Protocol Quality Metrics:**
    - **Quantitative Generalization:** The measure of clinical utility remains the protocol's generalization capability. The best-performing individual $\mathbf{x}^*_{E3}$ is applied to the different patients, and the resulting Pareto fronts and treatment plans are evaluated. The results of the treatment plans generated will be expanded in Appendix 7.
    - Qualitative Metric Analysis: We assess whether the EA autonomously converges close to standard clinical indices.
3. Clinical Expert Assessment: A senior medical physicist from VCU, blinded to the origin of the plan where possible, will perform a detailed review of the treatment plan generated by the optimized protocol $\mathbf{x}^*_{E3}$ for the validation patient.

4.6 Experimental Setup

Finally, mention that the experiments were performed on a Linux server. The hardware comprised a dual-socket configuration with two Intel Xeon Bronze 3206R CPUs, making a total of 16 physical cores available at a base clock speed of 1.90GHz. For hardware acceleration, the server was equipped with three NVIDIA RTX A5000 GPUs, each with 24 GB of dedicated memory. Both the evolutionary search process and the clinical plan evaluations were GPU-accelerated using the CUDA 11.7 toolkit.

The algorithmic framework was developed in Python 3.12. The core evolutionary optimization was driven by our implementation of GOMEA, and the generation and evaluation of candidate plans was done with the latest version of BRIGHT.

Chapter 5

Results

Following the experimental design and methodology detailed in Chapters 3 and 4, this chapter presents the findings derived from the different experiments.

We begin by explaining the thresholds that will act as a point of reference for the following experiments. Next, in Section 5.2 we present the analysis of the problem, observing the impact of stochasticity and evaluating the efficacy of the proposed normalization strategies. After that, in Section 5.3 we address the initial evaluation, assessing the algorithm's ability to optimize the aim values within a fixed protocol structure. Subsequently, Sections 5.4 and 5.5 present the outcomes of the partial and full protocol optimizations, testing the system's capacity to handle increased representational freedom. Finally, we present the insights obtained from discussing the results with a clinical expert.

Throughout this analysis, we revisit the research questions and hypotheses presented on previous chapters. We evaluate whether the obtained data support the hypotheses and discuss the fitness evolution and dose distribution improvements.

5.1 Baseline Protocol Performance

To evaluate the performance of the baseline protocol, we ran BRIGHT 100 times for all the patients in the training set and we reported the mean fitness over the 100 runs, as well as the best fitness value to be used as thresholds to indicate how our algorithm performed during each experiment. The results will be visualized on the plots for the following experiments.

5.2 Search Space Investigation Results

This experiment aims to characterize the fitness landscape of the CTV_{HR} objective, specifically focusing on the trade-off between computational cost (averaging runs k) and landscape smoothness (**H1.1**), as well as the potential for gradient steepening via normalization (**H1.2**).

It is important to note that the plots focus on the clinically relevant subset of the search space. Two regions of the parameter grid were omitted from the analysis for the following reasons: First, the upper-right region of the DVH, representing DVIs that combine high-

volume metrics with high-dose aims, consistently yielded infeasible treatment plans. Second, the lower-left region represented trivial constraints. These ”loose” DVIs exerted little optimization pressure on the planning software.

5.2.1 Impact of Stochastic Averaging

The topology of the fitness search space under varying levels of stochastic averaging ($k = 2$, $k = 10$, and $k = 50$) is presented in Figure 5.1. Each heatmap shows the fitness value $F^{(AFS)}$ across the DVH.

In the $k = 10$ and $k = 50$ plots (Figures 5.1b and 5.1c), a high-fitness region is clearly identifiable near $[94, 90]$ (corresponding to the DVI $D_{90\%}^{CTV_{HR}} < 94\%$). In contrast, in the $k = 2$ scenario, this region is still there but much more difficult to identify.

This confirms that even with low averaging, the global optimum is there. However, the texture of the landscape varies significantly:

- At $k = 2$ (Figure 5.1a): The landscape shows substantial noise. This creates numerous local optima, which poses a risk that the evolutionary algorithm will become trapped or oscillate between them.

- At $k = 10$ (Figure 5.1b): The landscape becomes remarkably smoother. The noise is mostly averaged out, revealing a clearer gradient.

- At $k = 50$ (Figure 5.1c): The landscape becomes slightly smoother than the previous scenario, demonstrating that additional computation efforts help, but the improvements get significantly lower.

While increasing k results in a cleaner signal, the computational cost scales linearly with k. The transition from $k = 2$ to $k = 50$ represents a 25-fold increase in runtime. The results suggest that, while $k = 2$ is noisy, it might still be good enough. Therefore, considering the already expensive fitness function, this is the averaging that will be used in Section 5.3 and Section 5.4.

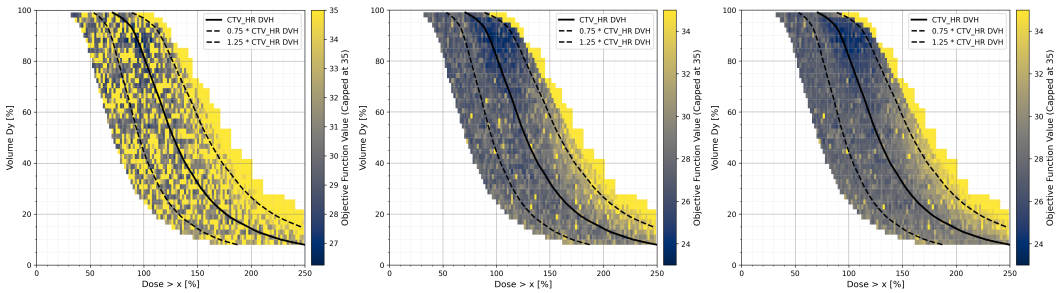

(a) $k = 2$ (b) $k = 10$ (c) $k = 50$

Figure 5.1: Heatmap representation of the fitness space for the DVH possible values for the CTV_{HR} with different levels of averaging.

5.2.2 Impact of Fitness Normalization

The effect of the normalization factor on the search gradient is depicted in Figure 5.2, which compares the non-normalized fitness landscape against the normalized version (using $k = 10$ for both to ensure clarity).

The non-normalized scenario shows a relatively flat plateau in regions far from the region surrounding the baseline protocol DVI (e.g., in very low-volume settings). In contrast, the normalized landscape shows a steeper gradient, particularly pronounced at the bottom-right of the DVH.

These results seem to support **H1.2**. The normalization factor effectively penalizes clinically irrelevant regions (where x_{vol} or x_{dose} are near zero), artificially creating a stronger selection pressure that guides the population toward the region we want to explore without hard-coded constraints. This steepening should accelerate the early phases of the evolutionary search. The quantitative validation of this convergence speedup is presented in the optimization results of the partial optimization experiment (see Section 5.4).

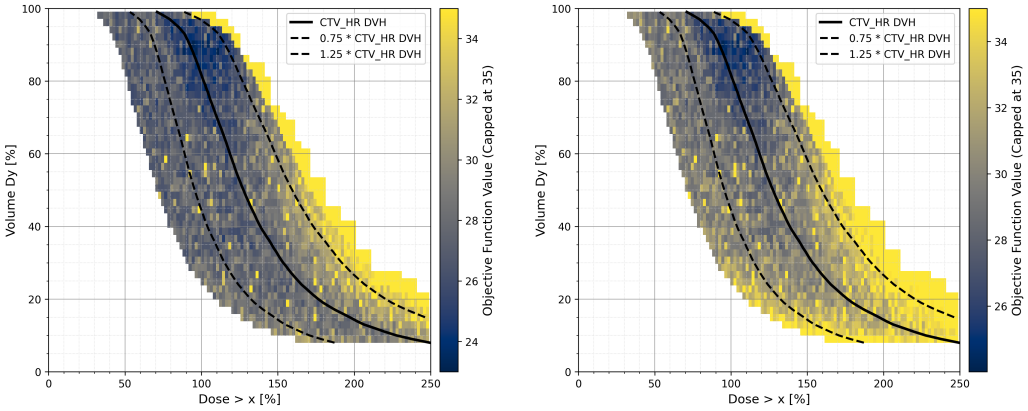

(a) without normalization factor

(b) with normalization factor

Figure 5.2: Heatmap representation of the fitness space for the DVH possible values for the CTV_{HR} with and without normalized fitness values.

Answering **RQ1.1**, the increased averaging works, as hypothesized, well with the problem, and it will be noticeable in Section 5.5. **RQ1.2** will be discussed in the following section.

5.3 Initial Evaluation Results

The optimization performance is visualized in Figure 5.3, which tracks the fitness improvement of the objective function $F^{(AFS)}$ across three independent runs of the algorithm. The results report the elite solution found in each generation.

Notably, even within the first generation, all three runs identified at least one protocol that resulted in a front of treatment plans superior to the baseline. Subsequent generations demonstrate the efficacy of the selection mechanism, which produced incrementally better plans. Despite the limited number of generations, the algorithm successfully navigated the search space, thereby validating **H1.1**.

Shifting our focus now to the best protocol obtained. Even if the optimization might still have room for improvement, as there are no signs of convergence, observing plans generated with this protocol should already show more similarity.

The protocol DVIs are the following:

$$
\begin{aligned}
D_{90\%}^{CTV_{HR}} &> 110.9\% \\
D_{90\%}^{PTV} &> 94.5\% \\
V_{65\%}^{Bladder} &< 2.02 \text{ cm}^3 \\
V_{65\%}^{Rectum} &< 2.75 \text{ cm}^3 \\
V_{65\%}^{Sigmoid} &< 2.74 \text{ cm}^3 \\
V_{90\%}^{Ring_{NT}} &< 3.04 \text{ cm}^3
\end{aligned}
$$

The resulting aims look more restrictive than the baseline protocol (Table 3.2). However, this is not a problem if it allows BRIGHT to generate plans with the same coverage and a more similar dose distribution. As we mentioned in **H2.2**, the resulting aims look reasonably similar to what the clinical experts would define.

To assess the impact of these parameters, we examined the dose distribution similarity for one of the patients. It is important to note that because the optimization minimizes an aggregate loss across the entire patient set, improvements may not be present for every patient. However, given the improved fitness score, some gains in plan similarity are expected for a subset of the patients.

Figure 5.4 presents a qualitative comparison for one of the patients. The figure displays the clinically approved plan (left) alongside the optimized plan (right), both plans selected with the same $D_{90}^{CTV_{HR}}$ coverage. The central panel illustrates the radiation differences, where green denotes areas of increased irradiation in the optimized plan and purple denotes decreased irradiation. Consequently, regions with lower color saturation indicate higher similarity in dose distribution.

Inspection of these comparison plots confirms that the optimized protocol produces treatment plans more similar to the clinically approved plan.

Note: For the lateral views, ROIs are delineated as follows: CTV_{HR} (blue), PTV (red), Bladder (yellow), Rectum (brown), and Sigmoid (green). Isodose lines correspond to 50%

(green), 75% (yellow), 100% (red), and 150% (blue) of the planned dose.

Further exploration of individual treatment plans will be done in Section 5.5.

Answering **RQ2.1** and **RQ2.2**, the algorithm optimized the structured protocol as we expected, and the obtained aims are of clinical relevance and easily interpretable by clinical experts.

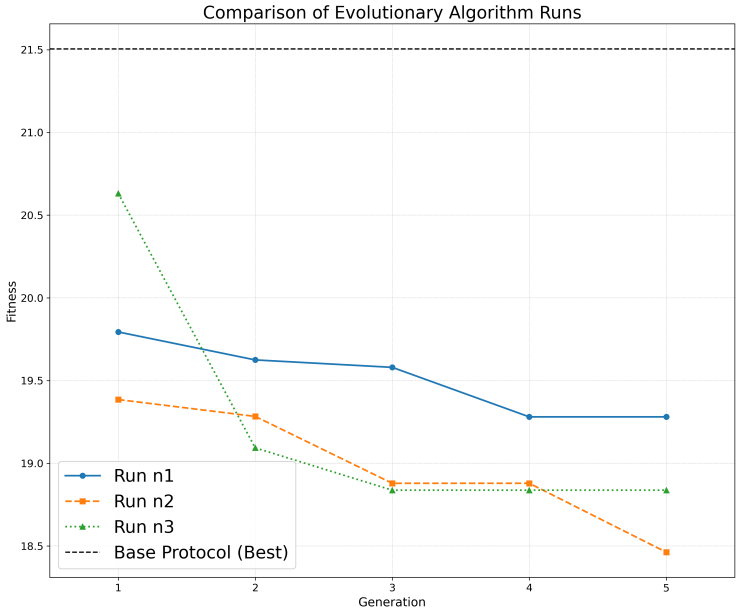

Figure 5.3: Optimization of Aim values using the same protocol structure as the baseline protocol. Fitness values of the best-so-far through generations of 3 experimental runs. (lower is better)

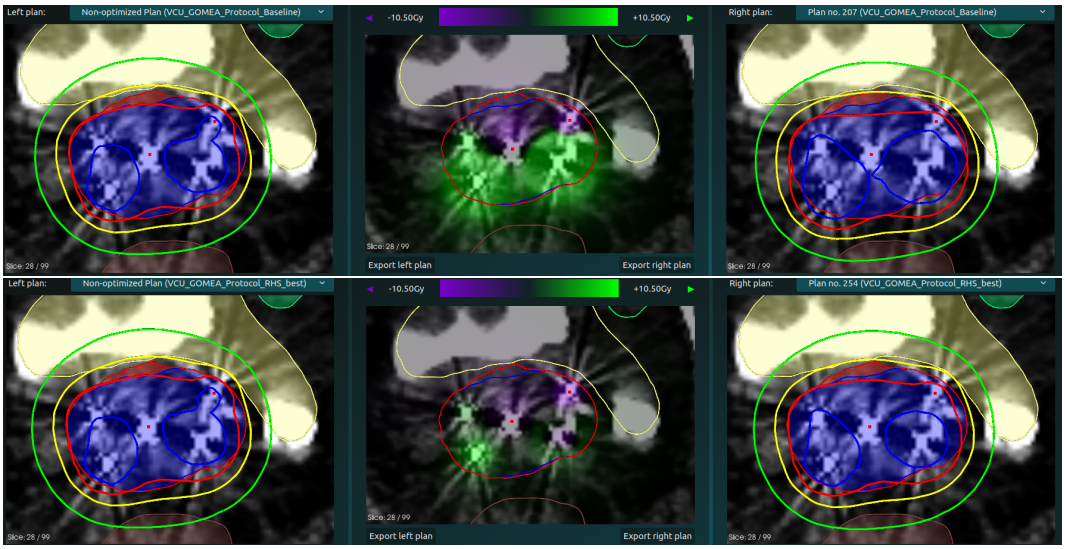

(a) Patient 6, CT slice 28.

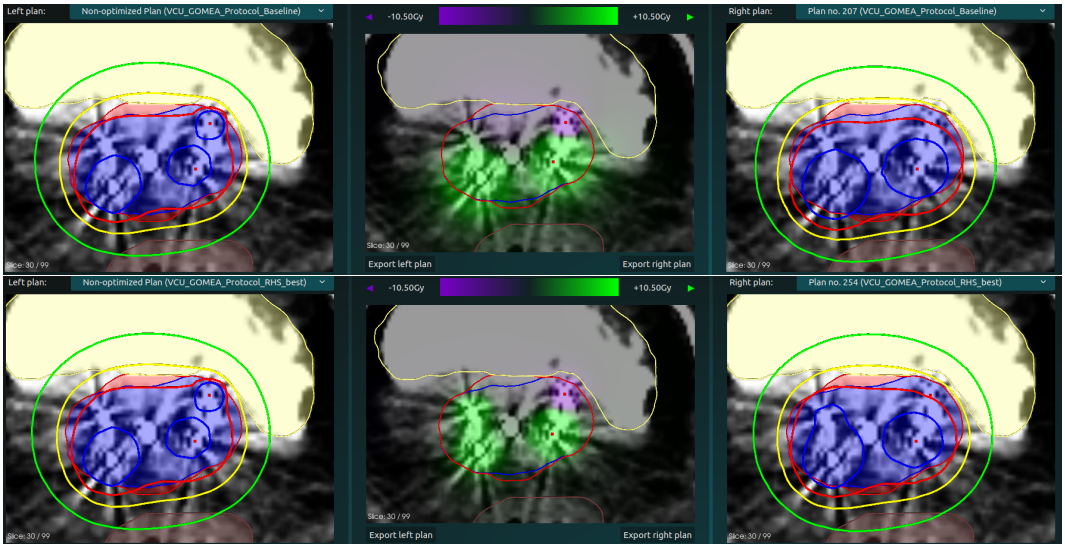

(b) Patient 6, CT slice 30.

Figure 5.4: Dose distribution comparison of two treatment plans for the same patient. Clinical plan (left), Optimized plan (right), Radiation comparison (middle). Baseline protocol (top), Protocol with the lowest error (bottom). (Less color in the middle plot is better)

5.4 Partial Optimization Results

The results for the single-DVI protocol optimizations are summarized in Figure 5.5. These plots illustrate the evolution of the single DVI experiments (CTV_{HR} and PTV), comparing performance under both normalized and non-normalized fitness formulations.

Figure 5.5a shows the average population fitness at each generation. These findings contradict the expectations outlined in Section 5.2.2. Contrary to **H1.2**, the runs utilizing the normalized fitness function failed to demonstrate a faster convergence rate compared to the non-normalized runs. Consequently, answering **RQ1.2**, the normalization factor was deemed ineffective for this specific problem. Based on these observations, the non-normalized fitness function was selected for subsequent experiments.

A detailed examination of the convergence trajectories for each population in each scenario (provided in the Appendix 7) reveals that all execution runs converged to a region proximal to the baseline protocol values. This is supported by the DVIs presented in Section 5.4.1, summarized in Table 5.1. The DVIs resemble the ones in the baseline protocol. Therefore, we can conclude that the resulting DVIs represent clinically intuitive aims consistent with standard practice. This indicates that the optimization successfully avoided convergence toward proportional yet clinically irrelevant DVIs, regardless of whether the normalization factor was used.

Regarding the elite solutions discovered presented in Figure 5.5b, improvements over the baseline protocol are observable but marginal compared to the gains achieved in the previous experiment. This reduced performance is likely due to the more limited search space. By limiting the protocol to explore DVI candidates for only a single volume, the optimization is unable to balance competing objectives. Specifically, the aim for one target volume may be dominated by the fixed DVI of the other.

This is supported by the results in Figure 5.6. When the search space is expanded to include both coverage DVIs ($D^{CTV_{HR}}$ and D^{PTV}), the optimization yields a solution approaching the fitness levels comparable to the optimization in Section 5.3.

5.4.1 Analysis of the Generated Treatment Plans

From all the individuals explored, it is hard to tell which one is the best one. This is because, even if the huge number of DCPs is a good estimate, when comparing thousands of individuals, the chosen seed can have a big impact on the evaluation of a solution. However, for practical reasons we can be certain that the protocol that obtained the best fitness is at least within the subset of solutions that are feasible, and close to the best.

First, let's take a look at the DVI values generated, which are shown in Table 5.1. We can see the values for CTV_{HR} and PTV come quite close to the baseline protocol.

Now, looking at the protocol that achieved the best fitness:

$$D^{CTV_{HR}}_{92.8\%} > 105.6\%$$
$$D^{PTV}_{92.7\%} > 86.12\%$$

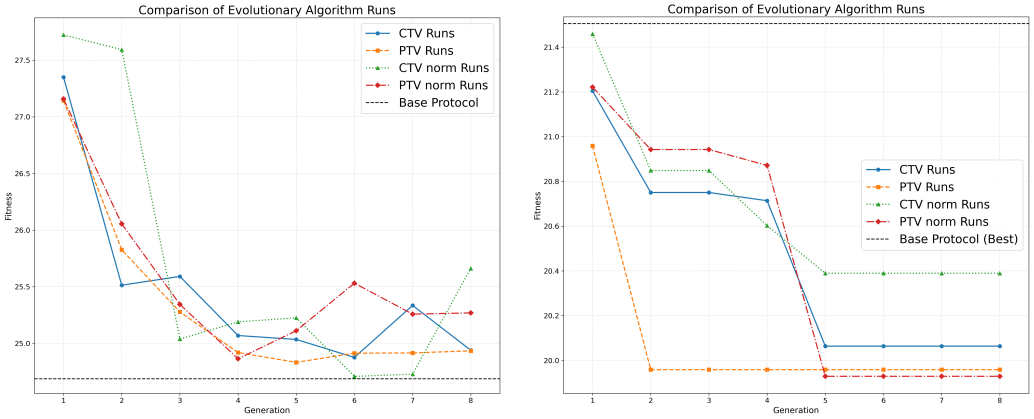

(a) Population's average fitness at each generation (b) Best-so-far individual at each generation

Figure 5.5: Optimization of a single DVI (CTV_{HR} or PTV) using $F^{(AFS)}$ or $F^{(AFS)}_{norm}$ (lower is better)

Table 5.1: Summary of the obtained DVIs across 6 runs of each scenario.

Experiment	**Best DVIs**
CTV_{HR} without normalization	$D^{CTV_{HR}}$ $89.12\% \pm 2.13\% > 100.37\% \pm 6.87\%$
CTV_{HR} with normalization	$D^{CTV_{HR}}$ $90.27\% \pm 3.65\% > 100.35\% \pm 6.10\%$
PTV without normalization	D^{PTV} $87.26\% \pm 10.06\% > 81.43\% \pm 9.56\%$
PTV with normalization	D^{PTV} $88.75\% \pm 6.80\% > 82.69\% \pm 4.66\%$
$CTV_{HR} + PTV$	$D^{CTV_{HR}}$ $91.70\% \pm 1.13\% > 108.21\% \pm 3.50\%$
	D^{PTV} $90.38\% \pm 2.43\% > 93.06\% \pm 6.94\%$
Baseline Protocol DVIs	$D^{CTV_{HR}}$ $90\% > 94\%$
	D^{PTV} $90\% > 80\%$

The target aims are more restrictive than those of the baseline protocol. Regardless, BRIGHT manages to generate treatment plans with the same $D_{90}^{CTV_{HR}}$ that we will compare for one of the patients. Additionally, looking at the generated Pareto approximation front for one of the patients in Figure 5.8, we can see how the algorithm tries to place the front as close to the clinical plan as possible and it never falling on the extremes, thus giving more margin for the clinical experts to go higher or lower on coverage, just focusing on maintaining the trade-off.

Figure 5.8 shows a comparison of dose distributions between the baseline protocol and a protocol with modified target aims corresponding to those last presented. What we observed through multiple CT scan slices is that the difference with the baseline protocol is very hard to notice. This might be due to the sparing aims being more forgiving that the intent during planning of the clinical treatment plan. A more comprehensive analysis of dose distributions will be done on the next experiment where the sparing aims will be a lot more restrictive.

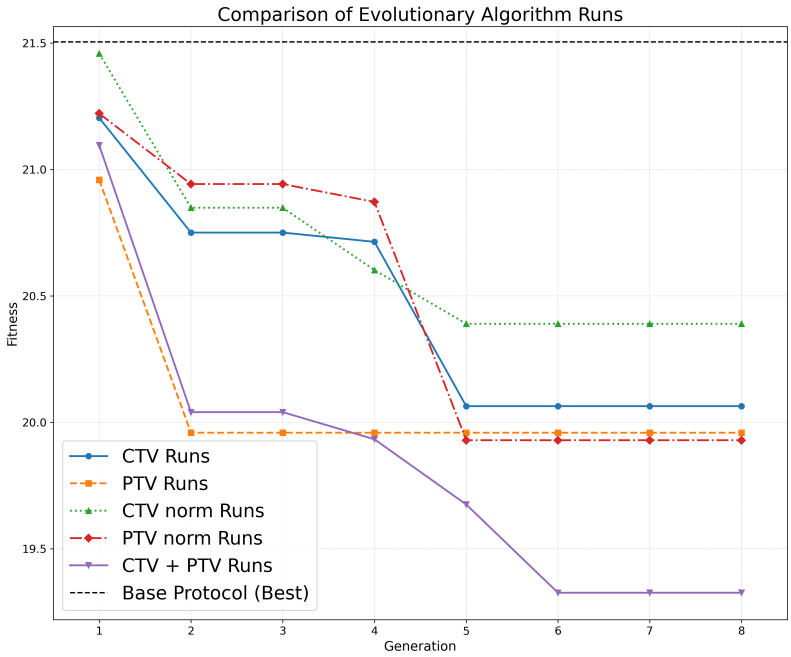

Figure 5.6: Optimization of two DVIs ($CTV_{HR} + PTV$) using $F^{(AFS)}$ vs. single-DIV optimizations. Best-so-far individual at each generation (lower is better)

Addressing the aforementioned hypotheses, **H3.1** seems correct since the average population fitness struggles more to reach the threshold level than in the previous experiment. However, this is not an issue since the elite solutions are of more relevance for our research. About **H3.2**, it was surprisingly incorrect. Observing the Table 5.1, it can be seen that the combined experiment of $D^{CTV_{HR}}$ and D^{PTV} had lower variance than the single-DVI experiments (Figure 5.6). We attribute that result to the low sample size of experimental runs. Finally, Table 5.1 confirms **H3.3**, resulting in a positive answer to the **RQ3**. The resulting DVIs did not fall into irrelevant proportional values, but instead were very similar to the structured protocol ones.

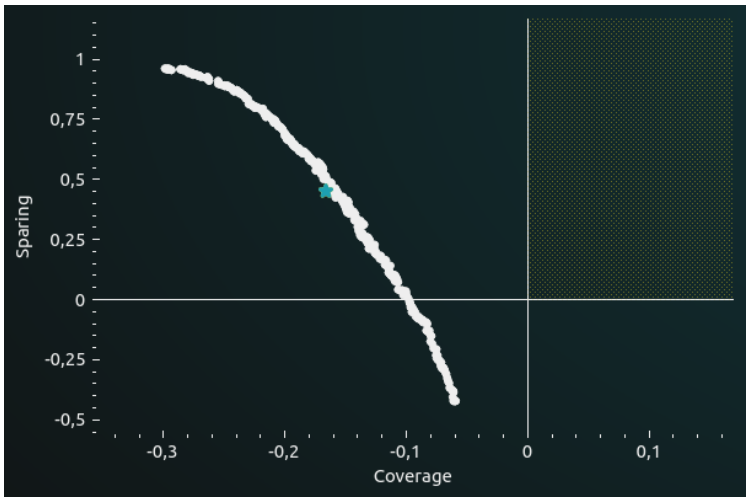

Figure 5.7: Pareto approximation front for one of the patients generated by BRIGHT using the best protocol optimizing $CTV_{HR} + PTV$. The blue star represents the clinical treatment plan.

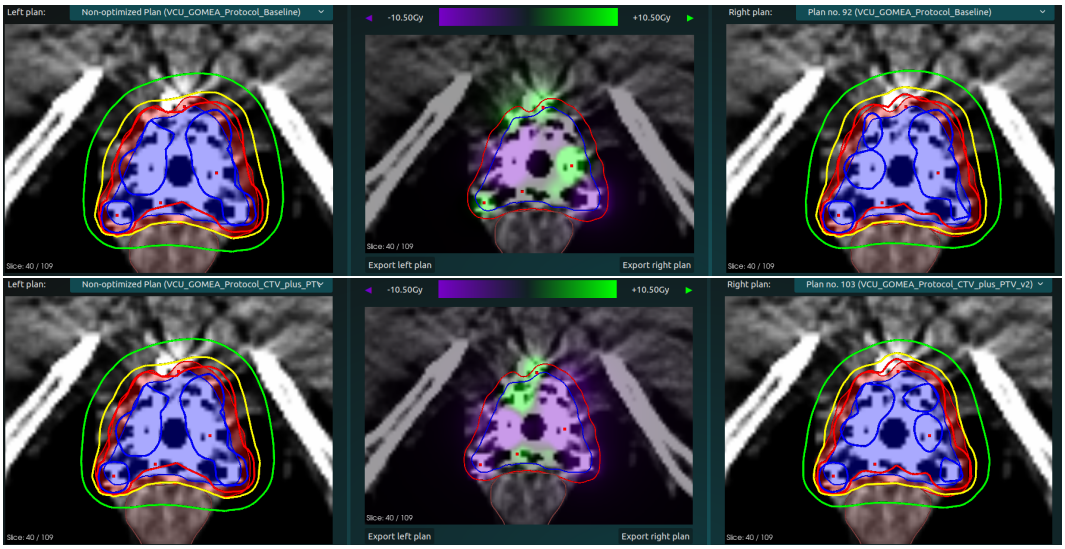

Figure 5.8: Patient 3, CT slice 40. Clinical plan (left), Optimized plan (right), Radiation comparison (middle). Baseline protocol (top), $CTV_{HR} + PTV$ protocol with the lowest error (bottom).

5.5 Complete Protocol Optimization

This experiment extends the optimization to include both metric definitions and aim values for all six DVIs simultaneously. The results are analyzed to evaluate the scalability of the framework (RQ 4.1) and the impact of parallelization on search efficiency (RQ 4.2).

The evolutionary trajectories for both the resource-constrained ($k = 2$) and parallelized ($k = 16$) scenarios are visualized in Figure 5.9.

Figure 5.9a reveals critical insights regarding the optimization process. Initially, the comprehensive optimization starts with a worse average fitness compared to the single-DVI experiments. This is expected, as the algorithm initializes from a random sampling of a significantly broader search space (12 dimensions), whereas the previous experiment optimized a single parameter while holding the rest of the protocol fixed. However, within seven generations, the population achieves an average fitness comparable to the single-DVI results, surpassing the elite individual of the previous experiment in less than six generations.

The expansion to 12 dimensions exponentially increases the volume of the search space. Consequently, the limited number of generations permitted by the computational budget was insufficient to concentrate the majority of the population into the region of better fitness than the baseline threshold. The continued downward trend of the average fitness curve at the point of termination suggests that the algorithm had not yet converged. Instead, the experiment was terminated while the search was still active due to the time constraints. Supporting **H4.1**.

The trajectory of the elite individuals, illustrated in Figure 5.9b, demonstrates the framework's efficacy. In both scenarios, the algorithm found protocol configurations that outperformed both the clinical baseline and the single-DVI results. This confirms that even within a high-dimensional, noisy search space, the automated framework is capable of finding clinically superior protocols. Quantitatively, the resulting plans show greater similarity to the clinical ground truth than those generated by the manually tuned baseline.

5.5.1 Impact of Computational Scaling ($k = 2$ vs. $k = 16$)

To address **RQ4.2** and **H4.2**, we compare the performance of the resource-constrained scenario ($k = 2$) with that of the massively parallelized scenario ($k = 16$).

As observed in Figure 5.9, the $k = 16$ scenario achieved superior fitness values in fewer generations compared to the $k = 2$ scenario. The increased averaging factor significantly smoothed the fitness landscape (as predicted in Section 5.2.1).

However, it is critical to acknowledge the trade-off between generational efficiency and computational cost. While the $k = 16$ runs converged faster in terms of generations, the computational cost per evaluation was eight times higher. Consequently, on fixed hardware, the total wall-clock time required to reach a specific fitness threshold would not necessarily decrease. Nevertheless, this result demonstrates the potential of parallelization: given sufficient hardware resources (e.g., a larger GPU cluster), increasing the stochastic averaging factor (k) provides a reliable pathway to accelerate convergence in complex, high-dimensional protocol optimizations.

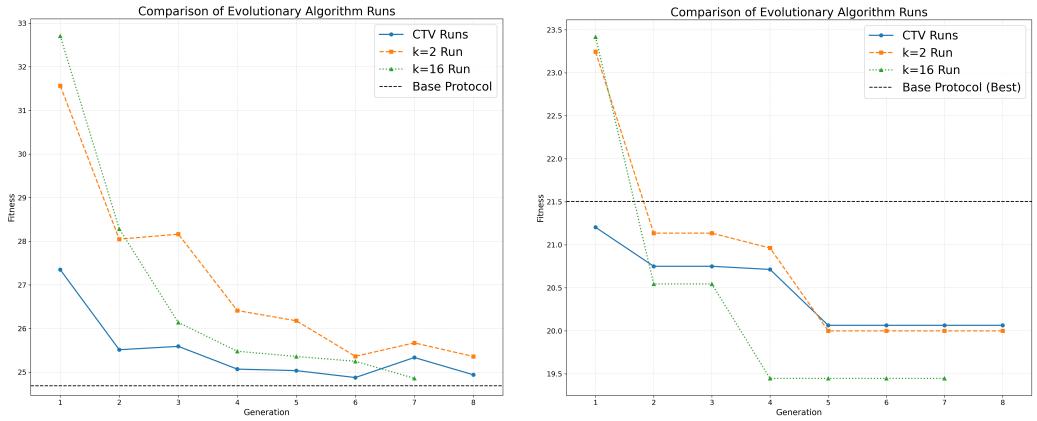

(a) Population's average fitness at each generation (b) Best-so-far individual at each generation

Figure 5.9: Optimization of the complete protocol for $k = 2$ and $k = 16$ (lower is better)

5.5.2 Analysis of the Generated Treatment Plans ($k = 16$)

Figure 5.10 shows the qualitative improvement in dose distribution through the generations. As protocols with better fitness are selected, the deviation from the clinical plan diminishes. By generation 6, the dose difference map displays significantly reduced color saturation, indicating a much higher similarity between the optimized and clinical doses at each sampled point. This suggests that the evolved protocol captures the complex trade-offs inherent in clinical practice more effectively than the baseline.

Visualizing these trade-offs across the entire geometry of the patient is challenging in a single summary figure. A comprehensive set of representative CT slices for all patients is provided in the Appendix 7.

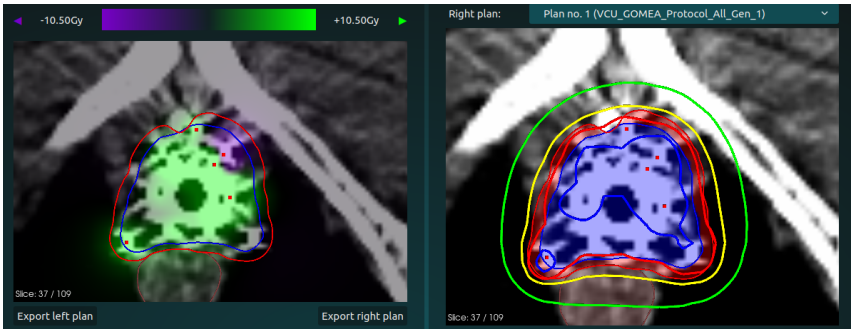

(a) Generation 1

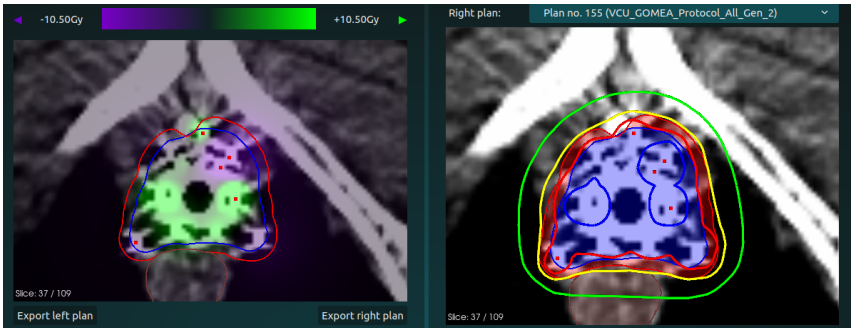

(b) Generation 2

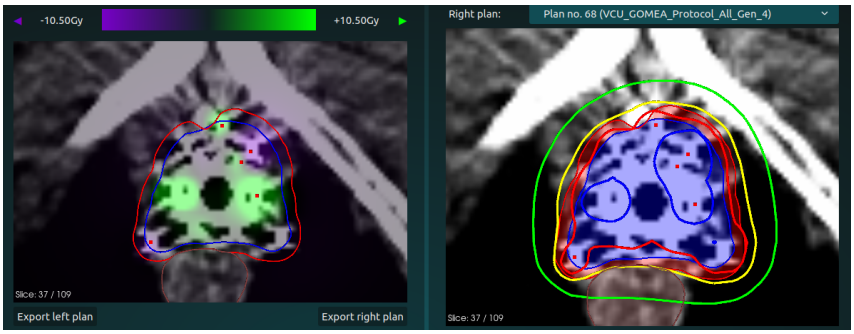

(c) Generation 4

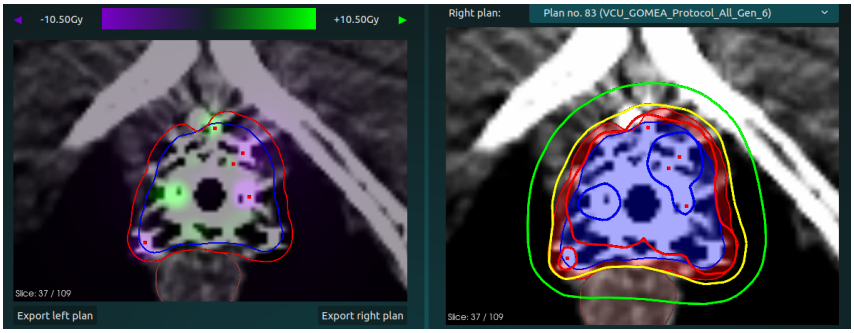

(d) Generation 6

Figure 5.10: Dose distribution and difference in radiation to the clinical plan for the best protocol at each generation for one patient. (Patient 3, CT Slice 37) (Less color on the left plot is better)

5.6 Automated Protocol Generation Results - Clinical Assessment

To validate the clinical utility of the automated protocol discovery framework, a qualitative assessment was conducted involving the aforementioned medical physicist. This assessment focused on the treatment plans generated by the optimized protocol $\mathbf{x}^*_{E3}$ for the validation patient.

The expert reviewed the Pareto approximation front generated by BRIGHT. The assessment focused on three key dimensions: clinical validity, the significance of dosimetric variance, and the operational utility of the system.

5.6.1 Clinical Validity and Generalization

The primary objective of this research was to generate a protocol capable of producing plans indistinguishable from the local standard of care. The expert review confirmed that the plans generated for the validation patient were clinically acceptable. Specifically, the expert noted that plans exhibiting coverage levels similar to the clinical ground truth could have been approved for patient treatment without modification.

Furthermore, the expert evaluated multiple treatment plans from the approximately 200 plans on the Pareto front. The assessment indicated that all generated plans fell within the range of validity for that specific patient geometry. This suggests that the discovered protocol $\mathbf{x}^*_{E3}$ successfully achieved our goal and further confirms **H4.2** regarding the system's ability to generalize to unseen anatomies.

5.6.2 Clinical Relevance of Dosimetric Variance

A quantitative view of the generated front showed a variance of approximately 10% in target coverage between the extreme ends of the Pareto front. While this mathematical variance might look significant, the clinical expert emphasized that it must be contextualized within the physical realities of brachytherapy delivery.

The discrepancy between the planned dose distribution (based on the planning CT) and the delivered dose is often dominated by organ motion during the time between them. Consequently, minor deviations in generated plans are frequently smaller than the physical uncertainties inherent in the treatment delivery. This insight suggests that the automated protocol is sufficiently robust (it achieves a level of precision where further mathematical optimization doesn't yield any clinical benefits).

5.6.3 Operational Utility and Limitations

The assessment highlighted a dichotomy between the system's efficiency and its usability. On the positive side, the expert highlighted the system's ability to rapidly generate high-quality plans that align with the oncologist's preferences. This directly addresses the "cold start" problem described in Chapter 1, demonstrating that the framework can significantly accelerate the configuration phase for new institutions.

However, the assessment also revealed a limitation regarding the output format. While BRIGHT provides a comprehensive set of trade-off solutions, the expert noted that presenting 200-300 options can lead to decision paralysis. In a time-constrained clinical environment, clinical experts often prefer a smaller selection of ”best” options rather than a vast Pareto front.

Finally, the expert remarked that because the ”optimal” dose for curing a specific patient while minimizing toxicity is not fully known, the system focuses on replicating historical practices ”what the doctor usually does” rather than discovering a biologically superior plan. While the proposed framework succeeds in its goal of automating the replication of clinical protocols, it does not necessarily challenge or improve upon the existing standard of care.

Chapter 6

Discussion and Future Work

The experimental results presented in Chapter 5 demonstrated that the proposed bi-level optimization framework is capable of autonomously discovering clinical protocols that replicate expert planning styles. Specifically, the system successfully evolved parameters that generated treatment plans similar to the clinical ground truth. However, translating a proof-of-concept into a robust clinical tool requires a critical examination of the system's current boundaries. This chapter interprets these findings within the context of the study's constraints. Section 6.1 discusses the limitations encountered and Section 6.2 outlines potential avenues for future research.

6.1 Discussion

Computational Constraints vs. Stochastic Noise

A big challenge of this research was the prohibitive computational cost due to the expensive fitness function. Because BRIGHT is a stochastic planning system, minimal noise evaluations require averaging over many runs. However, hardware limitations restricted most experiments to $k = 2$. Although RV-GOMEA was able to navigate this noisy landscape, the presence of stochastic noise likely hindered convergence speed and the precision of final solutions. Additionally, for the results to be significant, the experiments would have to be repeated more times and run for more generations than it was feasible within the context of this thesis.

Dataset Size and Overfitting

The optimization process was based on a relatively small training set of six patients. Although validation on the held-out patient indicated promising generalization, a larger and more diverse dataset is necessary. We couldn't even test strategies like k-folds, great for small datasets, since they were unfeasible within the constraints of this study. Therefore, there remains a risk that the discovered protocol is partially overfit to specific anatomical geometries.

Inter-Observer Variability and Ground Truth

Another challenge in this research is the reliance on a single clinical plan per patient as the ”ground truth”. In clinical practice, treatment planning is partly subjective; two expert clinicians within the same institution may produce distinguishable plans for the same patient, both of which are clinically acceptable. Optimizing the protocol to minimize the distance to a specific historical plan introduces the risk of overfitting to an individual doctor's preferences or to particular biases and minor inconsistencies present in that case.

A valid concern is that the algorithm may learn and propagate individual ”faults” rather than capturing the robust standard of care of the institution. However, since this optimization would not be the final setup of the automated TPS but just one step of the process, this would not be a big concern.

6.2 Future work

To advance this proof-of-concept into a usable tool, future research should focus on increasing the robustness of the fitness evaluation, studying the flexibility of the protocol representation, and increasing the computational scale.

Evolution of Protocol Structure via Genetic Programming

The current approach assumes a fixed number of objectives, with one per region of interest. A more advanced methodology would employ Genetic Programming (GP) to evolve the protocol's structure. Rather than optimizing parameters for a fixed set of six aims, a GP-based approach could dynamically add or remove aims, potentially revealing that a specific target requires multiple distinct DVIs for adequate coverage. This would transition the system from parameter tuning to genuine protocol architecture discovery.

Scalability via High-Performance Computing

Utilizing a larger GPU cluster would enable higher stochastic averaging ($k \geq 50$) without increasing wall-clock time. This would smooth the fitness landscape and likely allow the algorithm to converge in a few generations with greater precision.

Multi-Institutional Validation

Finally, testing the framework on datasets from multiple institutions with distinct planning styles would validate its ability to learn these local treatment styles. This would confirm its utility as a general-purpose tool for accelerating the deployment of automated planning systems to more hospitals.

Chapter 7

Conclusion

This thesis tested the feasibility of automating the discovery of clinical protocols for HDR brachytherapy treatment planning. By formulating the protocol configuration as a bi-level optimization problem and solving it using RV-GOMEA, this research aimed to streamline the onboarding of the BRIGHT treatment planning system at new medical institutions, reducing the manual labor required to introduce this technology.

We performed an analysis of the protocol discovery problem by progressively increasing the complexity of the search space. To make it feasible, a set of pruning strategies was implemented and validated to reduce the computational cost by reducing the number of fitness evaluations with minimal hard-coded rules. This approach ensures that the automated discovery process remains doctor-specific and data-centric, avoiding bias from preconceived notions of standard protocol values.

Based on the experimental results, the RV-GOMEA-based framework successfully identified protocol parameters that generated treatment plans quantitatively similar to, and subjectively as valid as, expert clinical plans. This confirms the fundamental hypothesis that implicit expert knowledge can be extracted and encoded into a clinical protocol using EAs. Additionally, we demonstrated that increasing the genotype flexibility by optimizing metric definitions alongside aim values yields superior results, finding protocols that allowed BRIGHT to generate treatment plans more similar to the ground truth compared to optimizing aim values alone.

Finally, this thesis addressed clinical applicability through collaboration with a domain expert. The generated protocols and resulting treatment plans were evaluated in consultation with an expert medical physicist. This qualitative feedback was useful for validating that the optimized protocols translated into clinically acceptable treatment plans, ensuring the research remained grounded in clinical practice.

Collectively, these contributions provide a robust proof-of-concept for the automated discovery of usable clinical protocols for the BRIGHT treatment planning system and answer the main research question. The evolutionary approach demonstrated potential to ease the initial involvement required for manual parameter tuning, increase clinician confidence by linking automated plans to historical expert data, and accelerate the adoption of automated planning tools in new clinical environments.

Acknowledgments

I extend my sincere thanks to Prof. Dr. Aurel Dorin Todor, Medical Physicist at Virginia Commonwealth University. I am grateful for his guidance during the thesis, which allowed me to deepen my knowledge of brachytherapy, and I appreciate his expert validation of the treatment plans. Most importantly, I thank him for providing the clinical patient data, without which this research would not have been possible.

Bibliography

[1] F. Bray, M. Laversanne, H. Sung, J. Ferlay, R. L. Siegel, I. Soerjomataram, and A. Jemal, "Global cancer statistics 2022: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries," *CA: A Cancer Journal for Clinicians*, vol. 74, pp. 229–263, December 2024.

[2] F. Guida, R. Kidman, J. Ferlay, J. Schüz, I. Soerjomataram, B. Kithaka, O. Ginsburg, R. B. M. Vega, M. Galukande, G. Parham, S. Vaccarella, K. Canfell, A. M. Ilbawi, B. O. Anderson, F. Bray, I. dos Santos-Silva, and V. McCormack, "Global and regional estimates of orphans attributed to maternal cancer mortality in 2020," *Nature Medicine*, vol. 28, p. 2563–2572, November 2022.

[3] C. Chargari, K. Peignaux, A. Escande, S. Renard, C. Lafond, A. Petit, D. L. C. Kee, C. Durdux, and C. Haie-Méder, "Radiotherapy of cervical cancer," *Cancer/Radiothérapie*, vol. 26, pp. 298–308, Feb–Apr 2022.

[4] Y. Qiao, H. Li, and B. Peng, "Neoadjuvant and adjuvant treatments compared to concurrent chemoradiotherapy for patients with locally advanced cervical cancer: A bayesian network meta-analysis," *Frontiers in Oncology*, vol. 12, Mar 2022. doi: 10.3389/fonc.2022.745522. PMID: 35372073; PMCID: PMC8966774.

[5] A. N. Viswanathan, S. Beriwal, J. F. D. L. Santos, D. J. Demanes, D. Gaffney, J. Hansen, E. Jones, C. Kirisits, B. Thomadsen, and B. Erickson, "American brachytherapy society consensus guidelines for locally advanced carcinoma of the cervix. part ii: High-dose-rate brachytherapy," *Brachytherapy*, vol. 11, pp. 47–52, January-February 2012.

[6] E. A. Suzumura, L. M. Gama, B. Jahn, A. G. Campolina, H. de Andrade Carvalho, and P. C. de Soárez, "Effects of 3d image-guided brachytherapy compared to 2d conventional brachytherapy on clinical outcomes in patients with cervical cancer: A systematic review and meta-analyses," *Brachytherapy*, vol. 20, pp. 710–737, July-Aug 2021.

[7] K. Derks, J. L. Steenhuijsen, H. A. van den Berg, S. Houterman, J. Cnossen, P. van Haaren, and K. D. Jaeger, "Impact of brachytherapy technique (2d versus 3d) on outcome following radiotherapy of cervical cancer," *Journal of Contemporary Brachytherapy*, vol. 10, pp. 17–25, 2018.

[8] S. Aouadi, A. S. khalid, S. Paul, S. Chandramouli, T. Torfeh, S. Yoganathan, S. Paloor, R. Hammoud, and N. Al-Hammadi, "Assessment of organ motion impact on brachytherapy treatment planning of cervix cancer," *International Journal of Clinical Images and Medical Reviews*, vol. 4, p. 1159, July 2024.

[9] L. R. Dickhoff, R. J. Scholman, D. L. Barten, E. M. Kerkhof, J. J. Roorda, L. A. Velema, L. J. Stalpers, B. R. Pieters, P. A. Bosman, and T. Alderliesten, "Keeping your best options open with ai-based treatment planning in prostate and cervix brachytherapy," *Brachytherapy*, vol. 23, no. 2, pp. 188–198, 2024.

[10] R. Pötter, K. Tanderup, C. Kirisits, A. de Leeuw, K. Kirchheiner, R. Nout, L. T. Tan, C. Haie-Meder, U. Mahantshetty, B. Segedin, P. Hoskin, K. Bruheim, B. Rai, F. Huang, E. V. Limbergen, M. Schmid, N. Nesvacil, A. Sturdza, L. Fokdal, N. B. K. Jensen, D. Georg, M. Assenholt, Y. Seppenwoolde, C. Nomden, I. Fortin, S. Chopra, U. van der Heide, T. Rumpold, J. C. Lindegaard, I. Jürgenliemk-Schulz, and E. C. Group, "The embrace ii study: The outcome and prospect of two decades of evolution within the gec-estro gyn working group and the embrace studies," *Clinical and Translational Radiation Oncology*, vol. 9, pp. 48–60, 2018.

[11] A. Elmali, F. Biltekin, S. Y. Sari, M. Gultekin, D. Yuce, and F. Yildiz, "Inter-observer variation of target volume delineation for ct-guided cervical cancer brachytherapy," *Journal of Contemporary Brachytherapy*, vol. 15, no. 4, pp. 253–260, 2023.

[12] Y. Otani, T. Ohno, K. Ando, K. Murata, S. Kato, S. ei Noda, K. Murofushi, H. Ushijima, D. Yoshida, N. Okonogi, F. Isohashi, M. Wakatsuki, and T. Nakano, "Dosimetric feasibility of computed tomography-based image-guided brachytherapy in locally advanced cervical cancer: a japanese prospective multi-institutional study," *Journal of Radiation Research*, vol. 62, no. 3, pp. 502–510, 2021.

[13] K. Kisling, T. D. Keiper, D. Branco, G. G.-Y. Kim, K. L. Moore, and X. Ray, "Clinical commissioning of an adaptive radiotherapy platform: Results and recommendations," *Journal of Applied Clinical Medical Physics*, vol. 23, December 2022.

[14] J. P. Tol, A. R. Delaney, M. Dahele, B. J. Slotman, and W. F. Verbakel, "Evaluation of a knowledge-based planning solution for head and neck cancer," *International Journal of Radiation Oncology*Biology*Physics*, vol. 91, no. 3, pp. 612–620, 2015.

[15] Y. Ge and Q. J. Wu, "Knowledge-based planning for intensity-modulated radiation therapy: A review of data-driven approaches," *Medical Physics*, vol. 46, no. 6, pp. 2760–2775, 2019.

[16] B. Wu, M. Kusters, M. Kunze-busch, T. Dijkema, T. McNutt, G. Sanguineti, K. Bzdusek, A. Dritschilo, and D. Pang, "Cross-institutional knowledge-based planning (kbp) implementation and its performance comparison to auto-planning engine (ape)," *Radiotherapy and Oncology*, vol. 123, no. 1, pp. 57–62, 2017.

[17] E. Podgorsak, *Radiation Oncology Physics: A Handbook for Teachers and Students*. International Atomic Energy Agency (IAEA), 2005.

[18] F. M. Khan, *The Physics of Radiation Therapy (5th ed)*. Lippincott Williams Wilkins, 2014.

[19] T. Ohno, M. Wakatsuki, T. Toita, Y. Kaneyasu, K. Yoshida, S. Kato, N. Ii, S. Tokumaru, H. Ikushima, T. Uno, S.-E. Noda, T. Kazumoto, Y. Harima, and the Working Group of the Gynecological Tumor Committee of the Japanese Radiation Oncology Study Group (JROSG), "Recommendations for high-risk clinical target volume definition with computed tomography for three-dimensional image-guided brachytherapy in cervical cancer patients," *Journal of Radiation Research*, vol. 58, p. 341–350, May 2017.

[20] A. W. M. Sharfo, P. W. Voet, S. Breedveld, J. W. M. Mens, M. S. Hoogeman, and B. J. Heijmen, "Comparison of vmat and imrt strategies for cervical cancer patients using automated planning," *Radiotherapy and Oncology*, vol. 114, no. 3, pp. 395–401, 2015.

[21] M. S. Gossman, S. S. Hancock, R. J. Kudchadker, P. R. Lundahl, M. Cao, and C. S. Melhus, "Brachytherapy dose-volume histogram commissioning with multiple planning systems," *Journal of Applied Clinical Medical Physics*, vol. 15, pp. 110–120, March 2014.

[22] N. H. Luong, T. Alderliesten, A. Bel, Y. Niatsetski, and P. A. Bosman, "Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment," *Swarm and Evolutionary Computation*, vol. 40, pp. 37–52, 2018.

[23] J. H. Holland, "Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence," 1992.

[24] P. A. Whigham, G. Dick, and J. Maclaurin, "On the mapping of genotype to phenotype in evolutionary algorithms," *Genet Program Evolvable Mach*, vol. 18, p. 353–361, 2017.

[25] T. Back and H.-P. Schwefel, "An overview of evolutionary algorithms for parameter optimization," *Evolutionary Computation*, vol. 1, pp. 1–23, March 1993.

[26] S. Katoch, S. S. Chauhan, and V. Kumar, "A review on genetic algorithm: past, present, and future," *Multimedia tools and applications*, vol. 80, p. 8091–8126, 2020.

[27] T. Blickle and L. Thiele, "A comparison of selection schemes used in evolutionary algorithms," *Evolutionary Computation*, vol. 4, no. 4, p. 361–394, 1996.

[28] P. J. Hancock, "An empirical comparison of selection methods in evolutionary algorithm," *AISB workshop on evolutionary computing*, p. 80–94, 1994.

[29] A. J. Umbarkar and P. D. Sheth, "Crossover operators in genetic algorithms: A review," *ICTACT journal on soft computing*, vol. 6, no. 1, 2015.

[30] A. C. Koenig, "A study of mutation methods for evolutionary algorithms," 2002.

[31] W. M. Spears, "Evolutionary algorithms: The role of mutation and recombination," *Springer Science Business Media*, 2000.

[32] R. Santana, "Gray-box optimization and factorized distribution algorithms: where two worlds collide," *Model-Based Evolutionary Algorithms workshop*, 2016.

[33] J. Stork, A. E. Eiben, and T. Bartz-Beielstein, "A new taxonomy of global optimization algorithms," *Natural Computing*, vol. 21, p. 219–242, 2022.

[34] D. Thierens and P. A. Bosman, "Optimal mixing evolutionary algorithms," *GECCO 2011: Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, p. 617–624, July 2011.

[35] A. Bouter, T. Alderliesten, C. Witteveen, and P. A. N. Bosman, "Exploiting linkage information in real-valued optimization with the real-valued gene-pool optimal mixing evolutionary algorithm," *Proceedings of the Genetic and Evolutionary Computation Conference*, p. 705–712, July 2017.

[36] P. A. Bosman and D. Thierens, "Linkage neighbors, optimal mixing and forced improvements in genetic algorithms," *GECCO '12: Genetic and Evolutionary Computation Conference*, pp. 585 – 592, 2012.

[37] P. A. Bosman and D. Thierens, "On measures to build linkage trees in ltga," *Lecture Notes in Computer Science*, vol. 7491, pp. 276–285, 2012.

[38] P. A. N. Bosman, J. Grahl, and D. Thierens, "Benchmarking parameter-free amalgam on functions with and without noise," *Evolutionary Computation*, p. 445–469, September 2013.

[39] A. Bouter, N. H. Luong, C. Witteveen, T. Alderliesten, and P. A. N. Bosman, "The multi-objective real-valued gene-pool optimal mixing evolutionary algorithm," *Proceedings of the Genetic and Evolutionary Computation Conference*, p. 537–544, July 2017.

[40] N. H. Luong, T. Alderliesten, A. Bel, Y. Niatsetski, and P. A. Bosman, "Application and benchmarking of multi-objective evolutionary algorithms on high-dose-rate brachytherapy planning for prostate cancer treatment," *Swarm and Evolutionary Computation*, vol. 40, pp. 37–52, June 2018.

[41] D. L. Barten, B. R. Pieters, A. Bouter, M. C. van der Meer, S. C. Maree, K. A. Hinnen, H. Westerveld, P. A. Bosman, T. Alderliesten, N. van Wieringen, and A. Bel, "Towards artificial intelligence-based automated treatment planning in clinical practice: A prospective study of the first clinical experiences in high-dose-rate prostate brachytherapy," *Brachytherapy*, vol. 22, no. 2, pp. 279–289, 2023.

[42] A. Bouter and P. A. N. Bosman, "Gpu-accelerated parallel gene-pool optimal mixing in a gray-box optimization setting," *GECCO '22: Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 675 – 683, 2022.

[43] L. R. M. Dickhoff, E. M. Kerkhof, H. H. Deuzeman, D. L. J. Barten, L. A. Velema, L. J. A. Stalpers, B. R. Pieters, C. L. Creutzberg, P. A. N. Bosman, and T. Alderliesten, "The versatility of evolutionary intelligent tri-objective treatment planning for cervical cancer brachytherapy," *Medical Physics*, vol. 52, August 2025.

Evolutions of Single-DVI Experiment

Evolution through 8 generations for all runs of a single DVI experiment: CTV HR without normalization factor, CTV HR with normalization factor, PTV without normalization factor, PTV with normalization factor.

The region of the search space with a better fitness according to the search-space exploration is represented in blue. Notice that all 4 scenarios have different heatmaps as the fitness landscape is different.

The black dots represent the individuals in the populations that start randomly sampled and end up converging.

CTV_{HR} not normalized

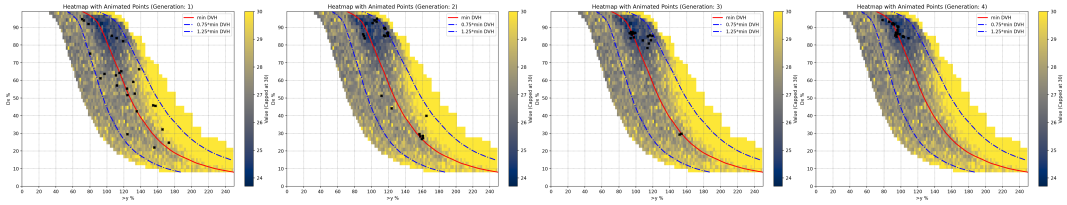

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

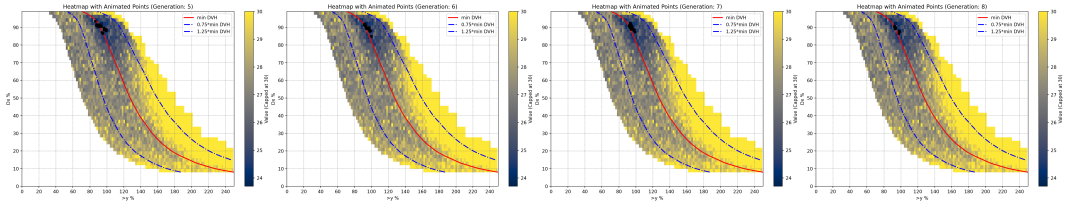

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 1: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 1

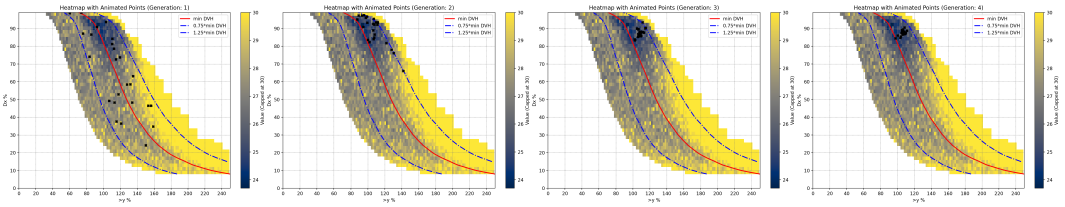

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

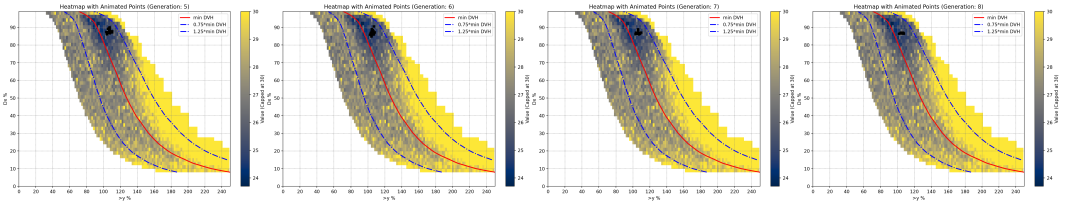

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 2: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 2

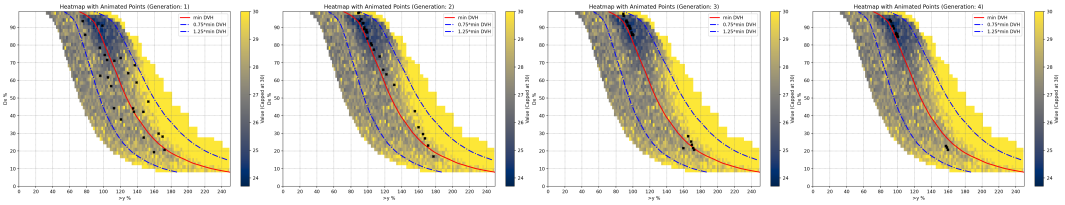

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

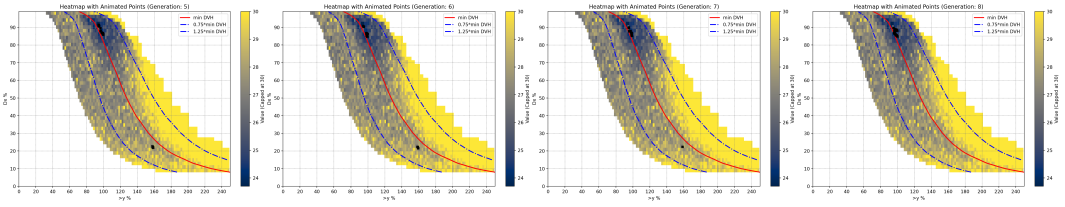

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 3: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 3

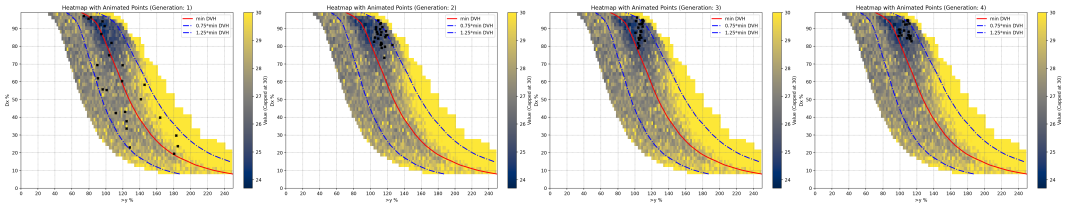

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

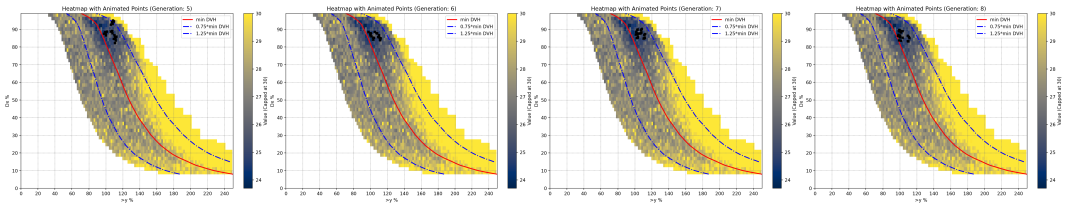

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 4: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 4

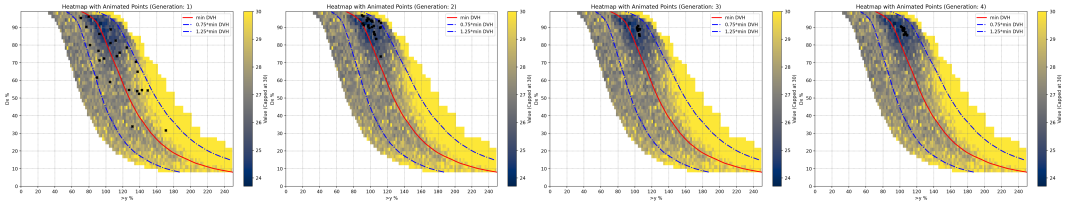

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

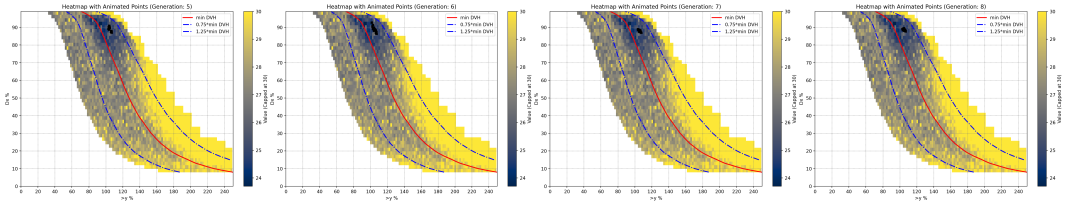

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 5: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 5

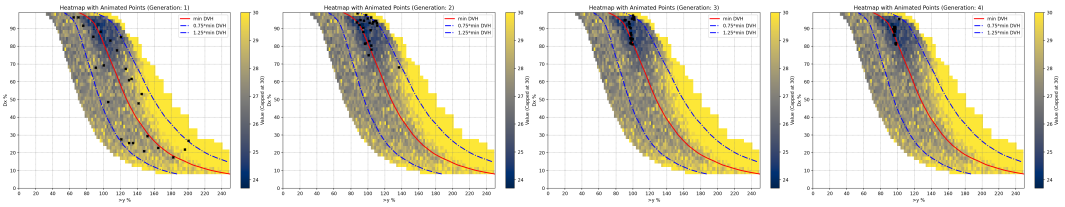

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

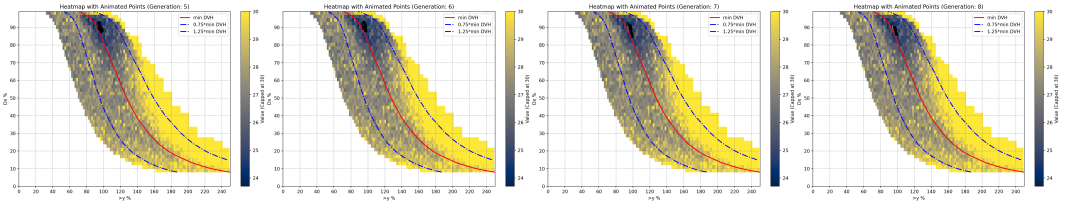

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 6: Evolution of the single DVI experiment, CTV_{HR} without normalization factor, run 6

CTV_{HR} normalized

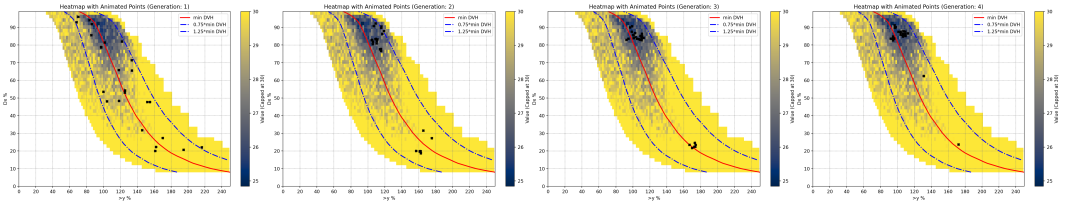

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

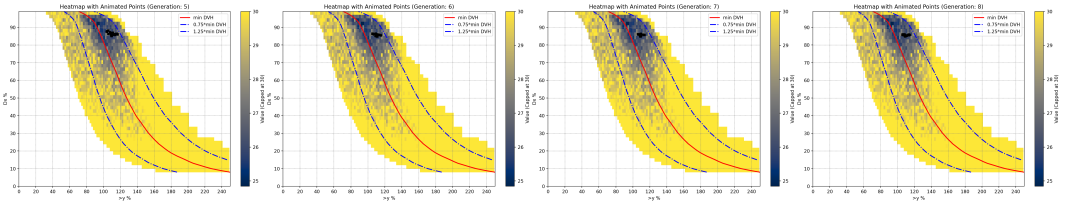

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 7: Evolution of the single DVI experiment, CTV_{HR} with normalization factor, run 1

PTV not normalized

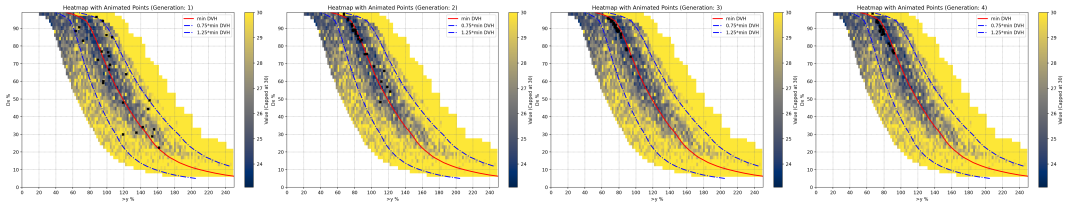

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

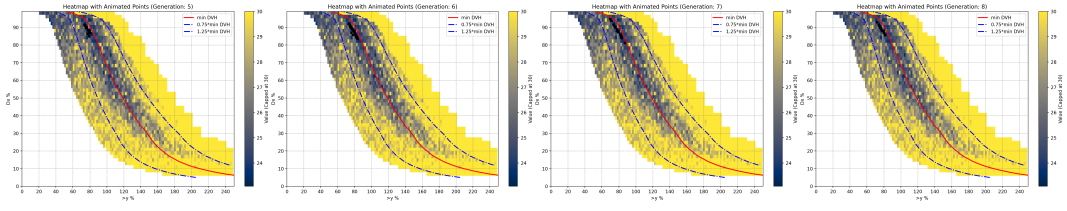

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 8: Evolution of the single DVI experiment, *PTV* without normalization factor, run 1

PTV normalized

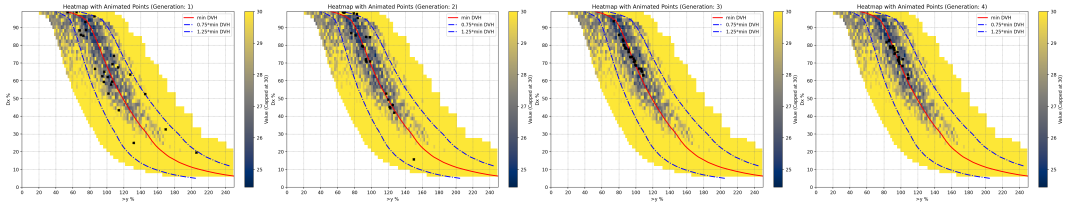

(a) Generation 1 (b) Generation 2 (c) Generation 3 (d) Generation 4

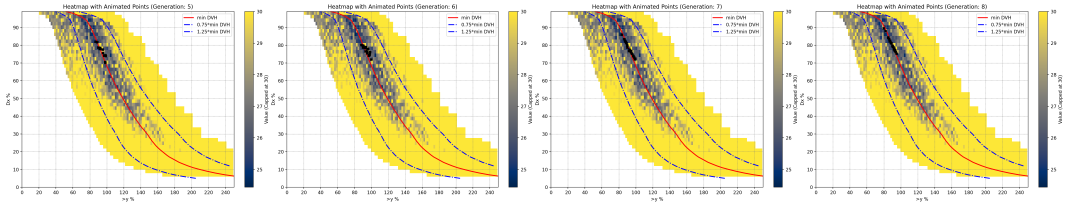

(e) Generation 5 (f) Generation 6 (g) Generation 7 (h) Generation 8

Figure 9: Evolution of the single DVI experiment, *PTV* with normalization factor, run 1

Comparison of generated treatment plans for single DVI optimized protocols with and without the normalization factor

To generate the boxplots, we ran BRIGHT 100 times for each scenario and took the plan that had the closest radiation around the Ring of normal tissue. Then reported the distance from the clinical plan (dotted line).

Single DVI Optimization for the CTV_{HR} volume.

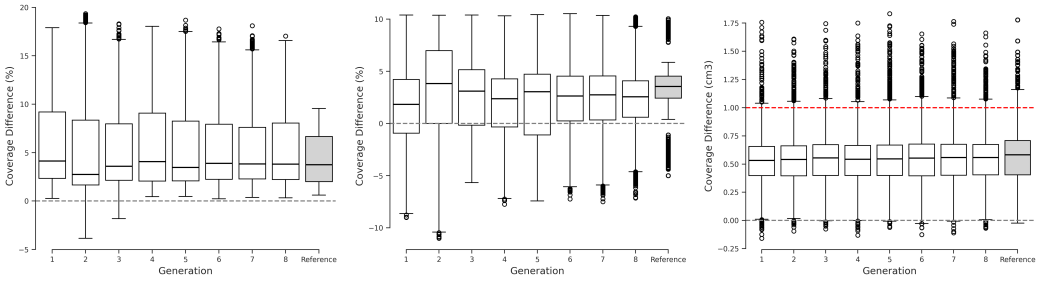

(a) CTV HR

(b) PTV

(c) Bladder

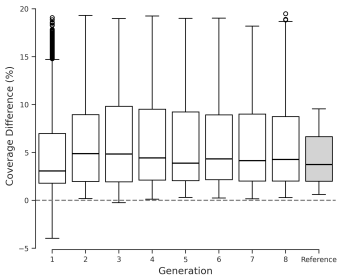

(d) CTV HR

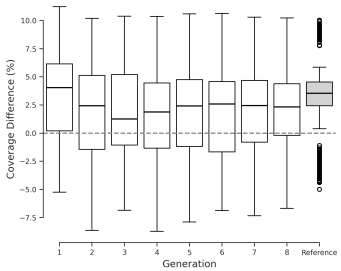

(e) PTV

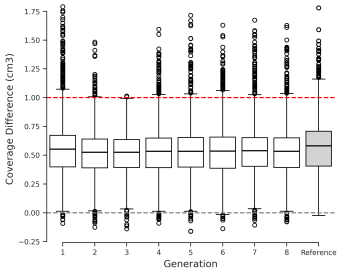

(f) Bladder

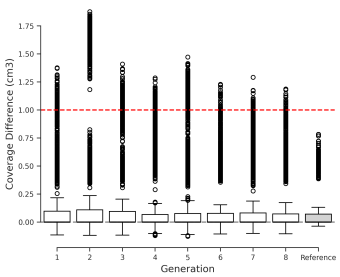

(g) Rectum

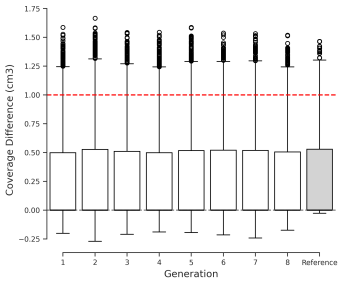

(h) Sigmoid

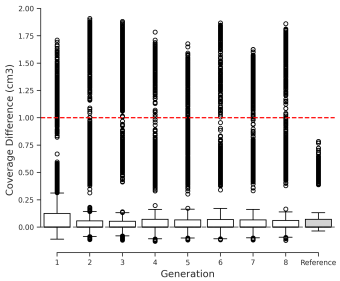

(i) Rectum

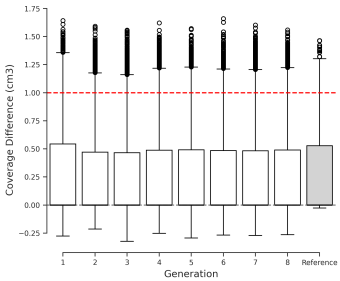

(j) Sigmoid

Figure 10: Generated plans for the single DVI optimization for the CTV_{HR} without normalization (upper) vs with normalization factor (lower) for the validation patient.

Single DVI Optimization for the *PTV* volume.

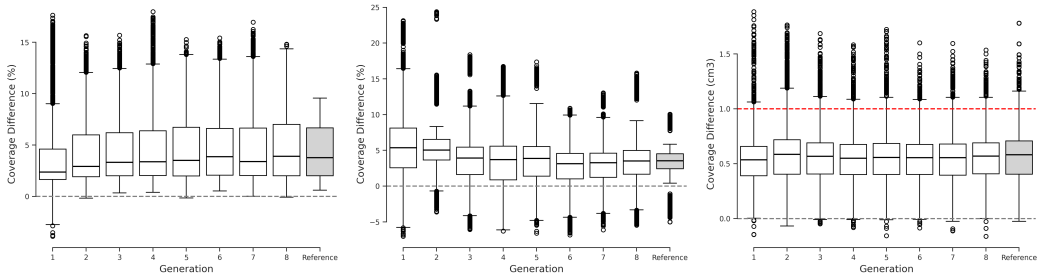

(a) CTV HR (b) PTV (c) Bladder

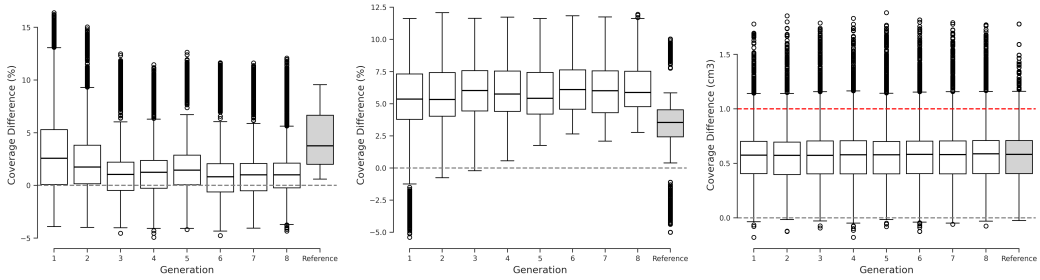

(d) CTV HR (e) PTV (f) Bladder

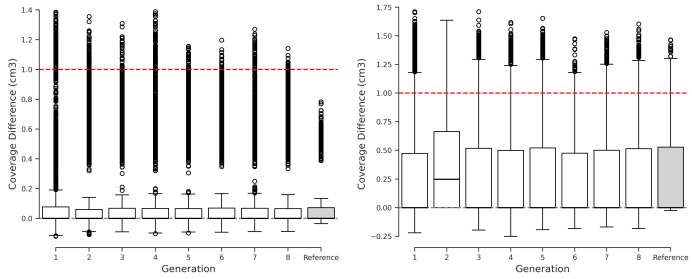

(g) Rectum (h) Sigmoid

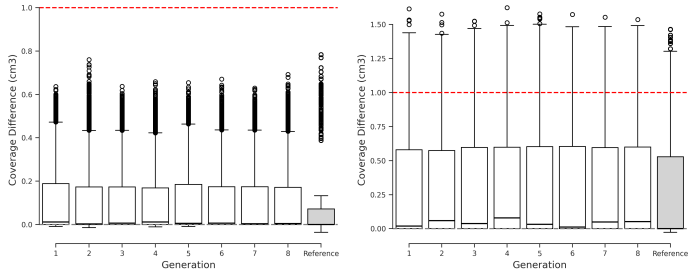

(i) Rectum (j) Sigmoid

Figure 11: Generated plans for the single DVI optimization for the *PTV* without normalization (upper) vs with normalization factor (lower) for the validation patient.

Dose distributions

This Appendix shows the dose distributions for 5 different CT slides for each patient using the protocols that showed improvement. Therefore, we analyzed the best protocol from generations 1, 2, 4, and 6.

The colored lines on the right picture correspond to the isodose lines for: 50%, 75%, 100%, and 150% of the nominal radiation dose (7 Gy per fraction). The colored regions indicate the ROI: CTV_{HR} (blue), PTV (red), $Bladder$ (yellow), *Sigmoid* (green), *Rectum* (brown).

The color-shaded areas in the left picture depict differences in radiation with respect to the clinical treatment plan. Green means greater radiation, and purple indicates lower. For this research, either worsens the plan. Thus, less color means better.

All images were taken from BRIGHT.

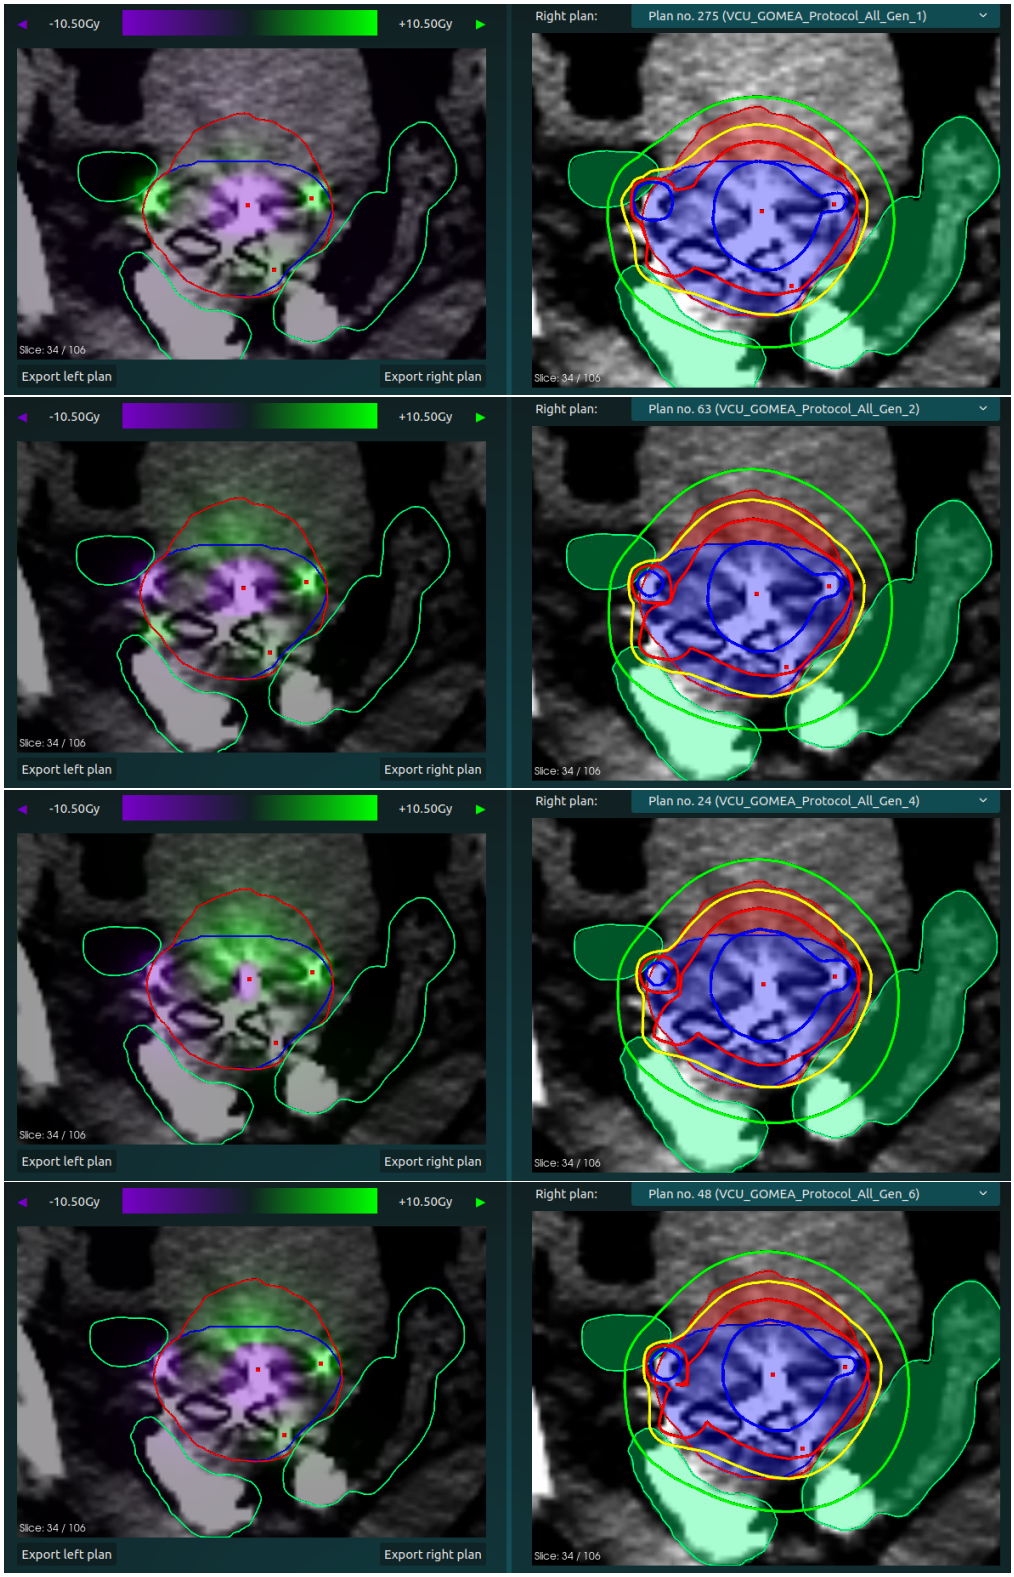

Figure 12: Patient 1, CT Slice 34

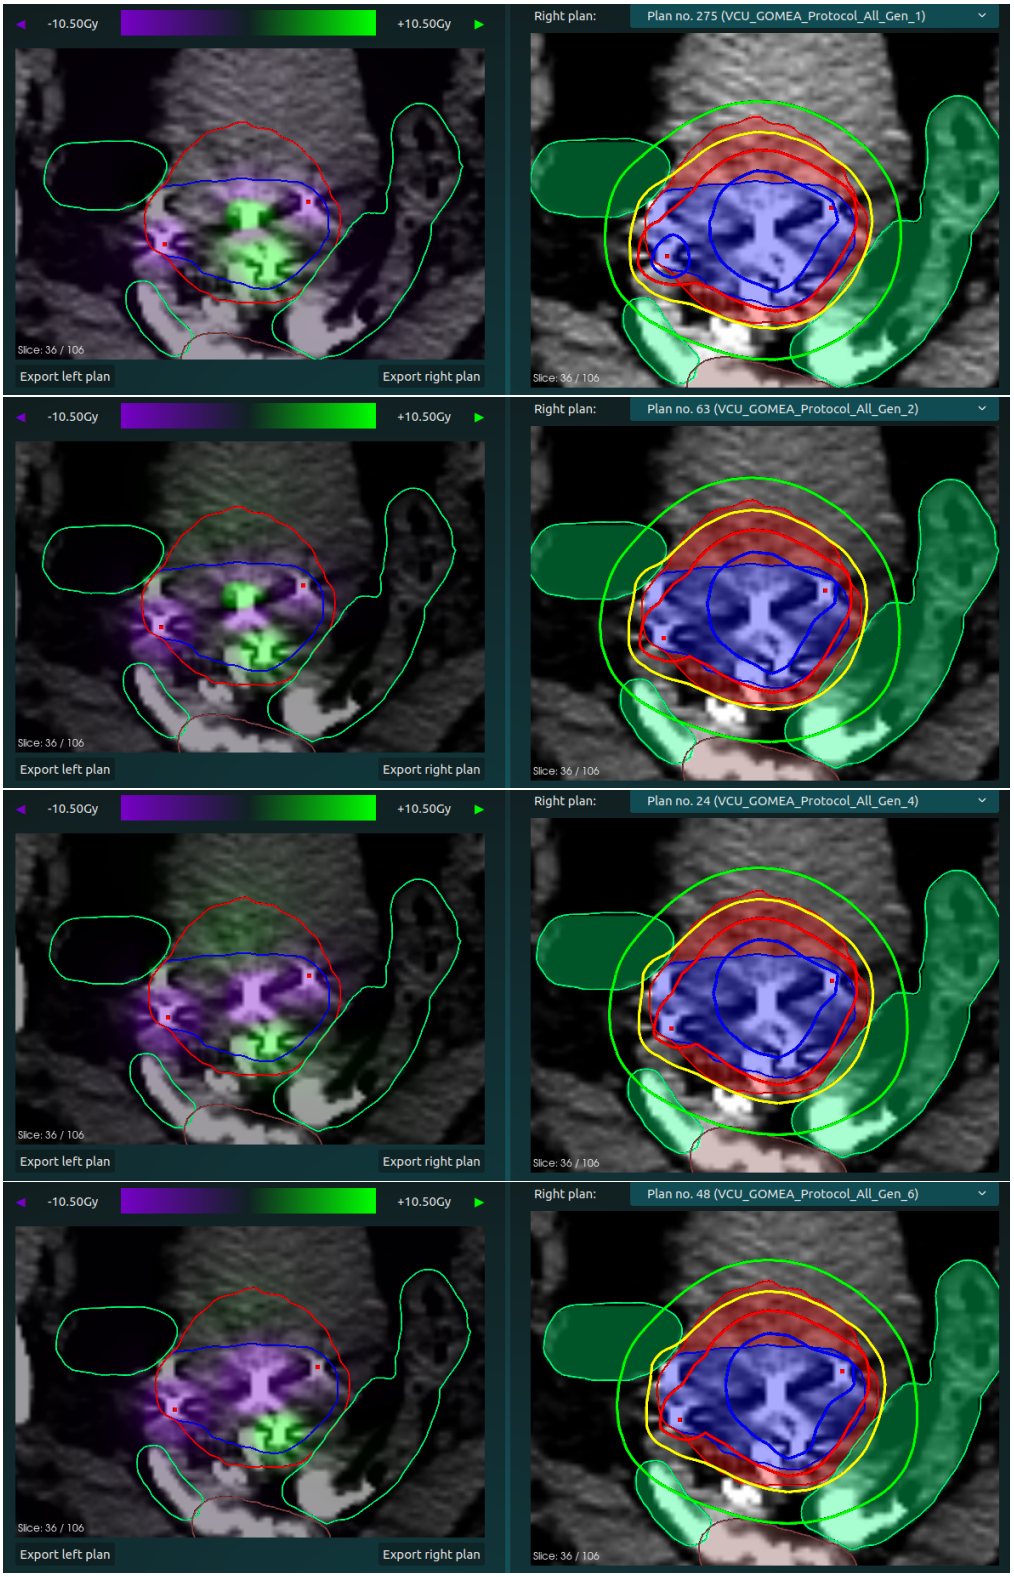

Figure 13: Patient 1, CT Slice 36

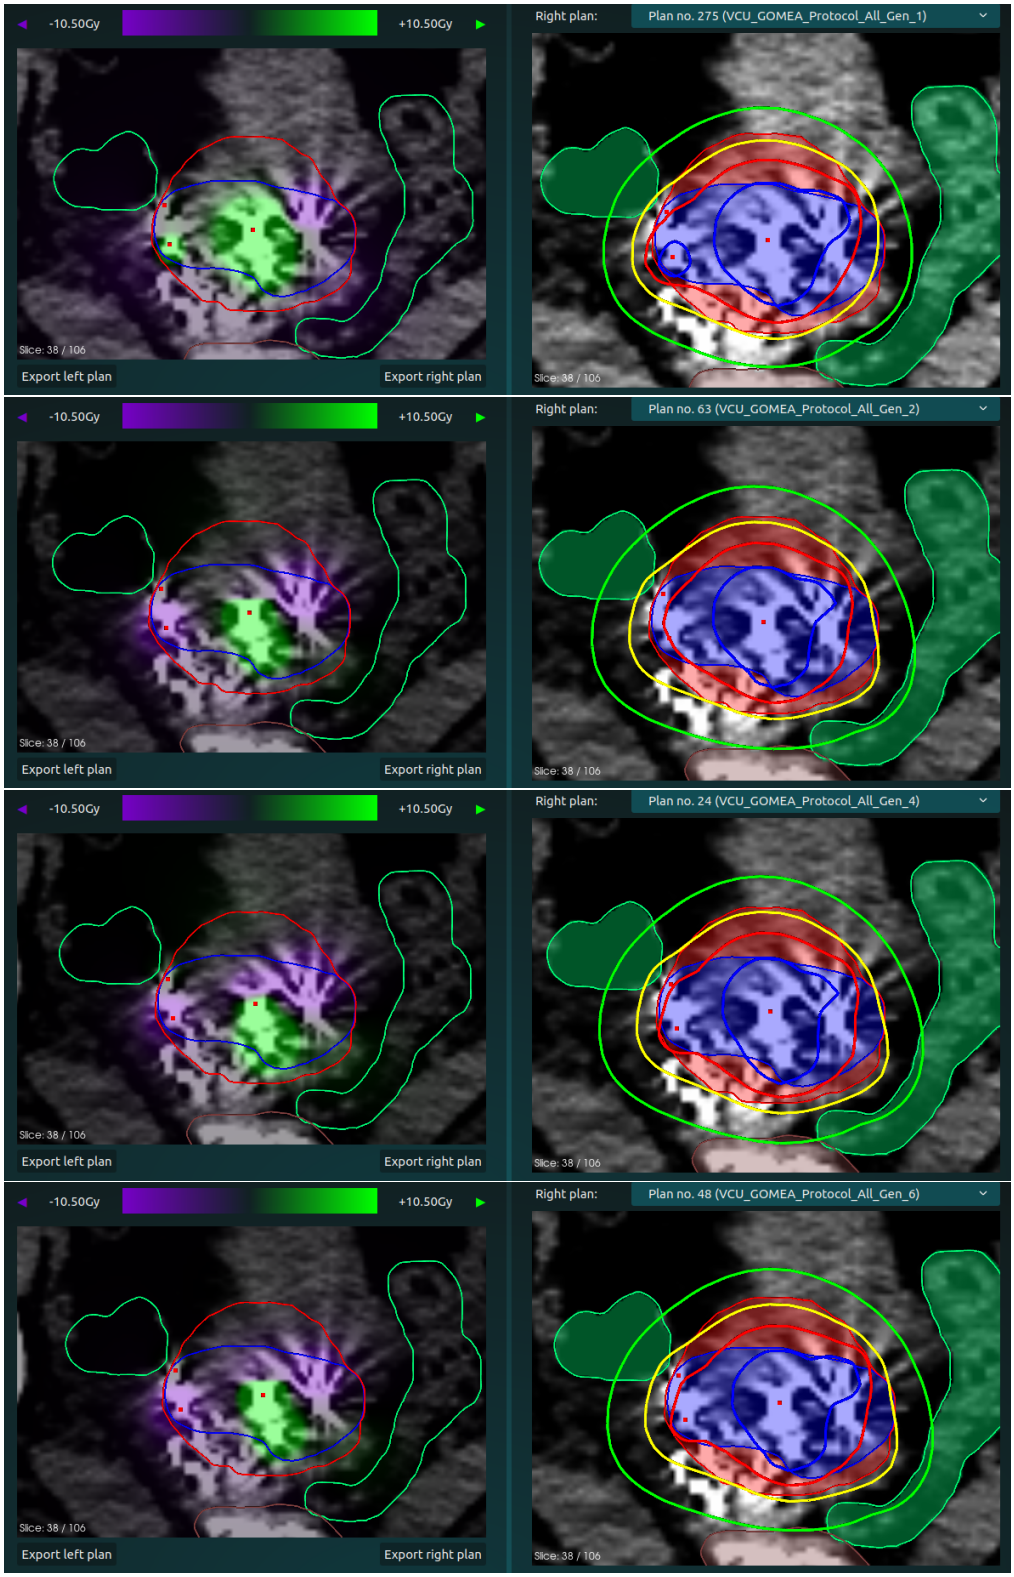

Figure 14: Patient 1, CT Slice 38

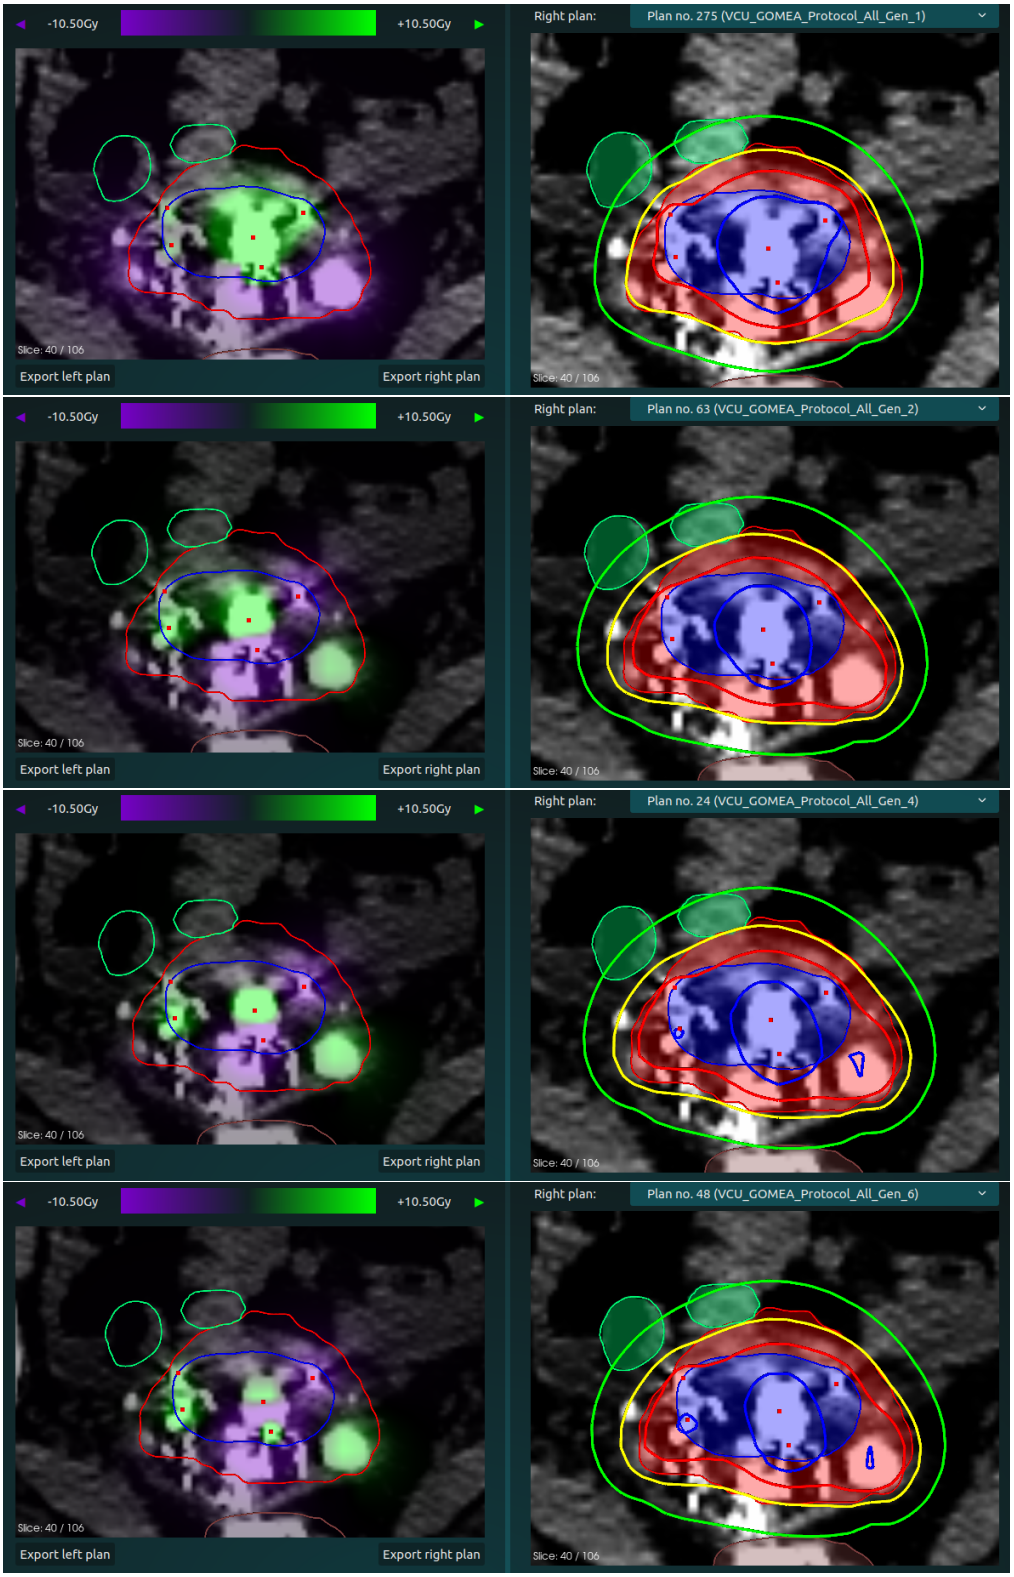

Figure 15: Patient 1, CT Slice 40

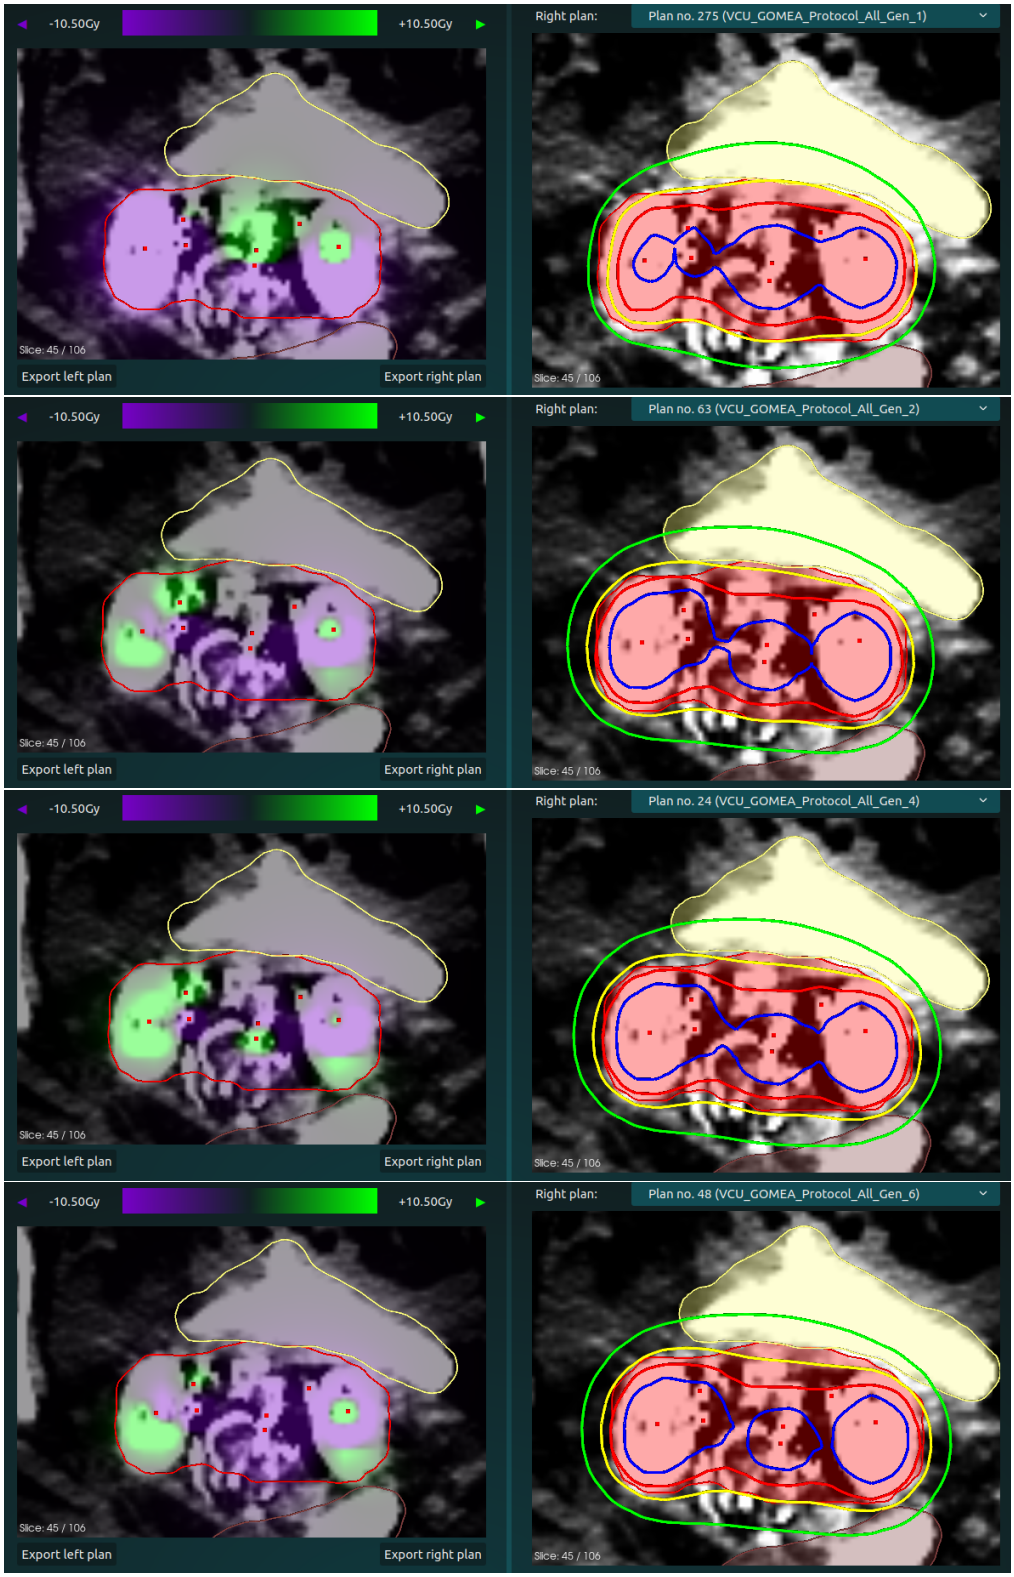

Figure 16: Patient 1, CT Slice 45

Figure 17: Patient 2, CT Slice 23

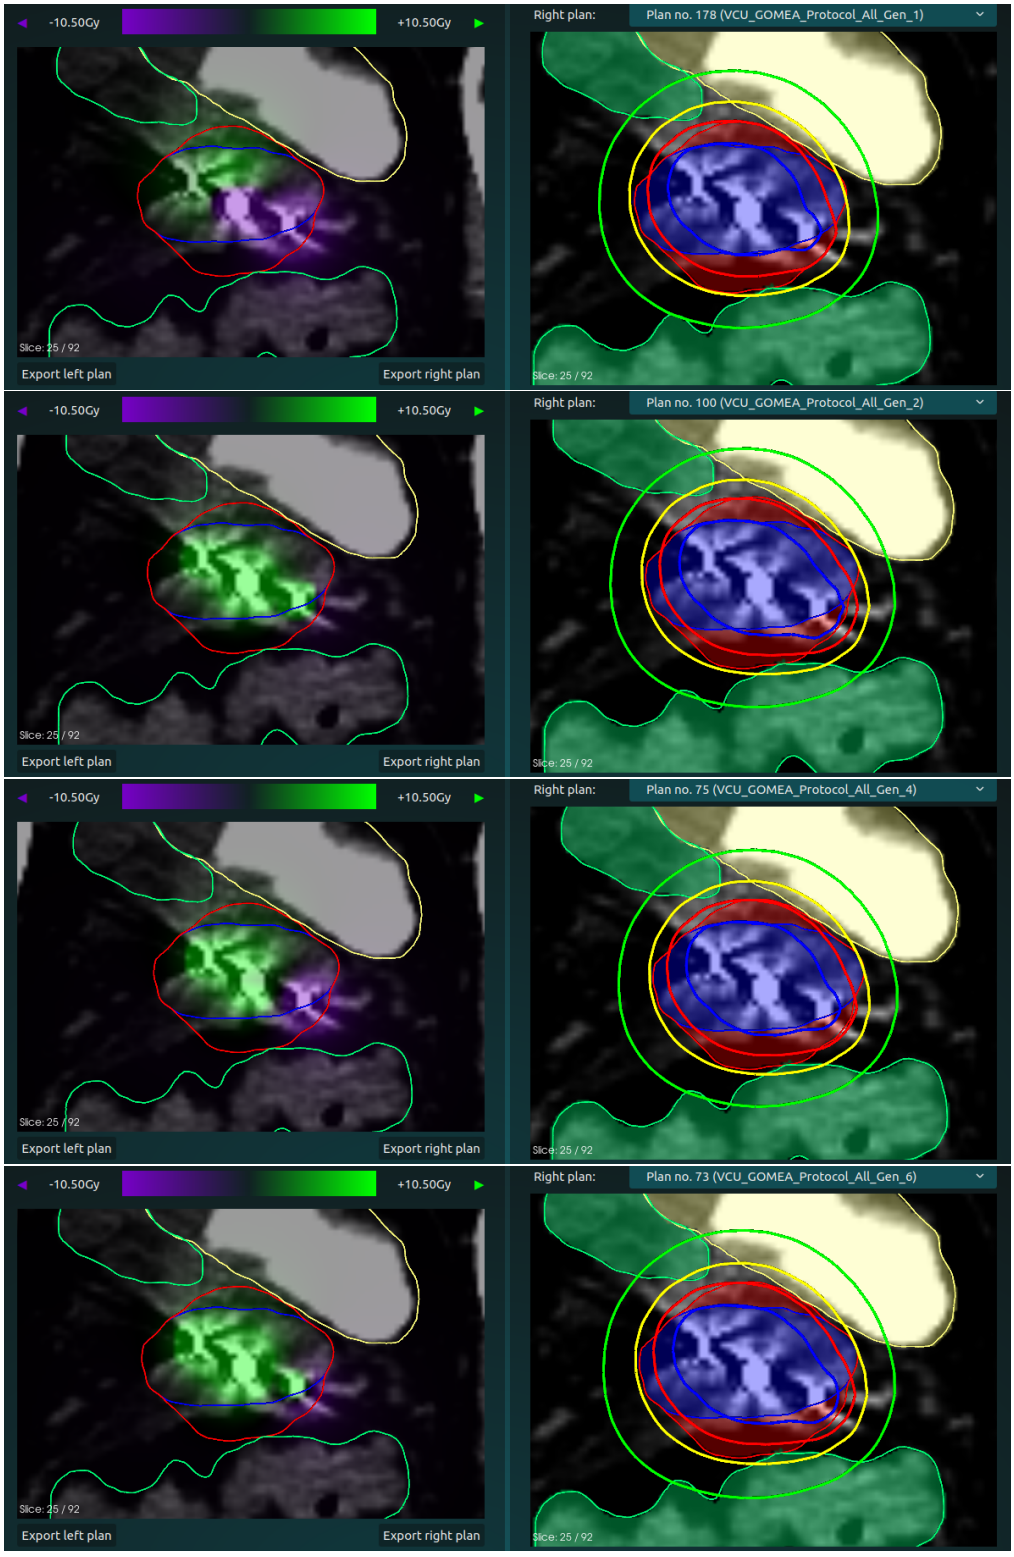

Figure 18: Patient 2, CT Slice 25

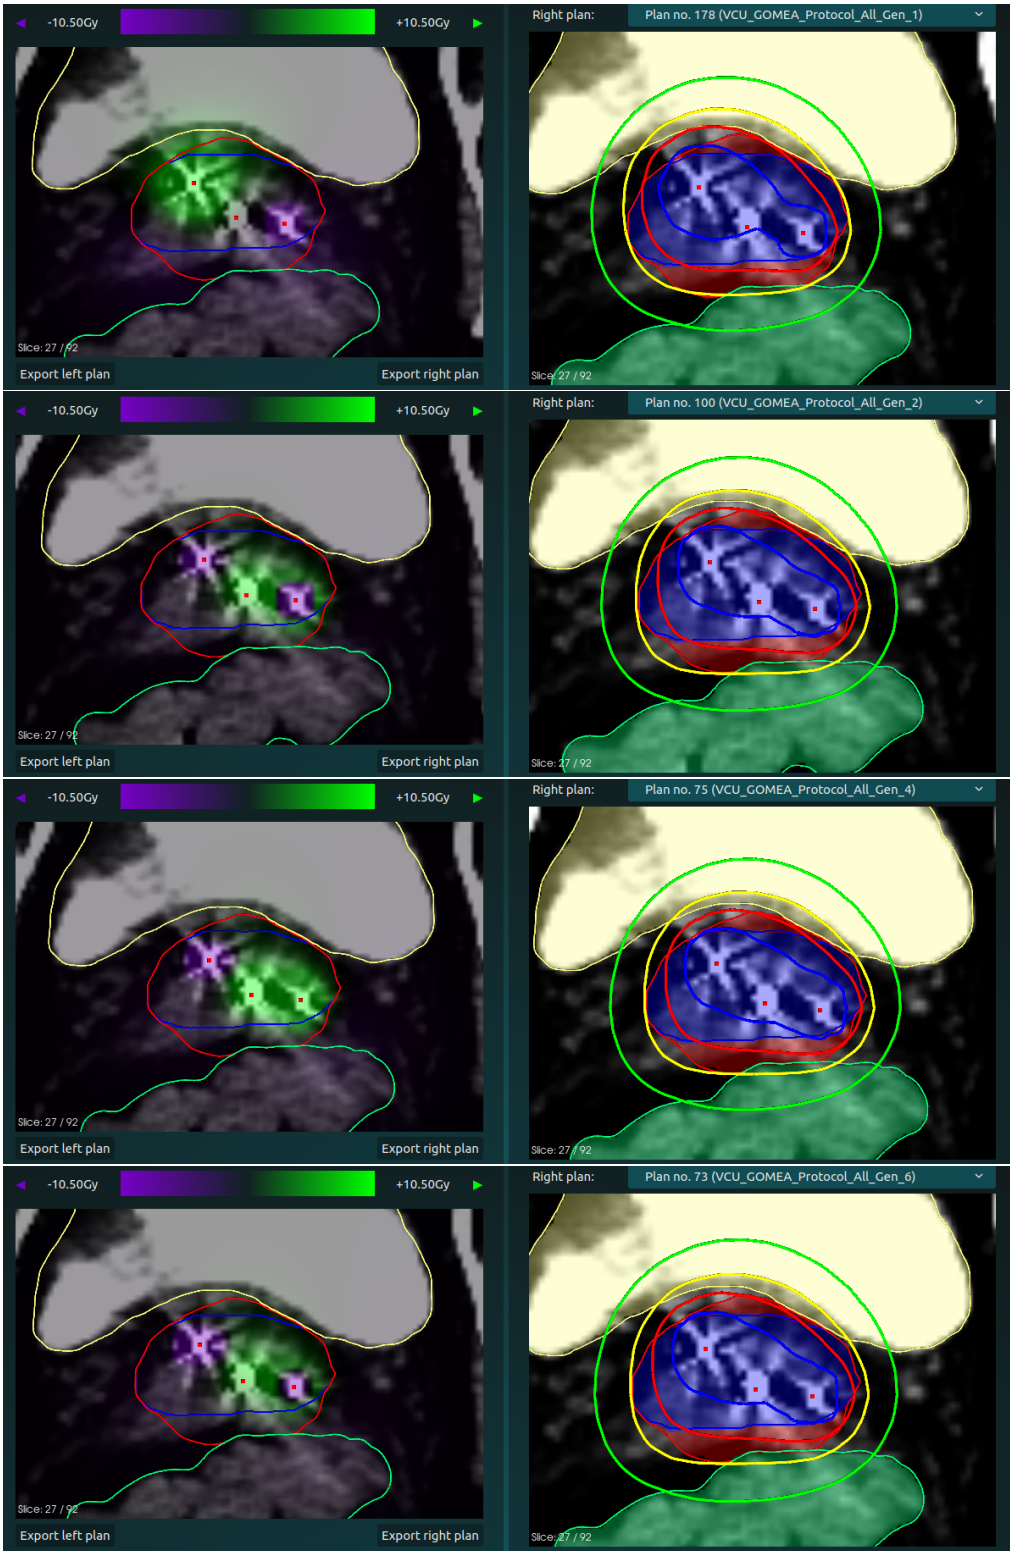

Figure 19: Patient 2, CT Slice 27

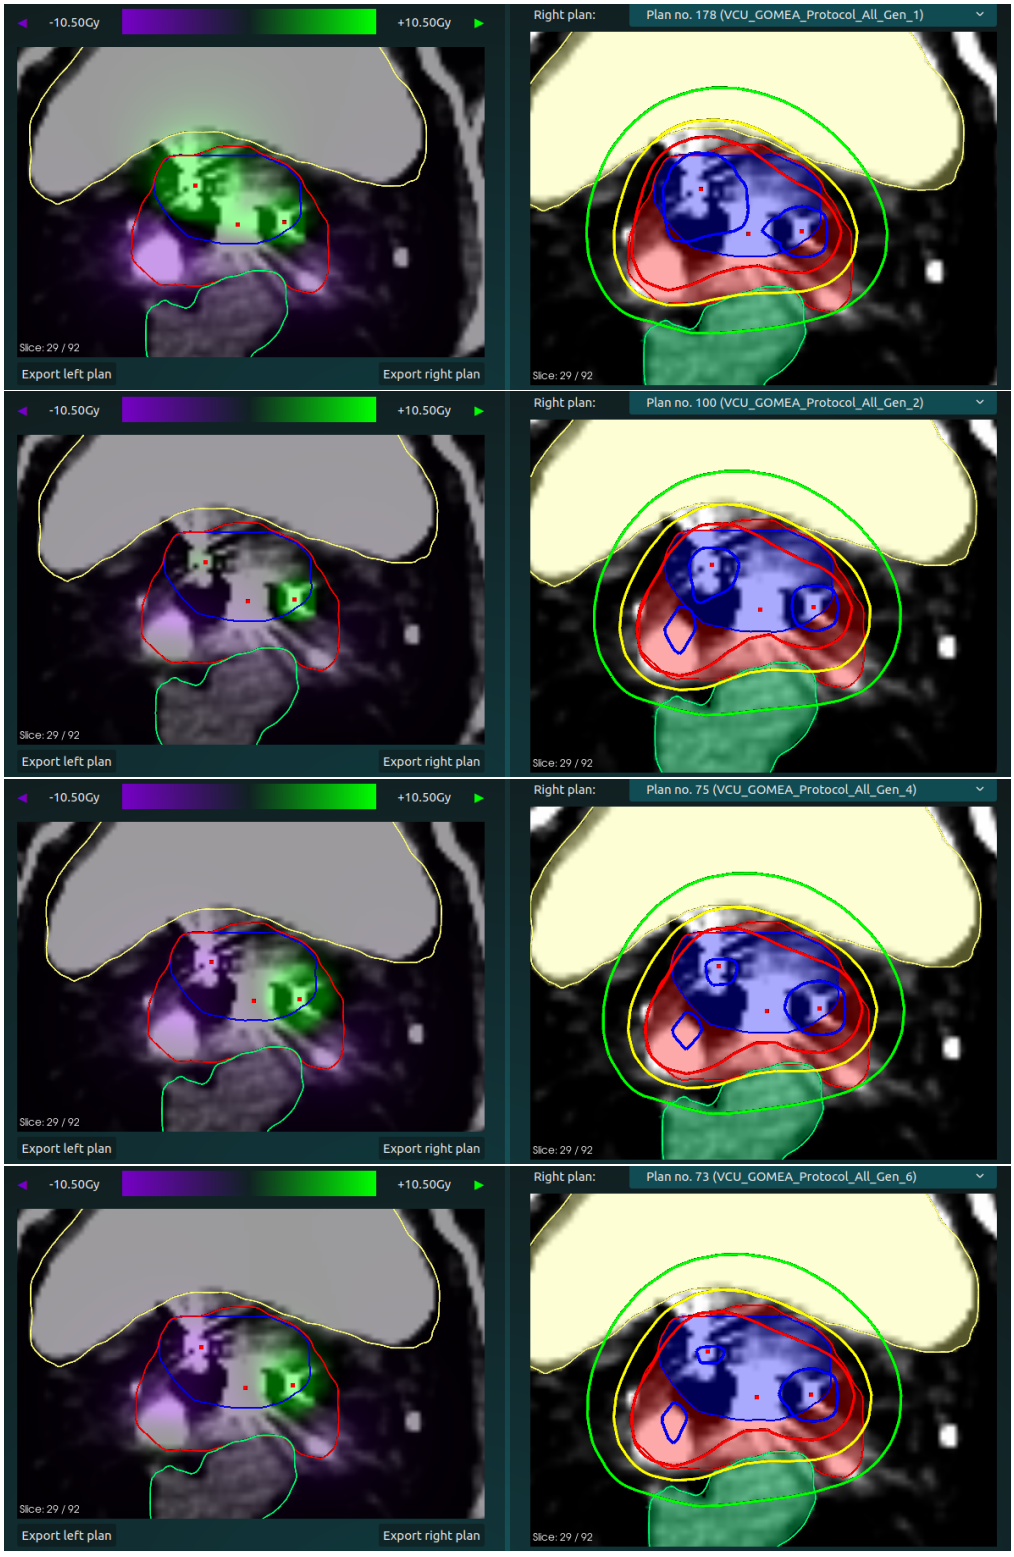

Figure 20: Patient 2, CT Slice 29

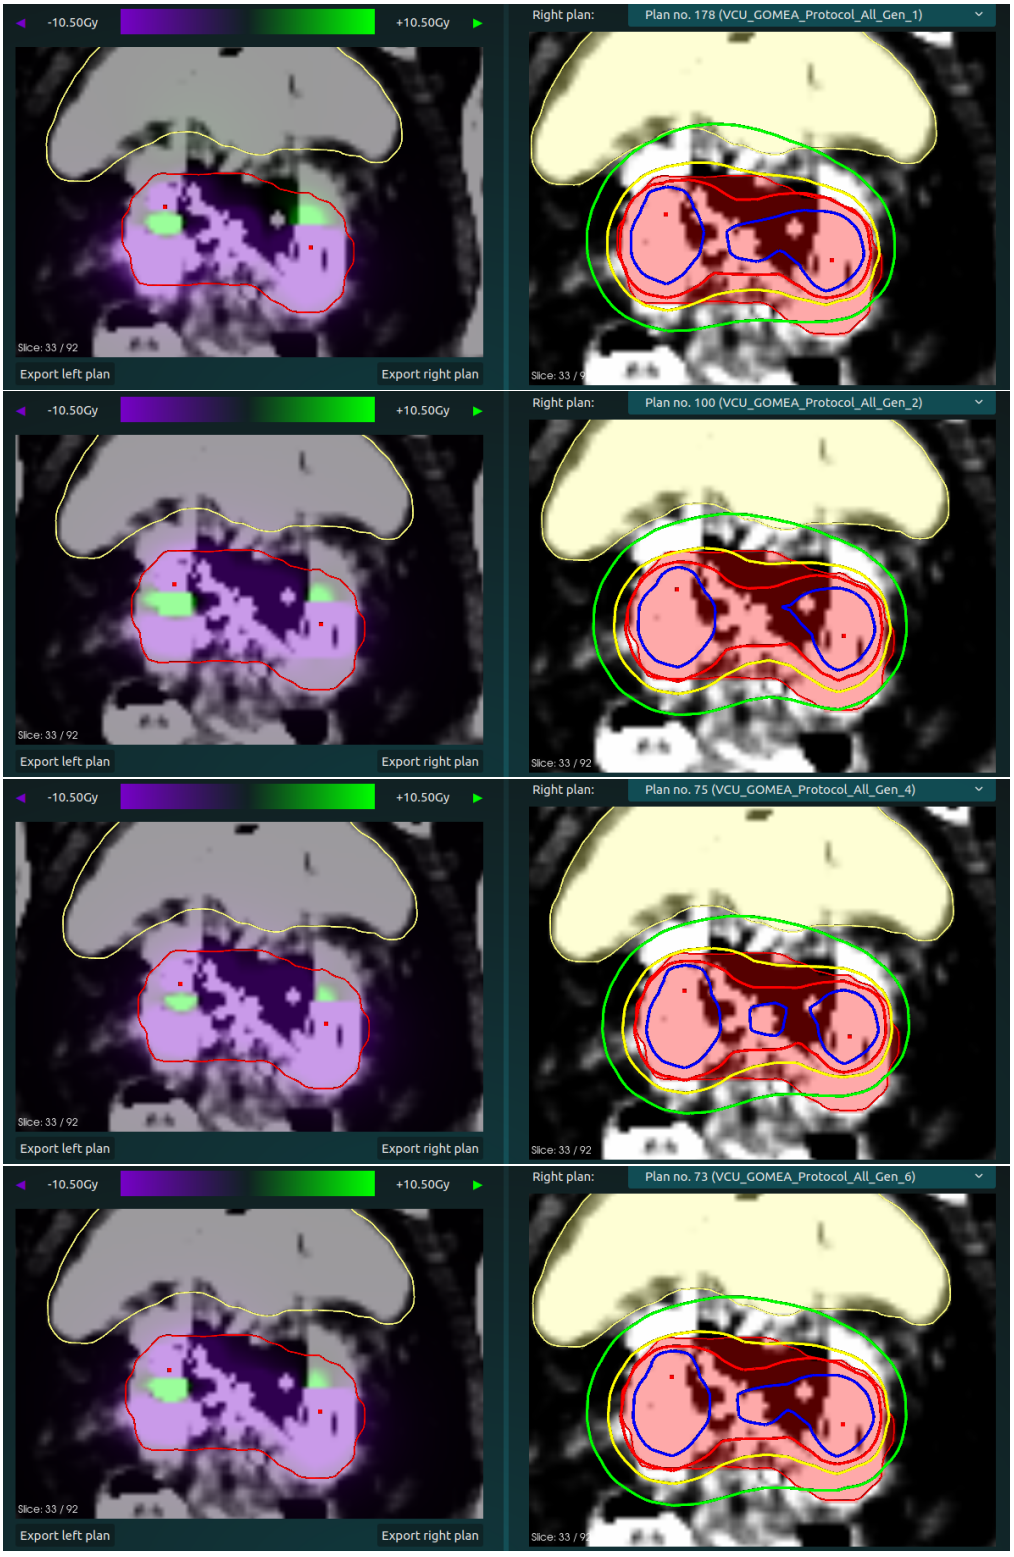

Figure 21: Patient 2, CT Slice 33

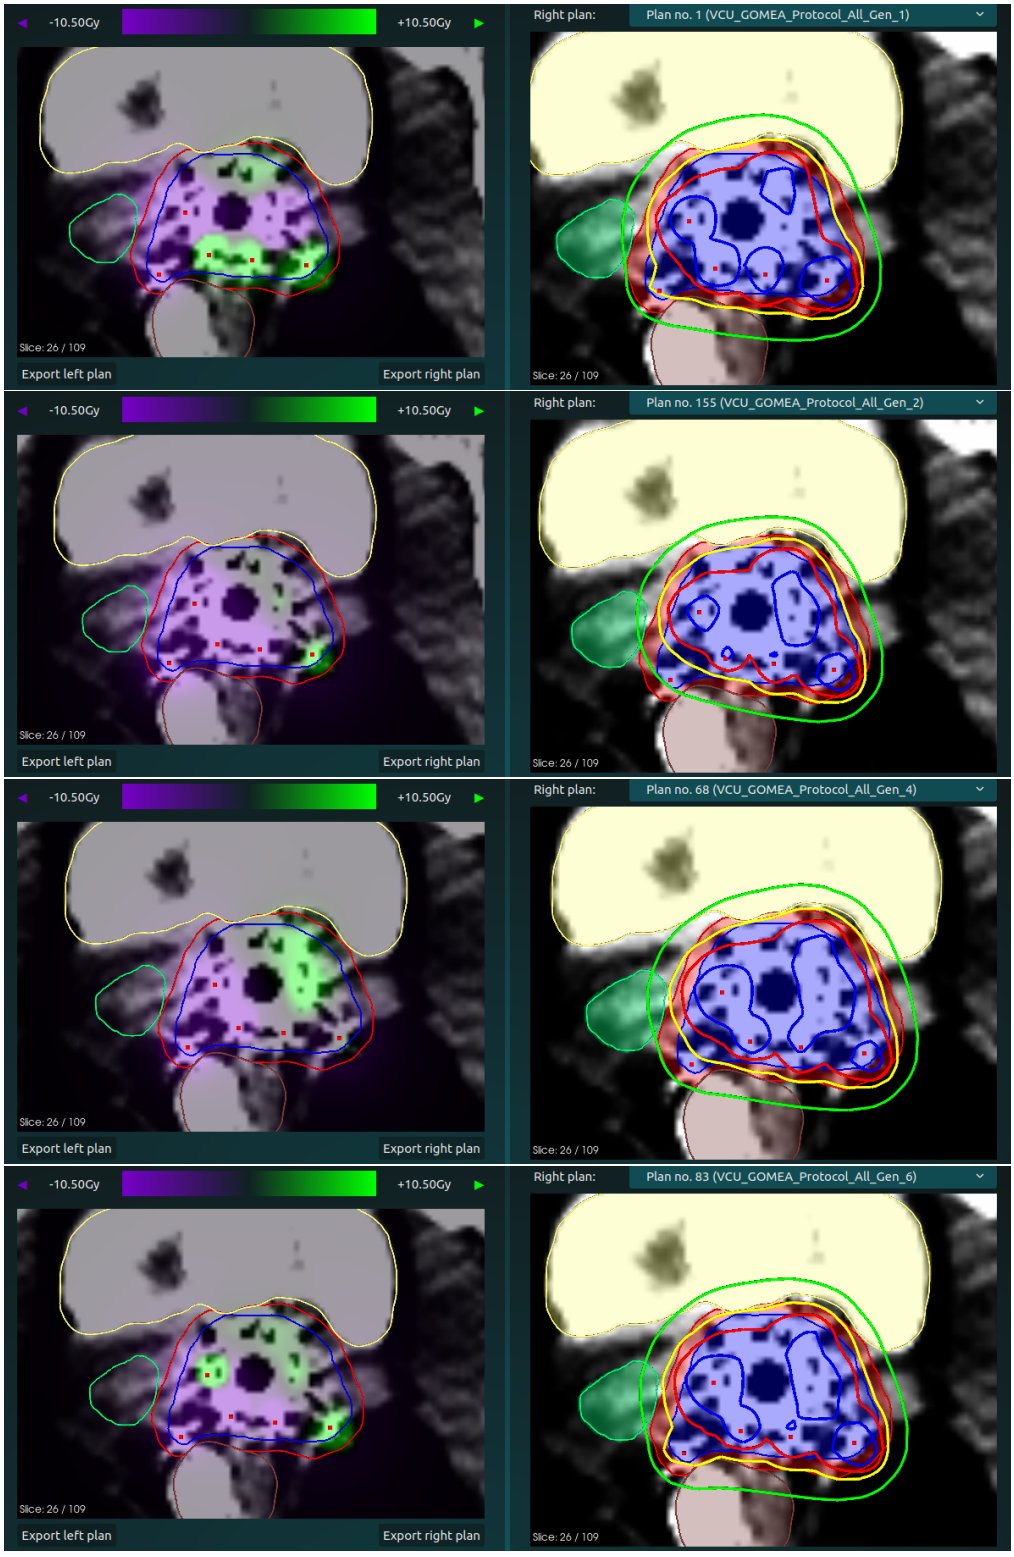

Figure 22: Patient 3, CT Slice 26

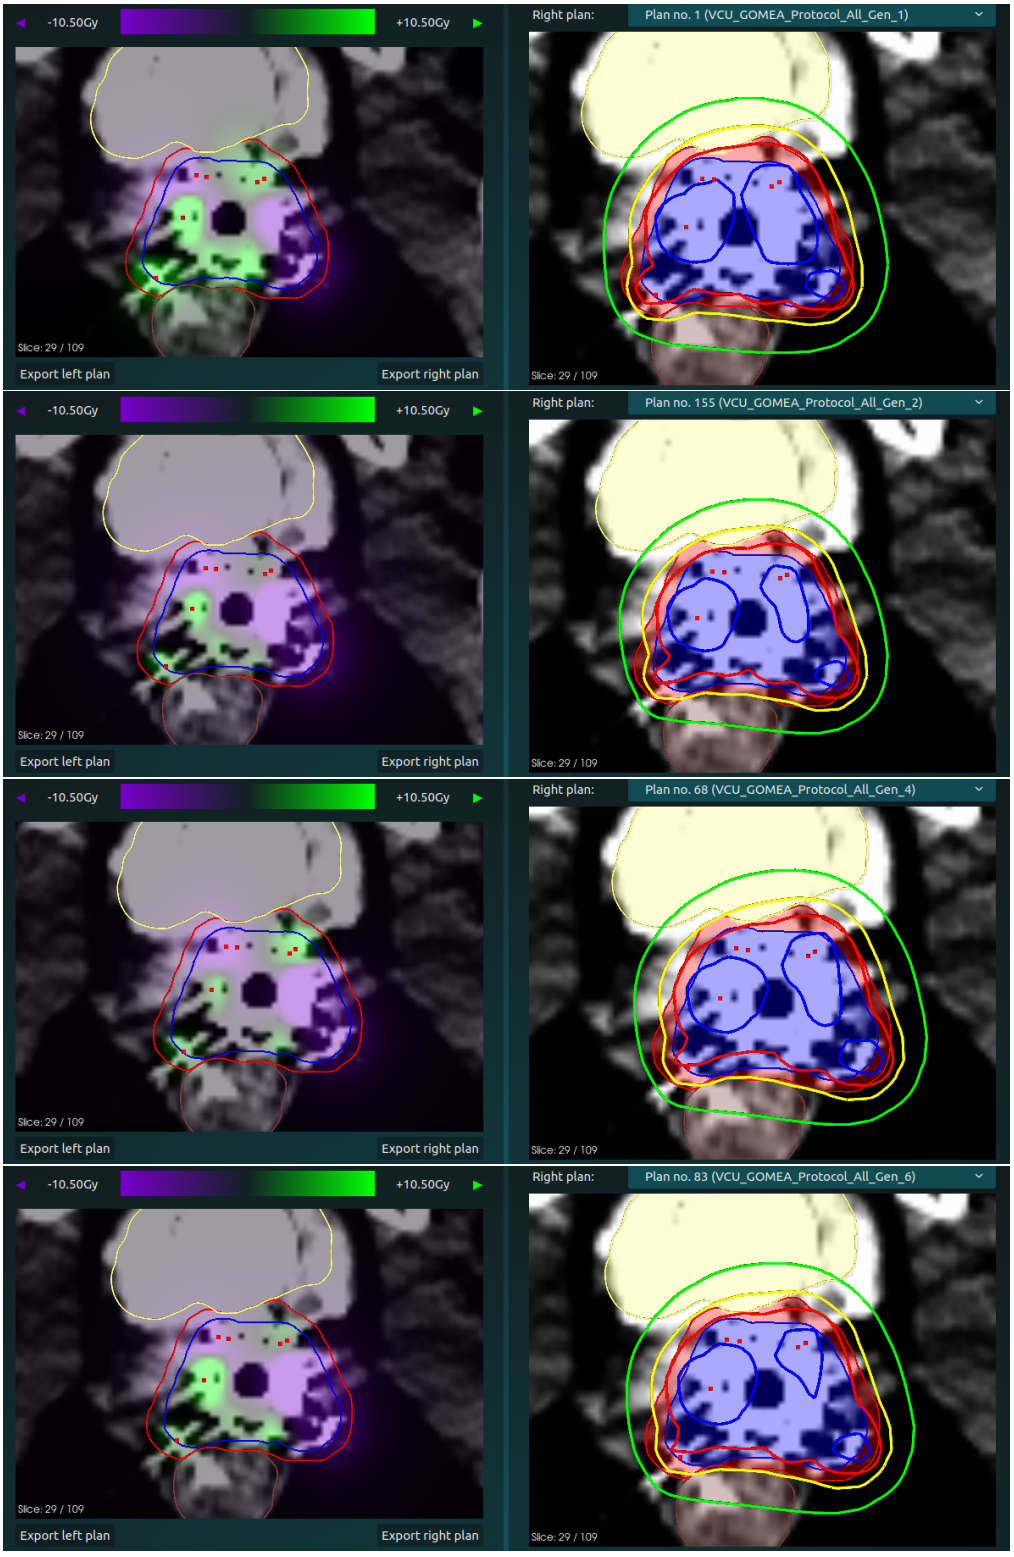

Figure 23: Patient 3, CT Slice 29

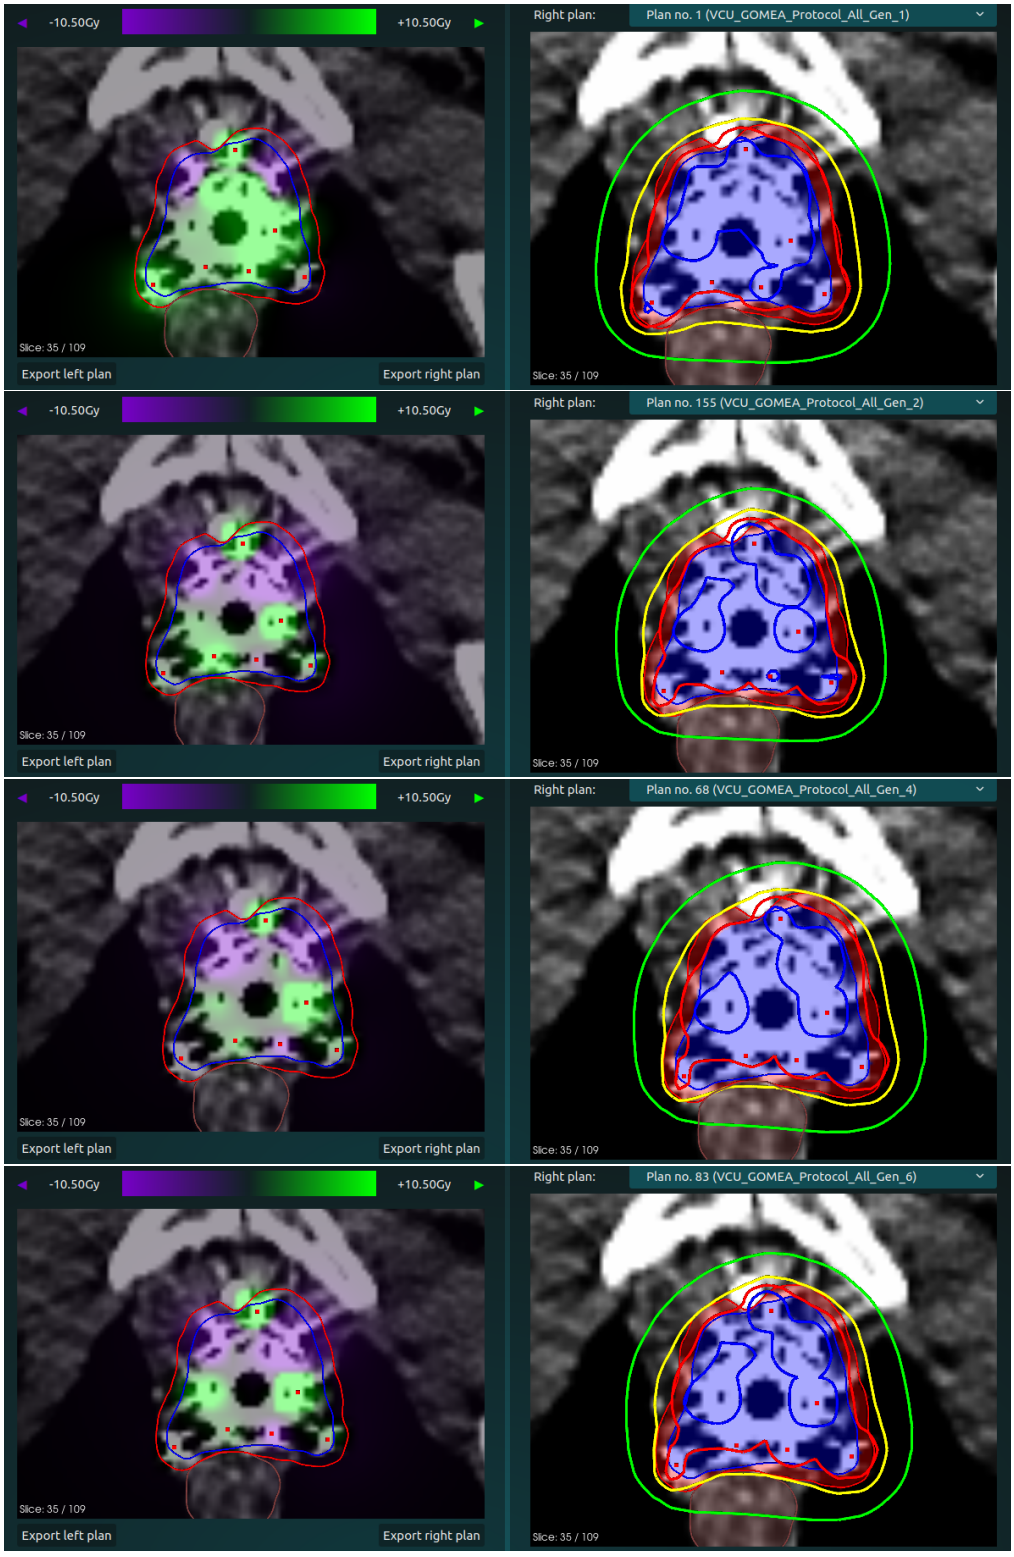

Figure 24: Patient 3, CT Slice 35

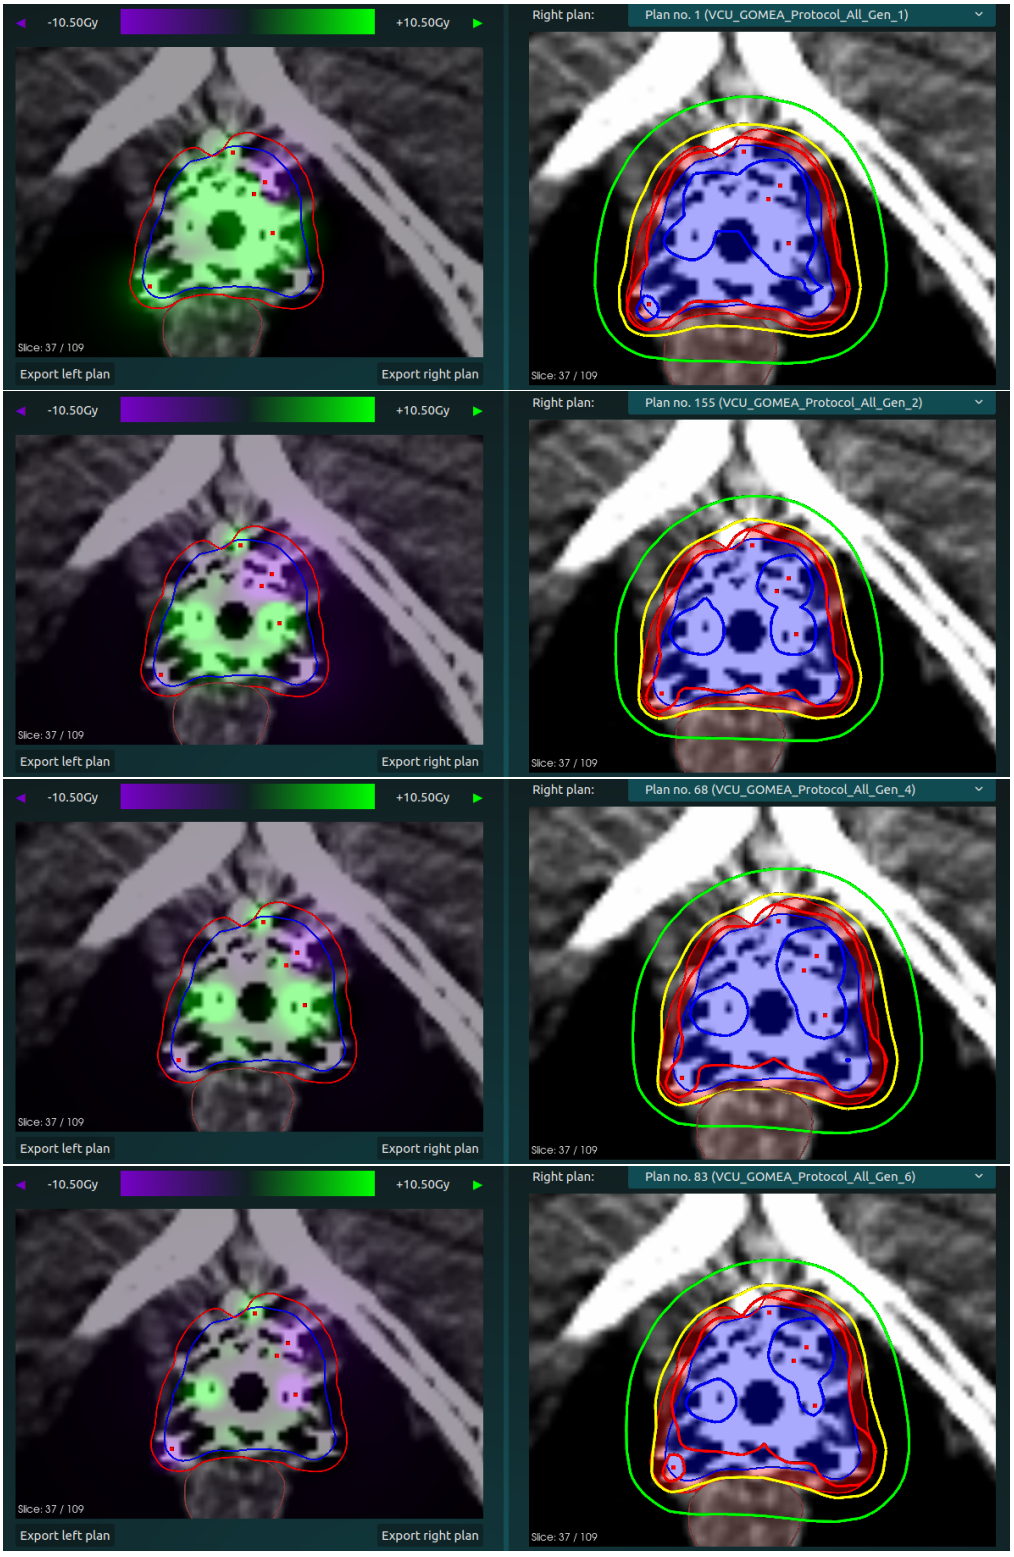

Figure 25: Patient 3, CT Slice 37

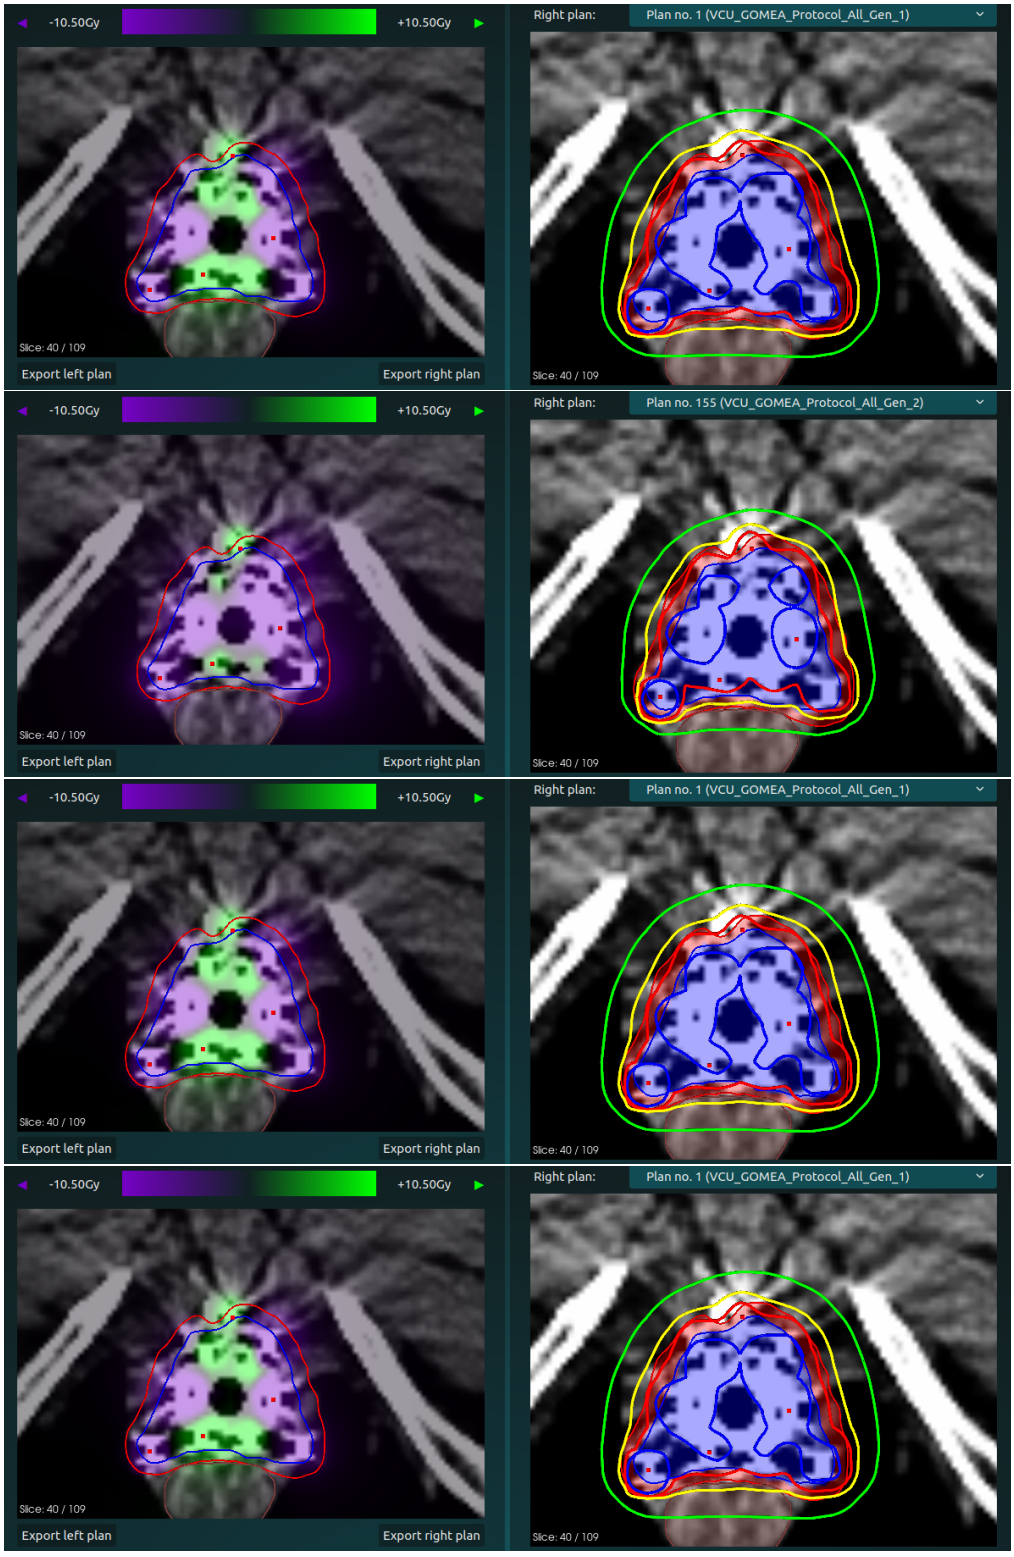

Figure 26: Patient 3, CT Slice 40

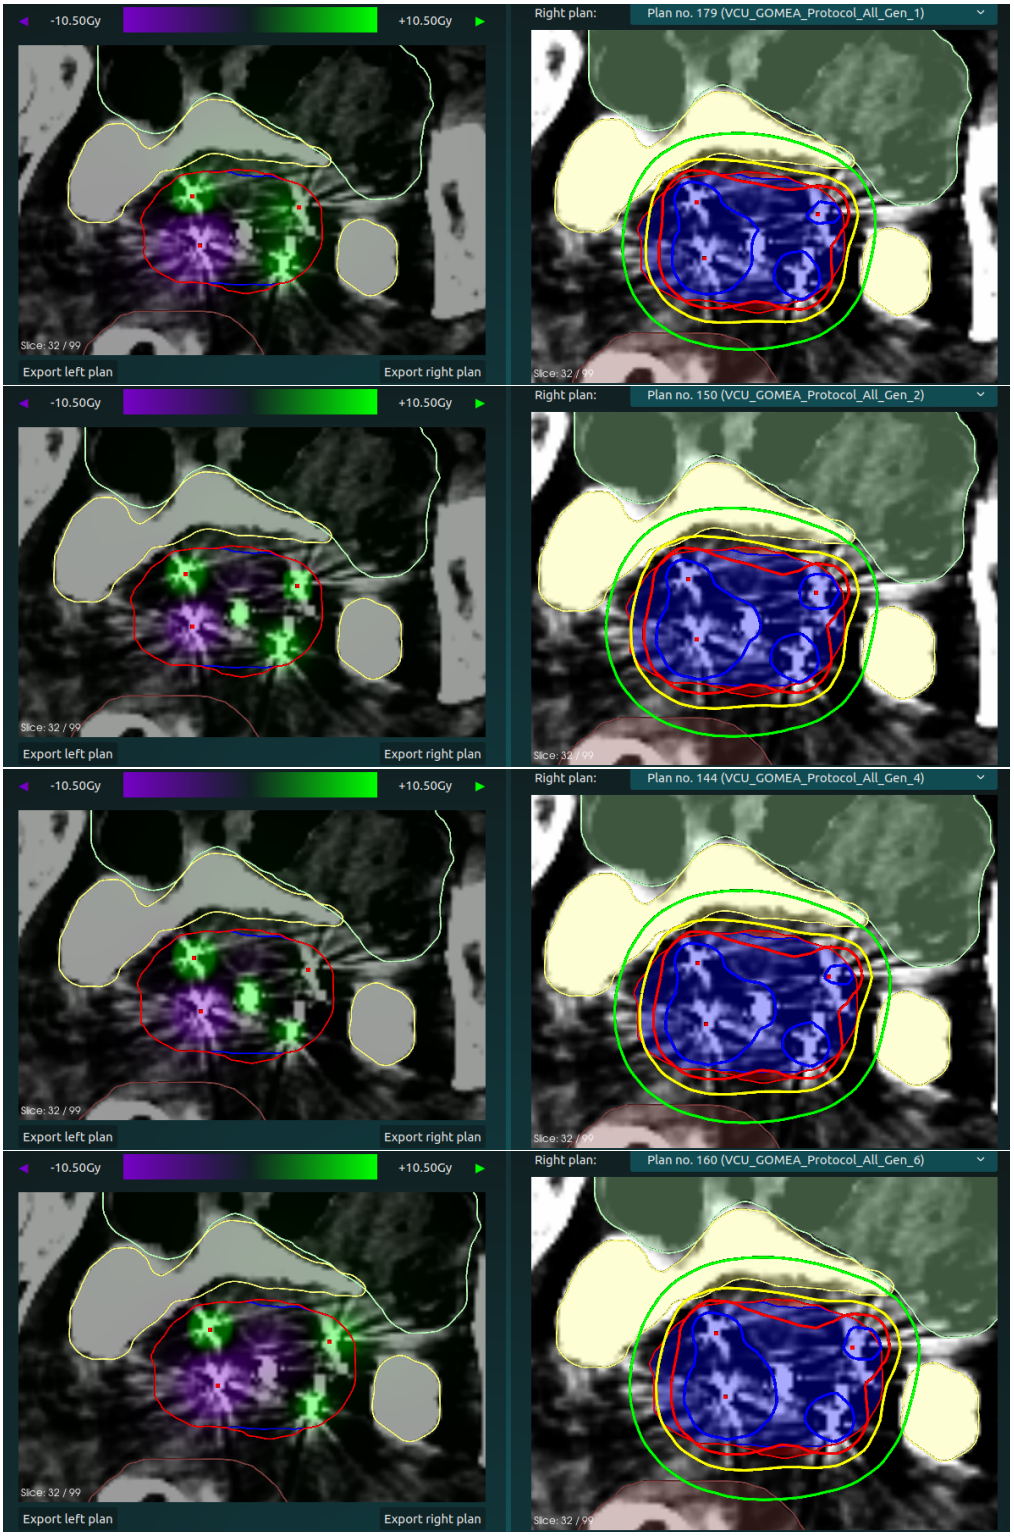

Figure 27: Patient 4, CT Slice 32

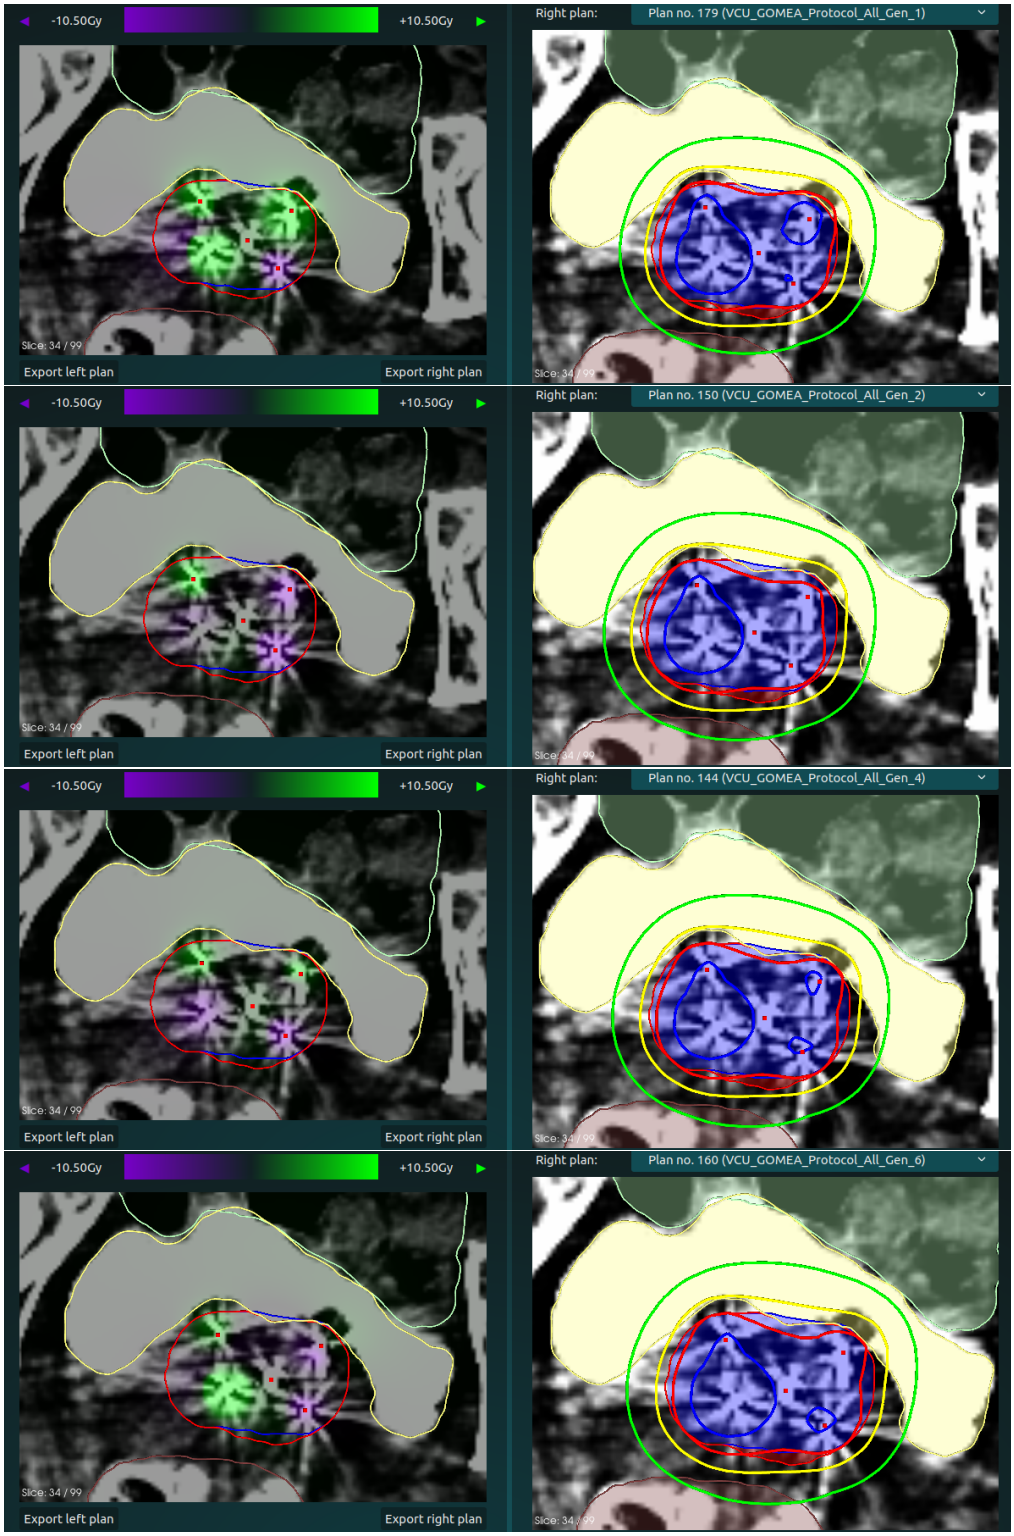

Figure 28: Patient 4, CT Slice 34

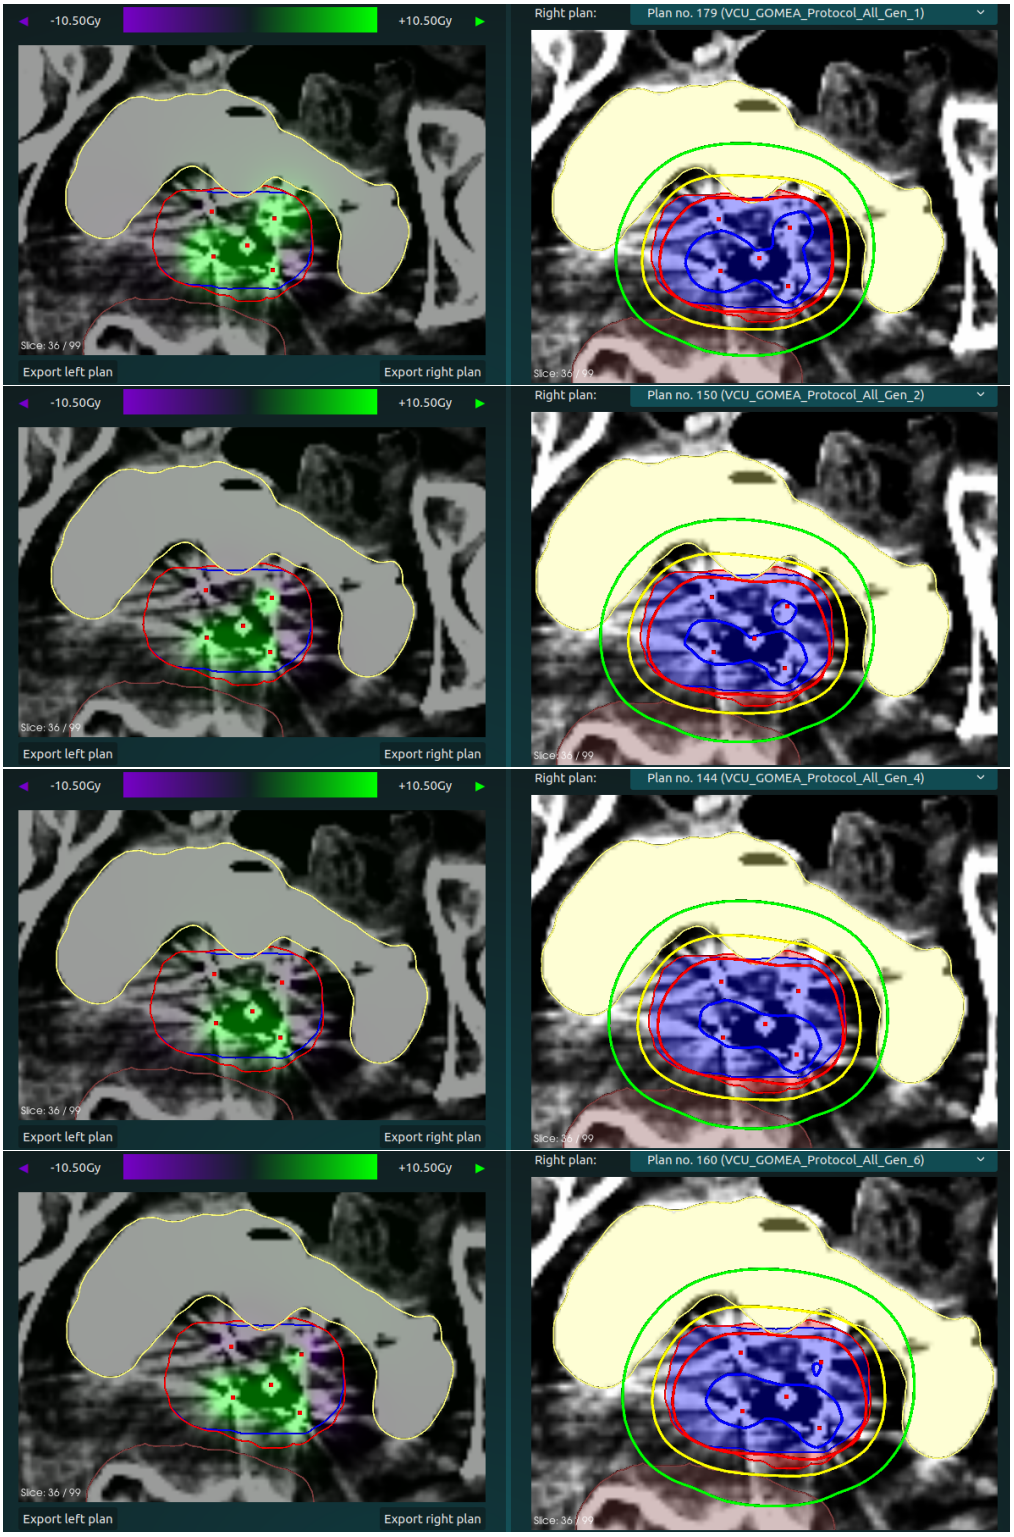

Figure 29: Patient 4, CT Slice 36

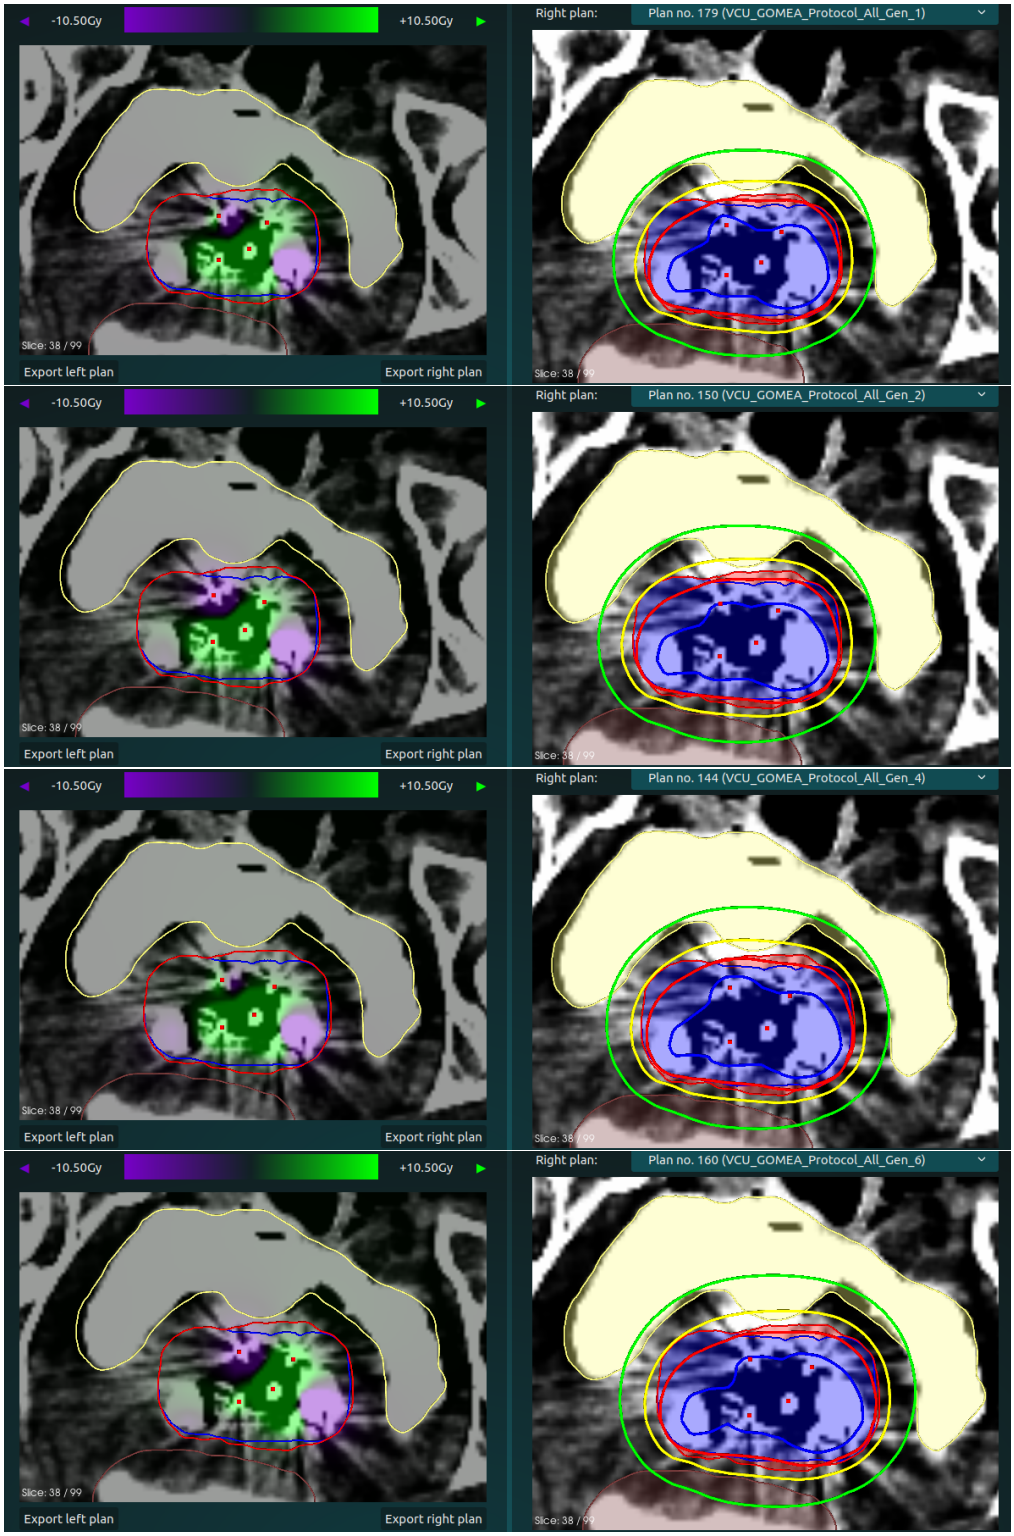

Figure 30: Patient 4, CT Slice 38

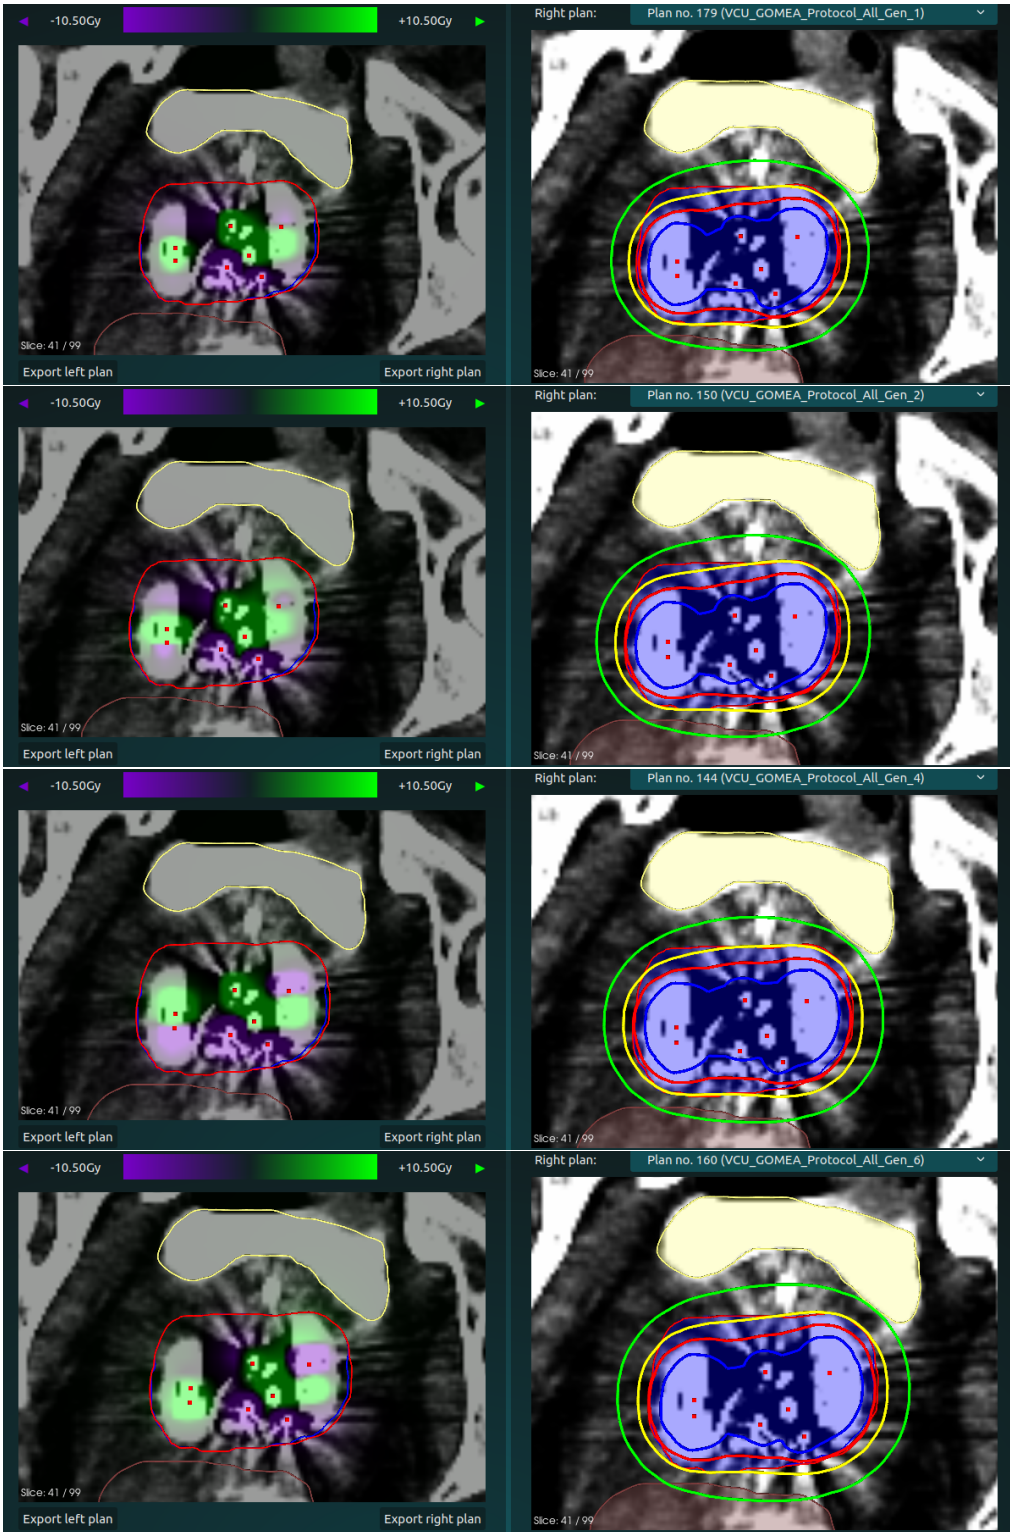

Figure 31: Patient 4, CT Slice 41

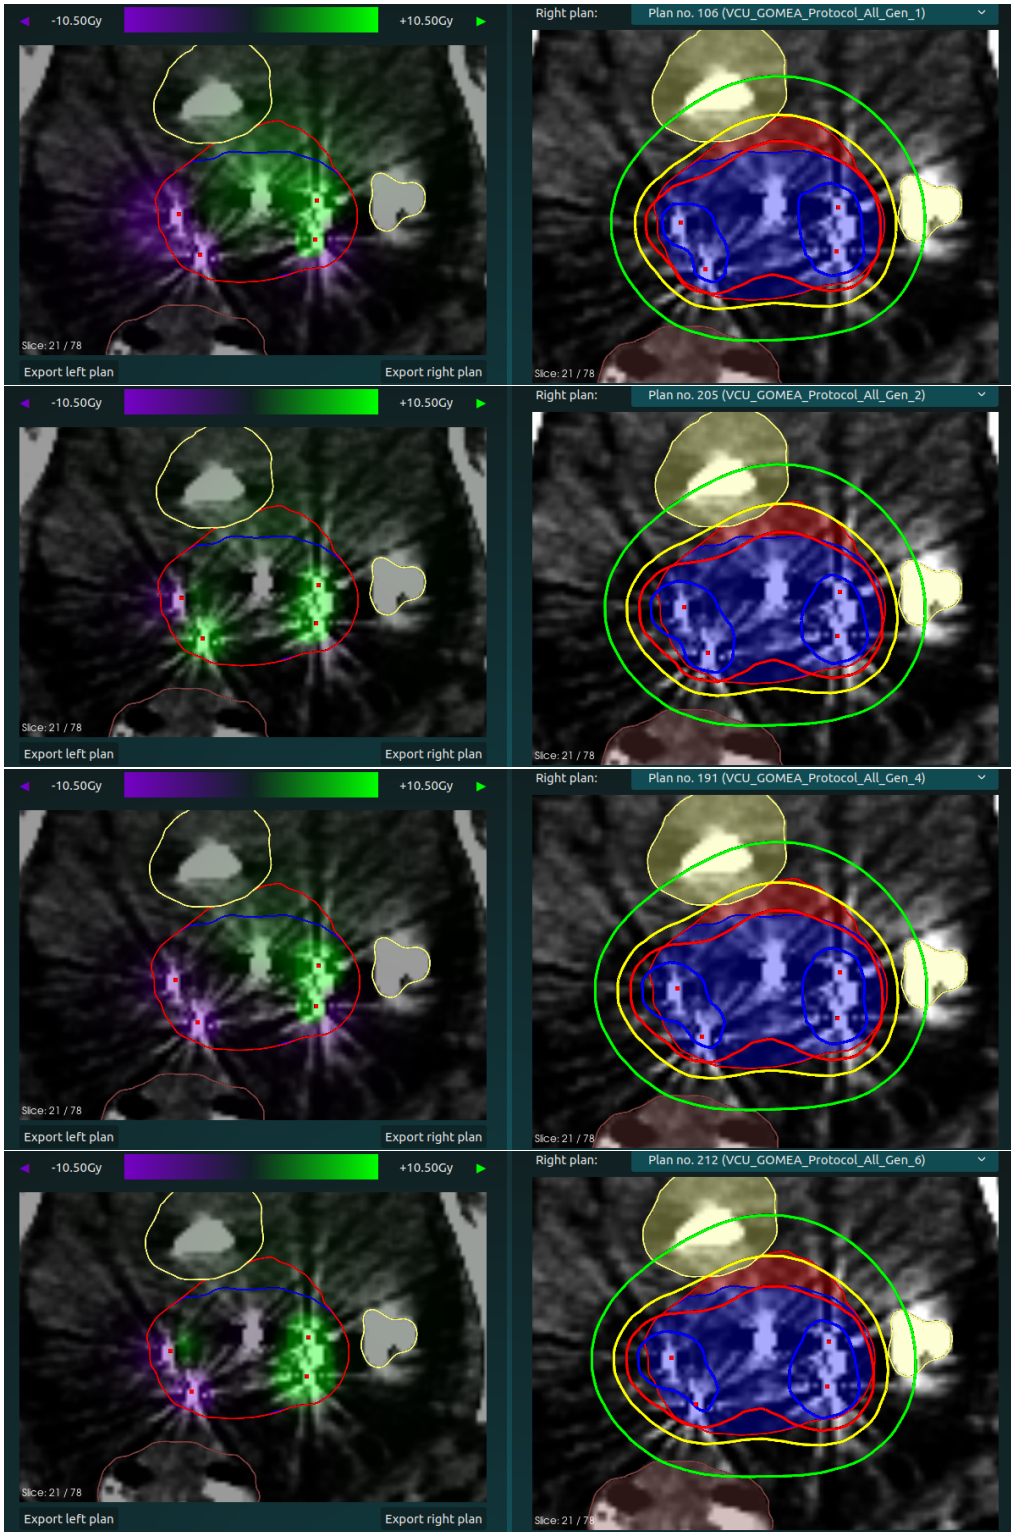

Figure 32: Patient 5, CT Slice 21

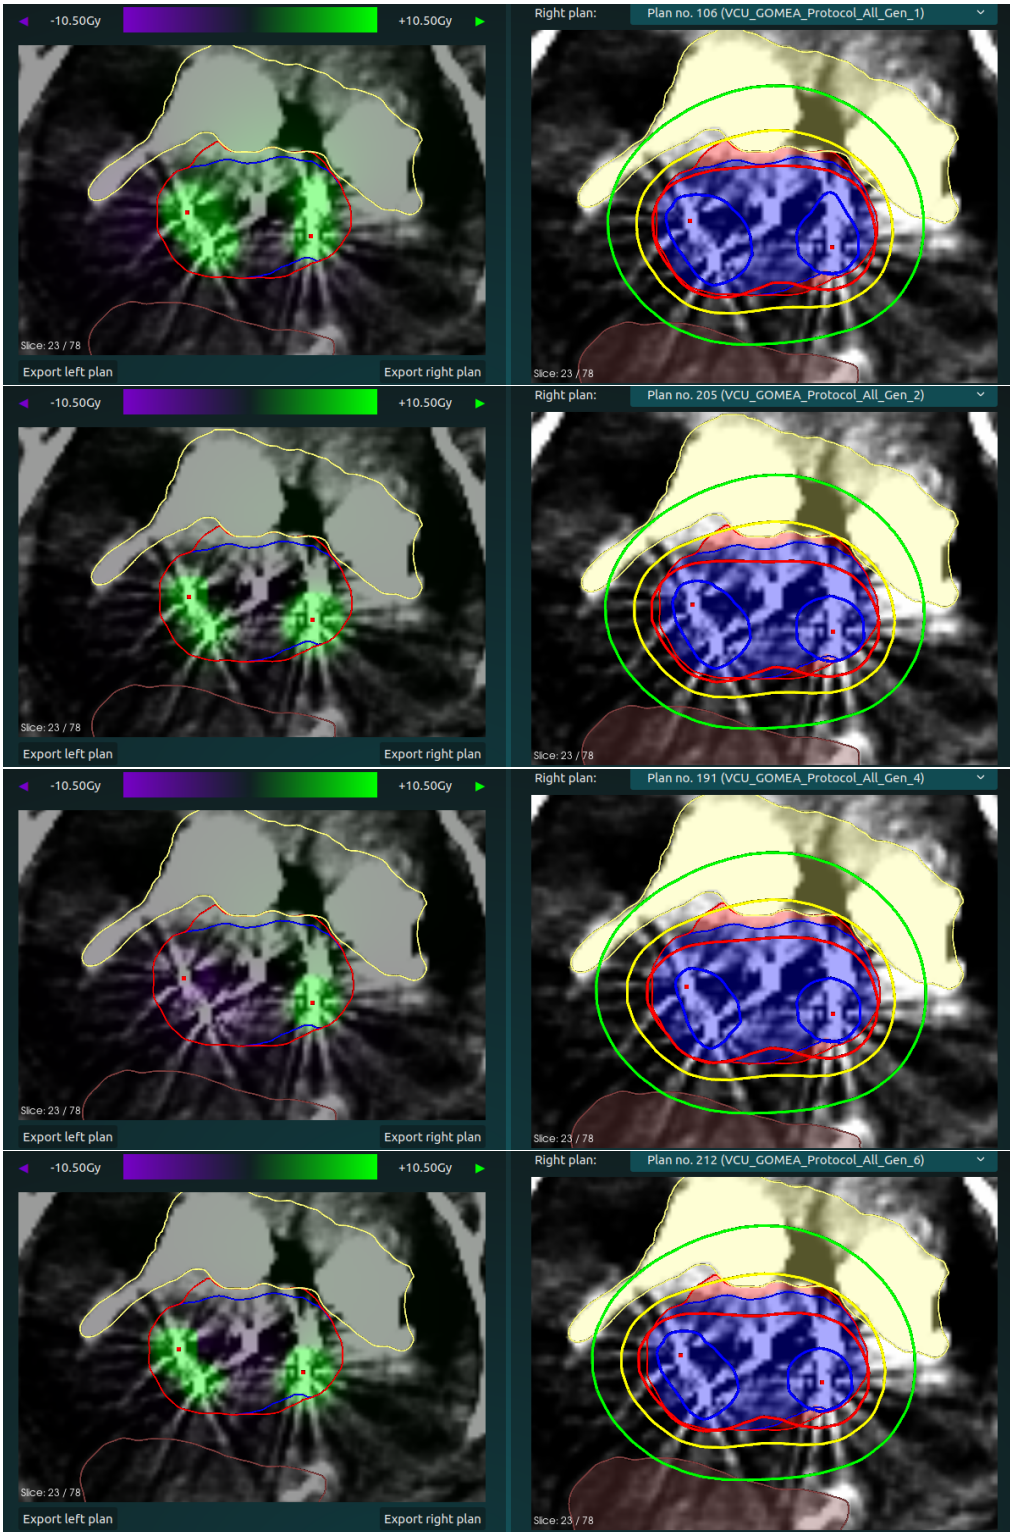

Figure 33: Patient 5, CT Slice 23

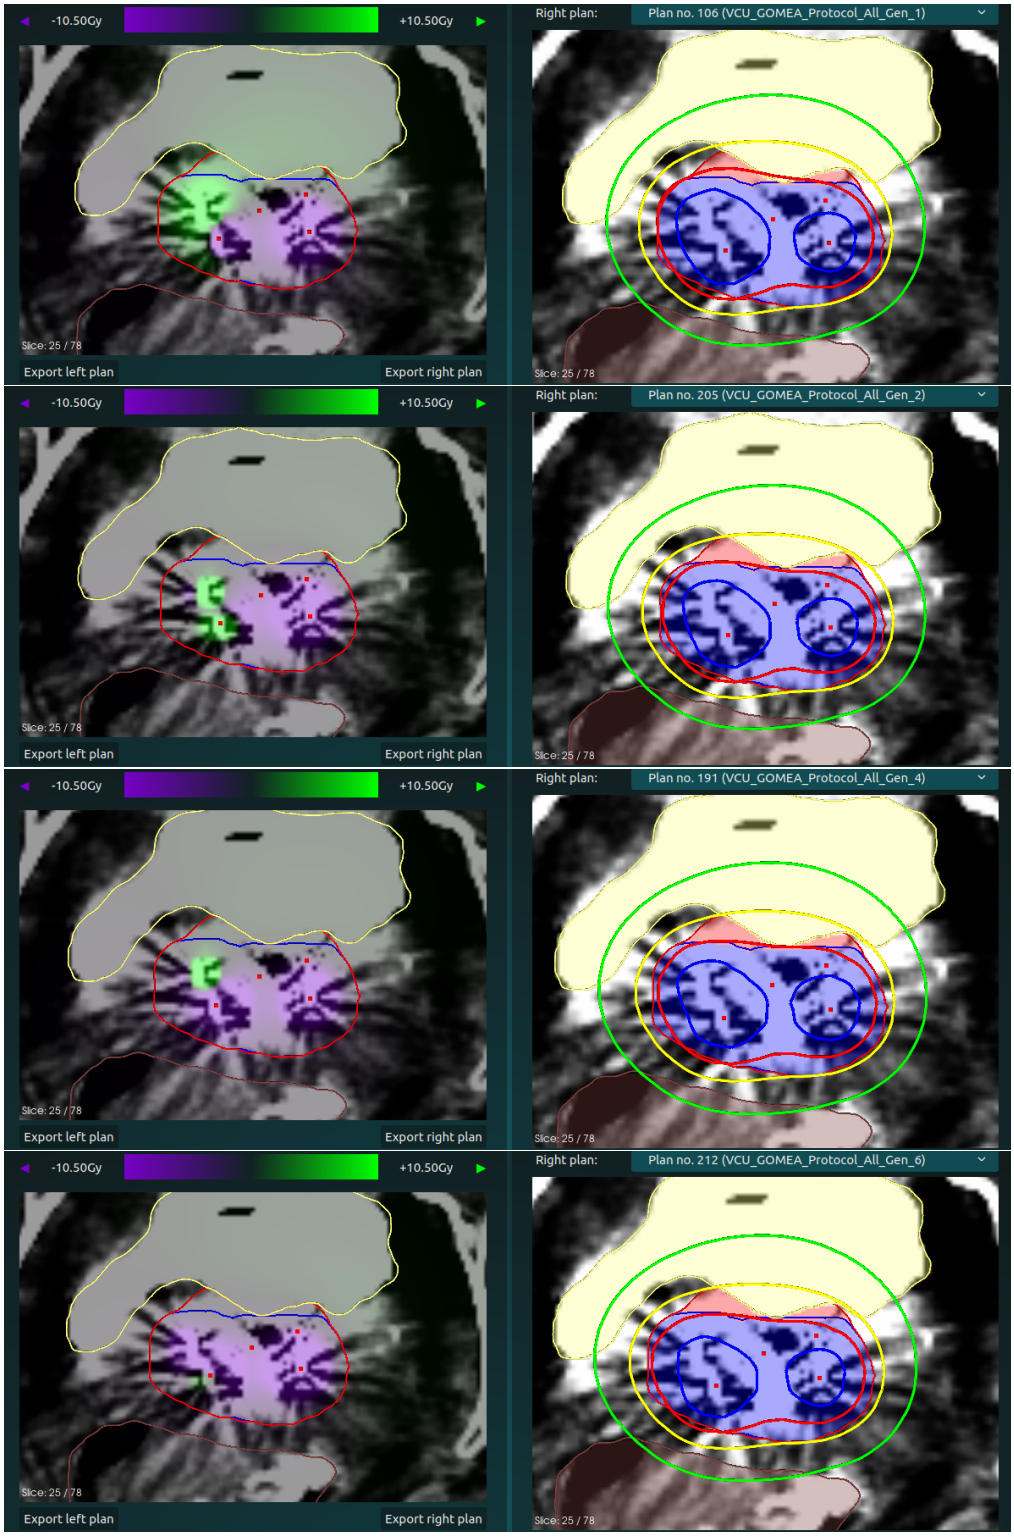

Figure 34: Patient 5, CT Slice 25

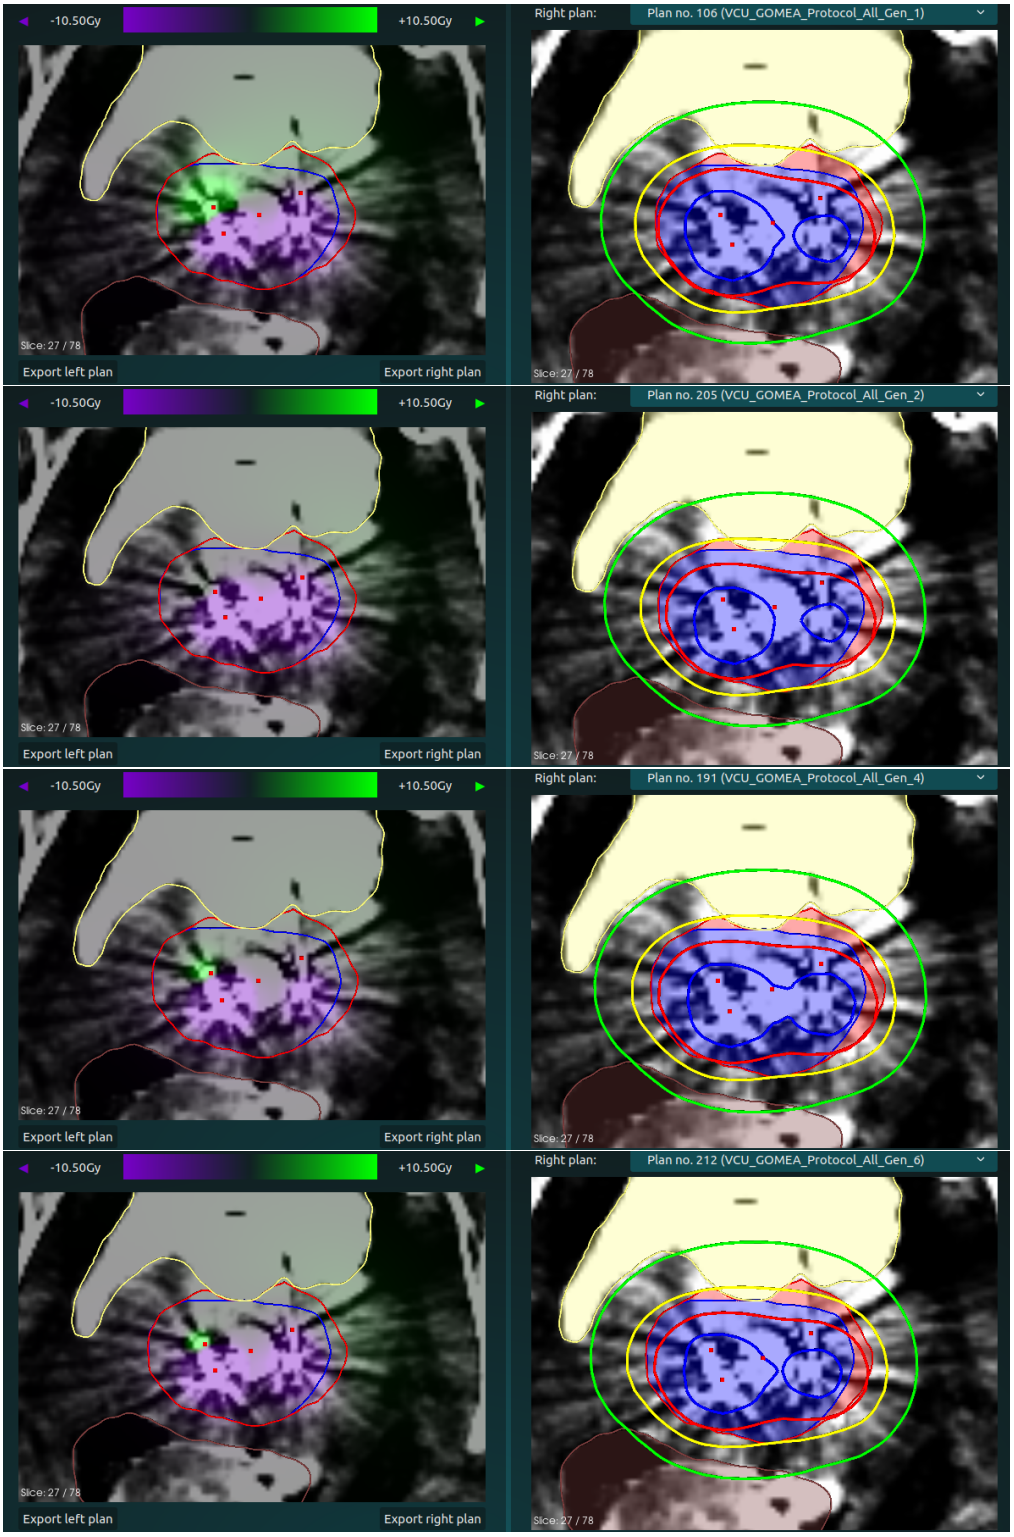

Figure 35: Patient 5, CT Slice 27

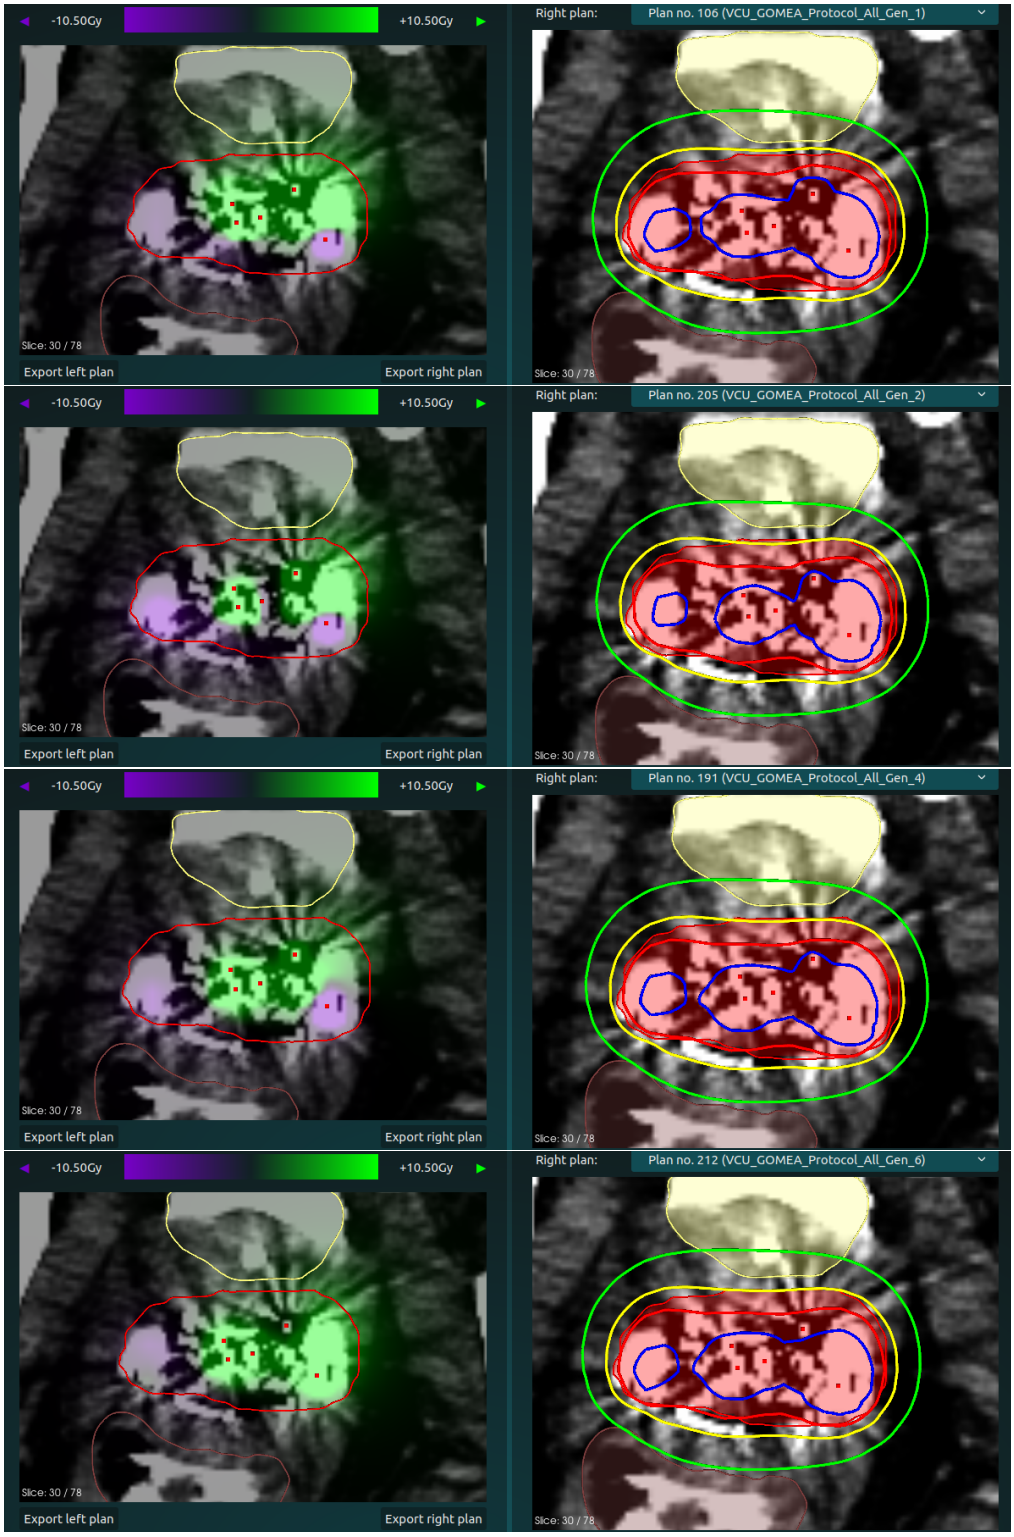

Figure 36: Patient 5, CT Slice 30

Figure 37: Patient 6, CT Slice 26

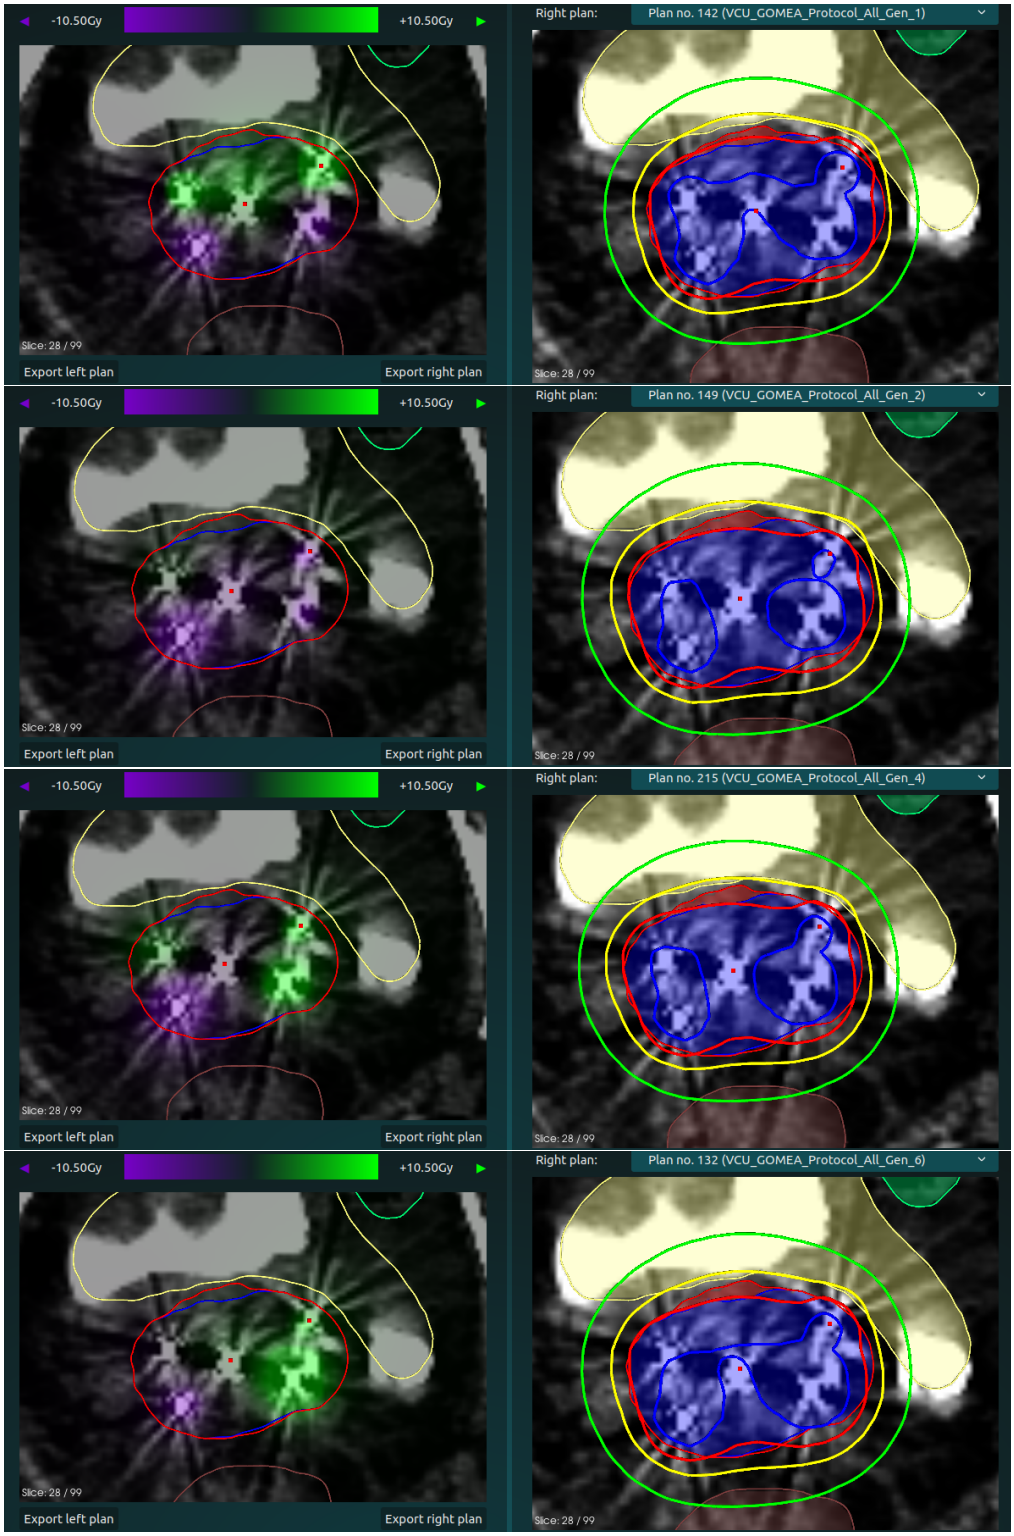

Figure 38: Patient 6, CT Slice 28

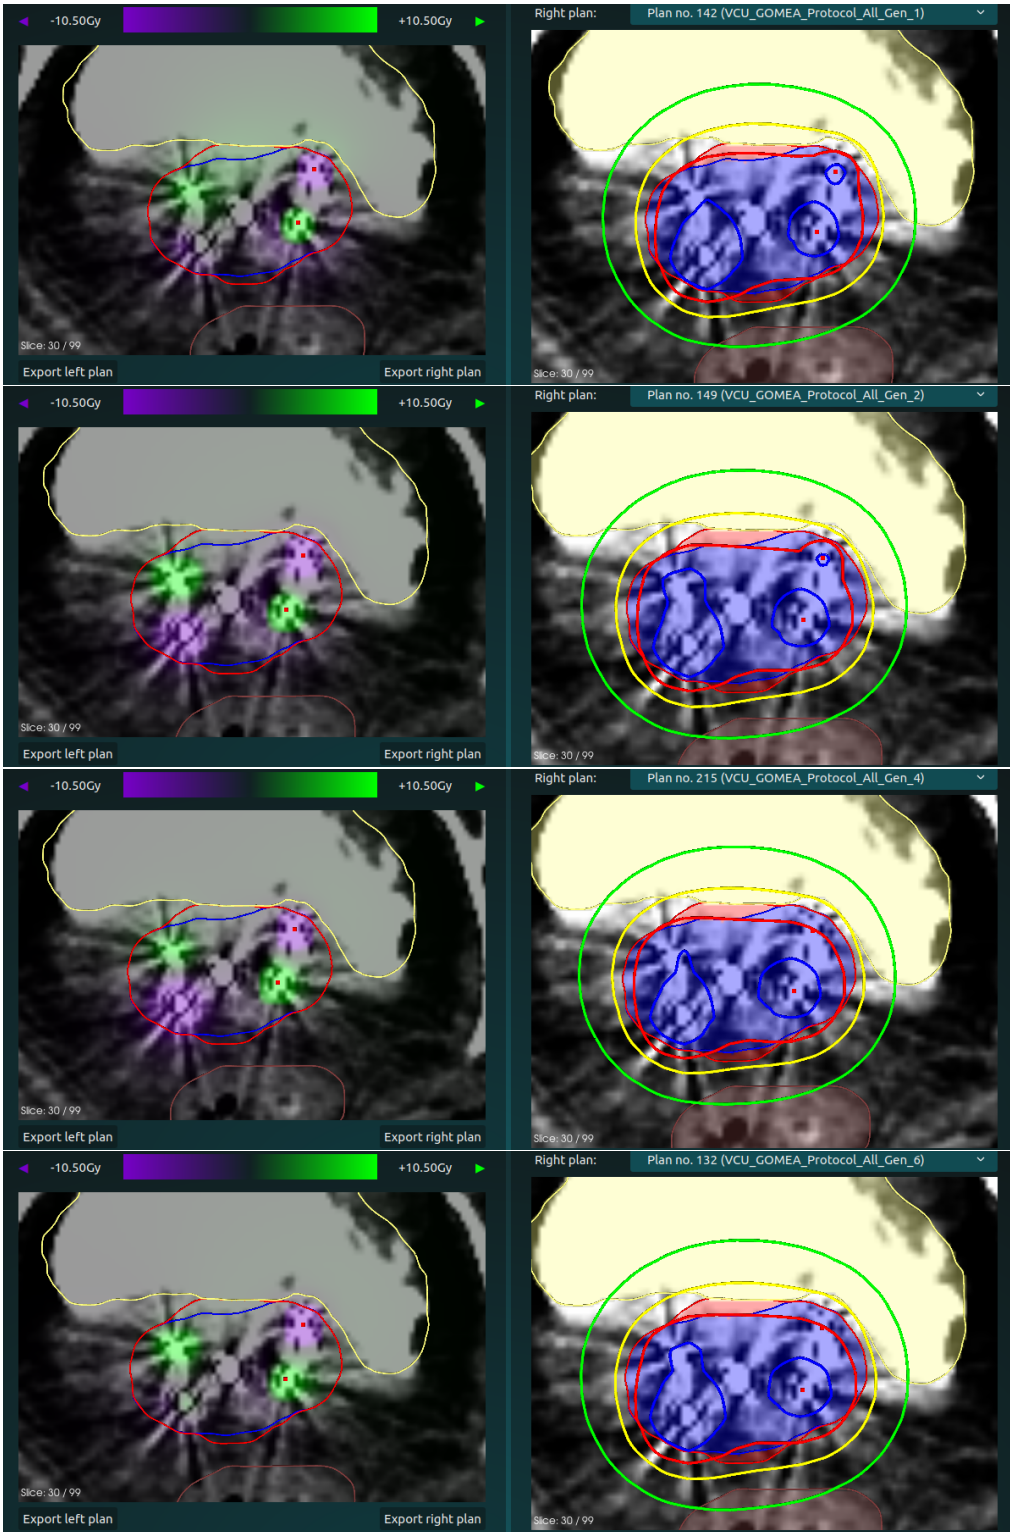

Figure 39: Patient 6, CT Slice 30

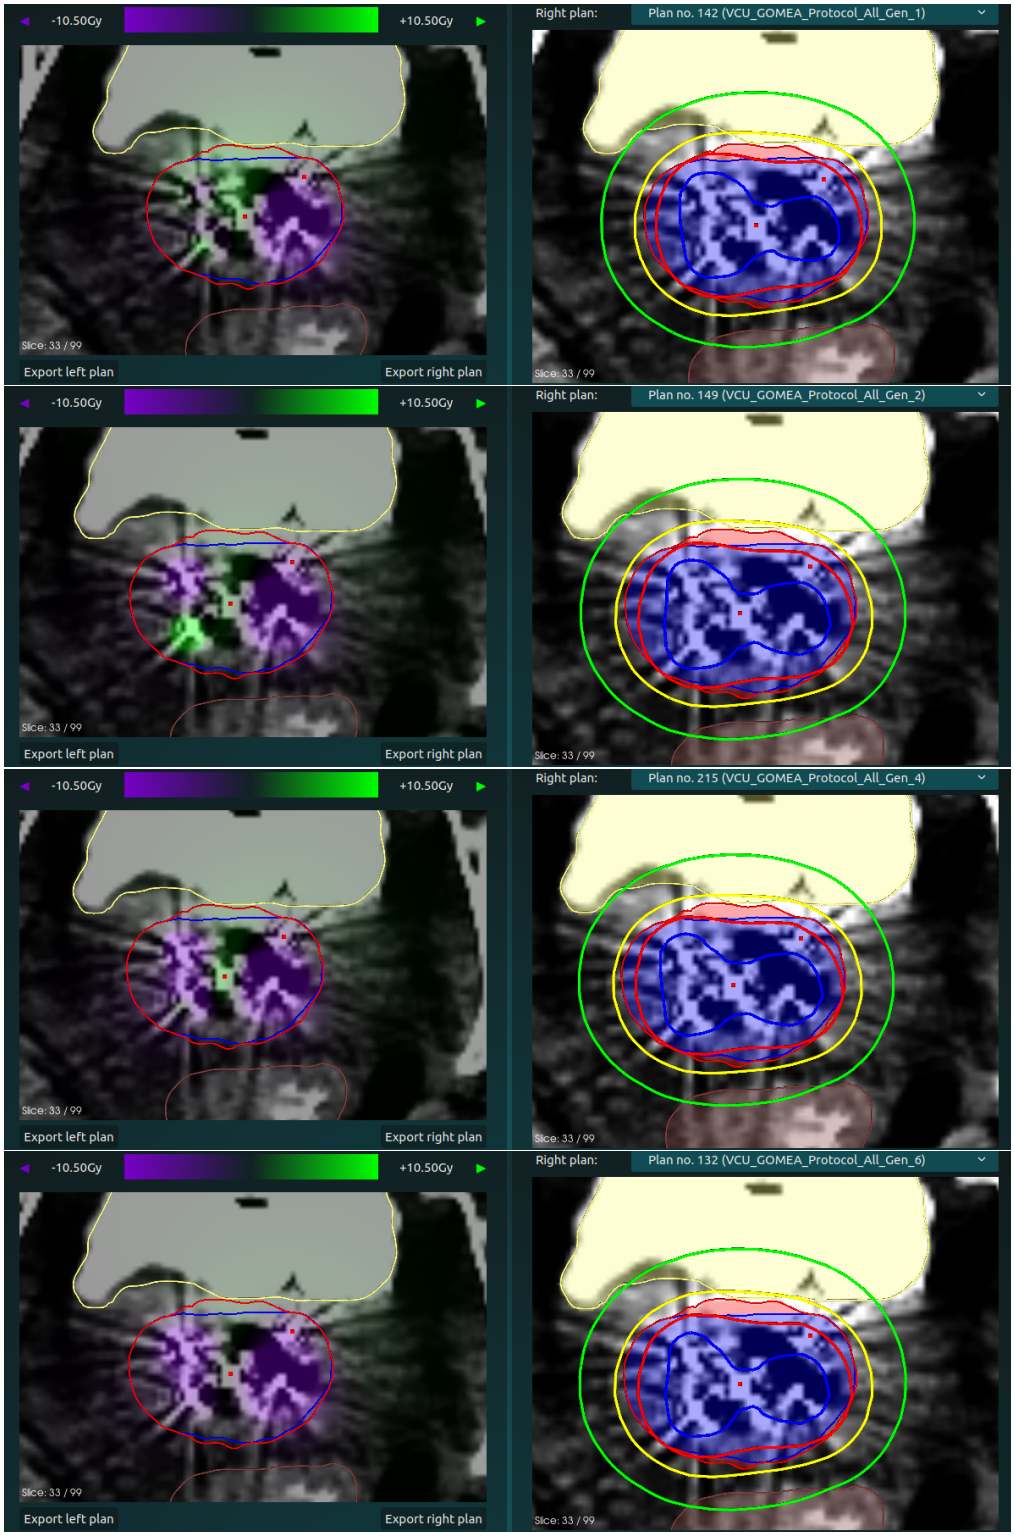

Figure 40: Patient 6, CT Slice 33

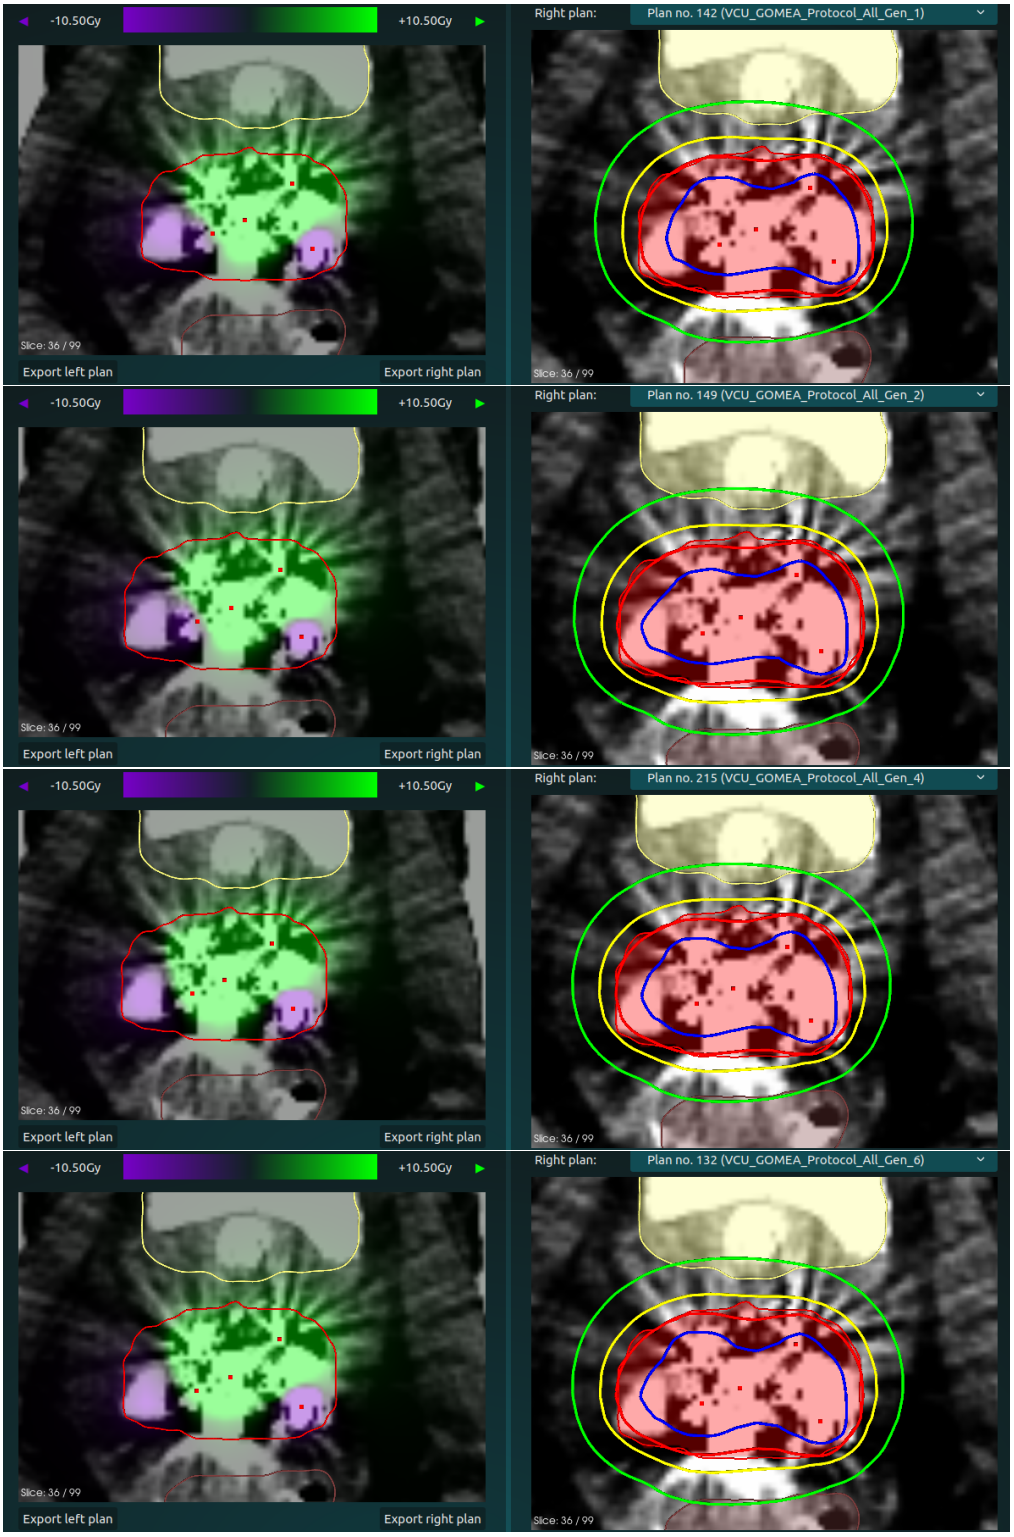

Figure 41: Patient 6, CT Slice 36

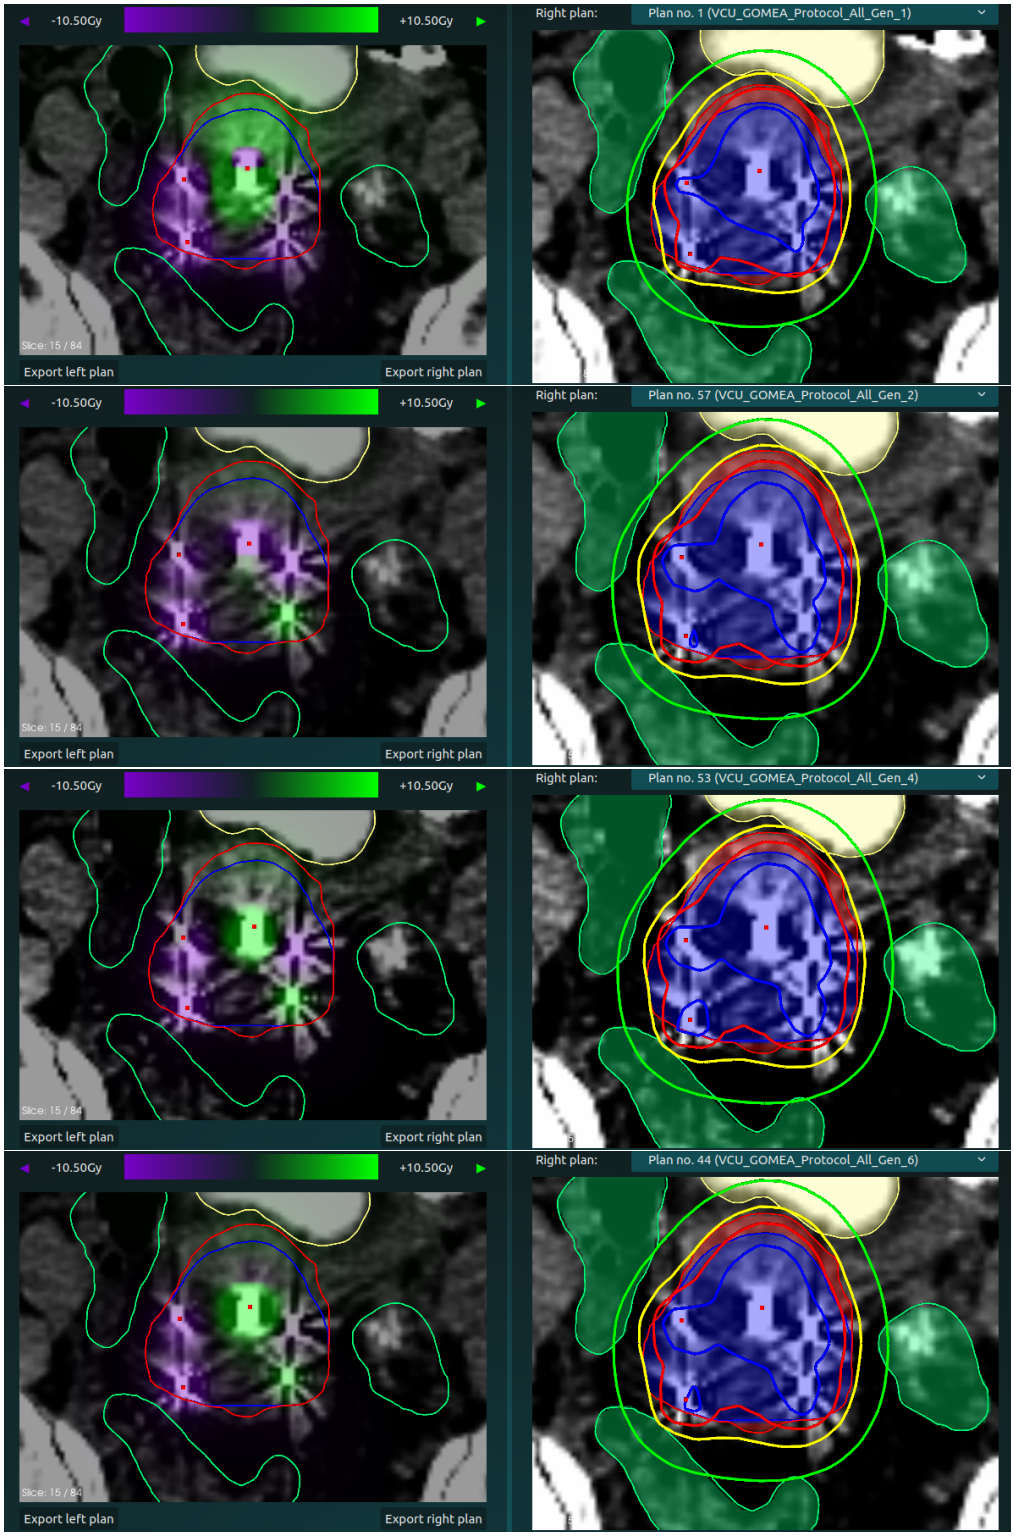

Figure 42: Validation Patient, CT Slice 15

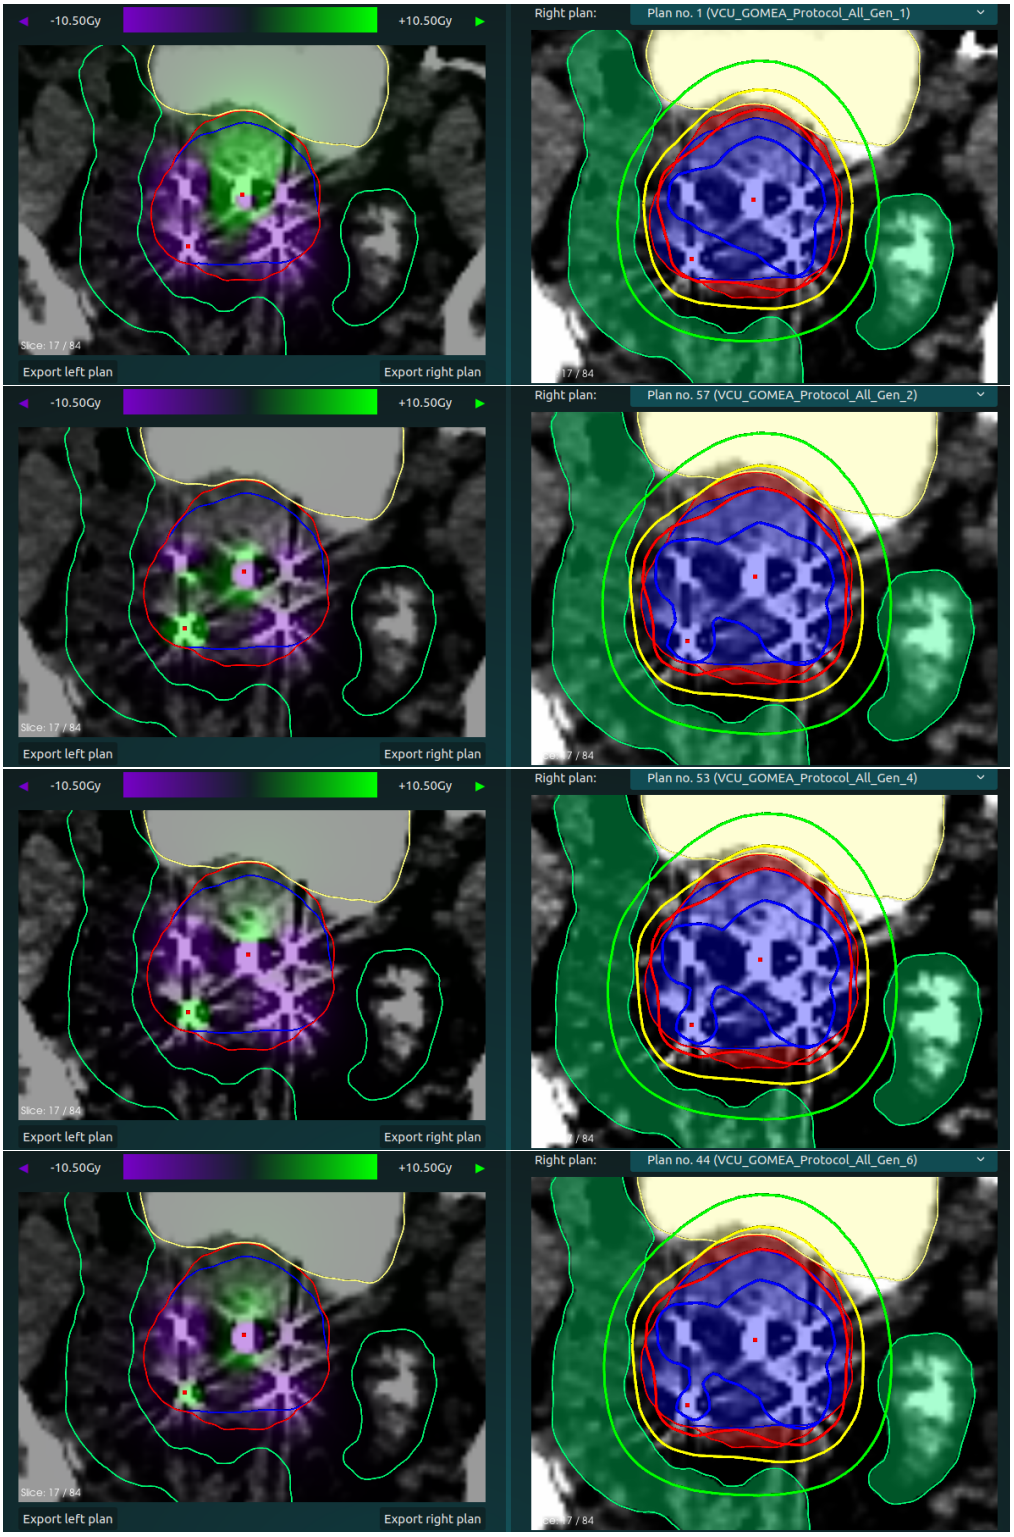

Figure 43: Validation Patient, CT Slice 17

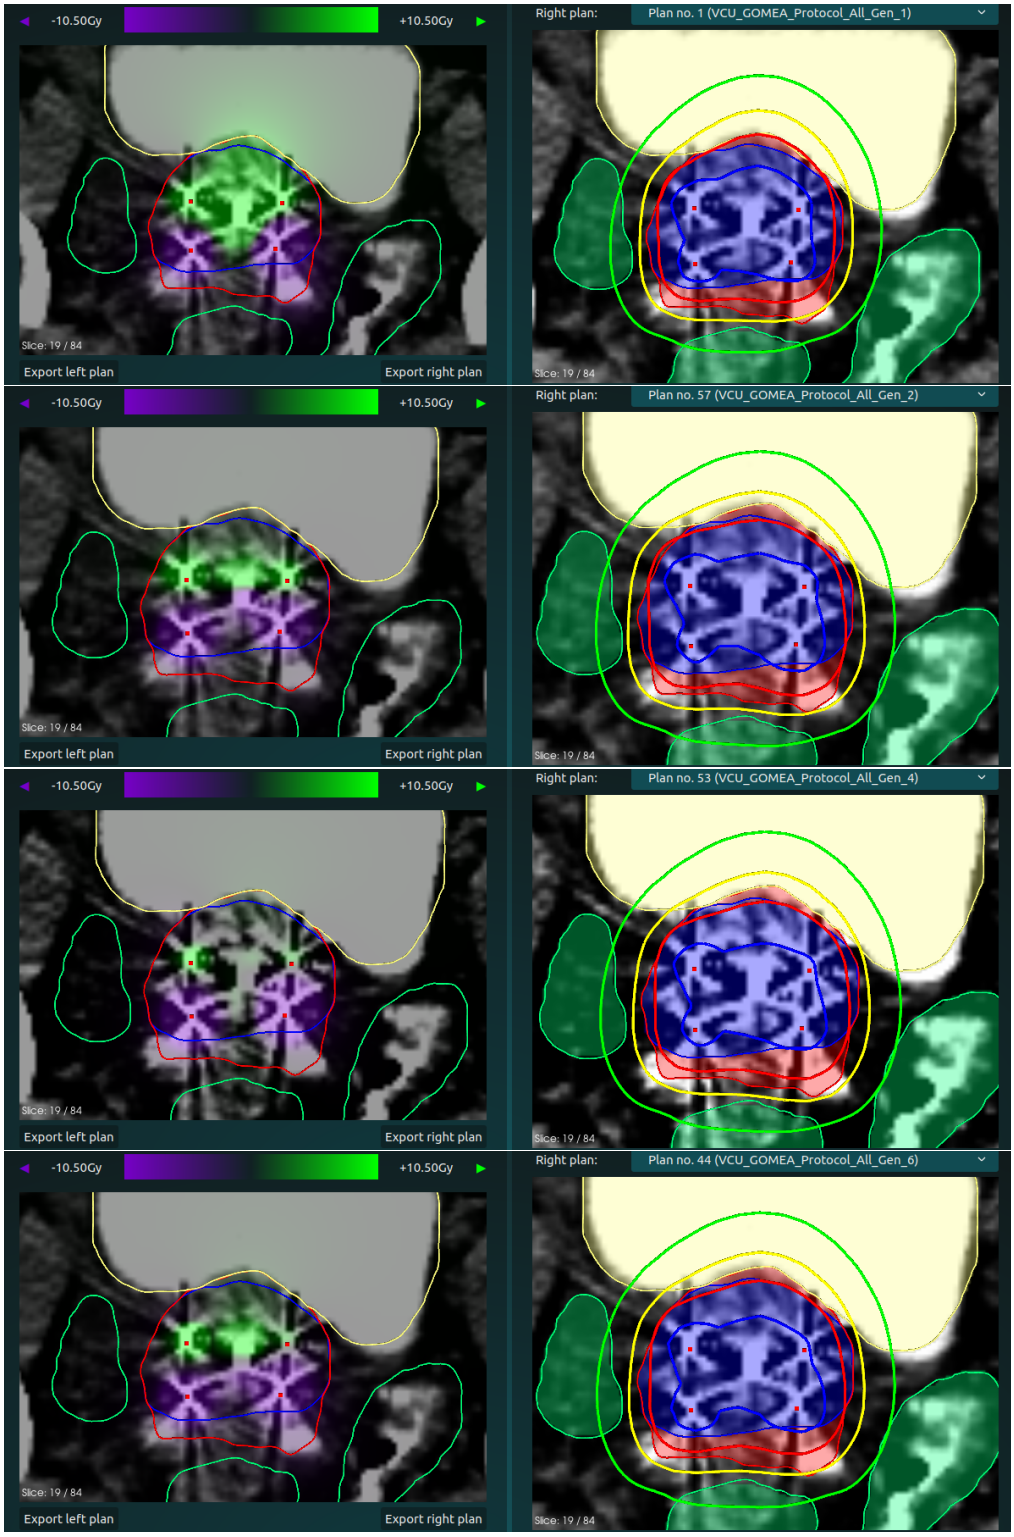

Figure 44: Validation Patient, CT Slice 19

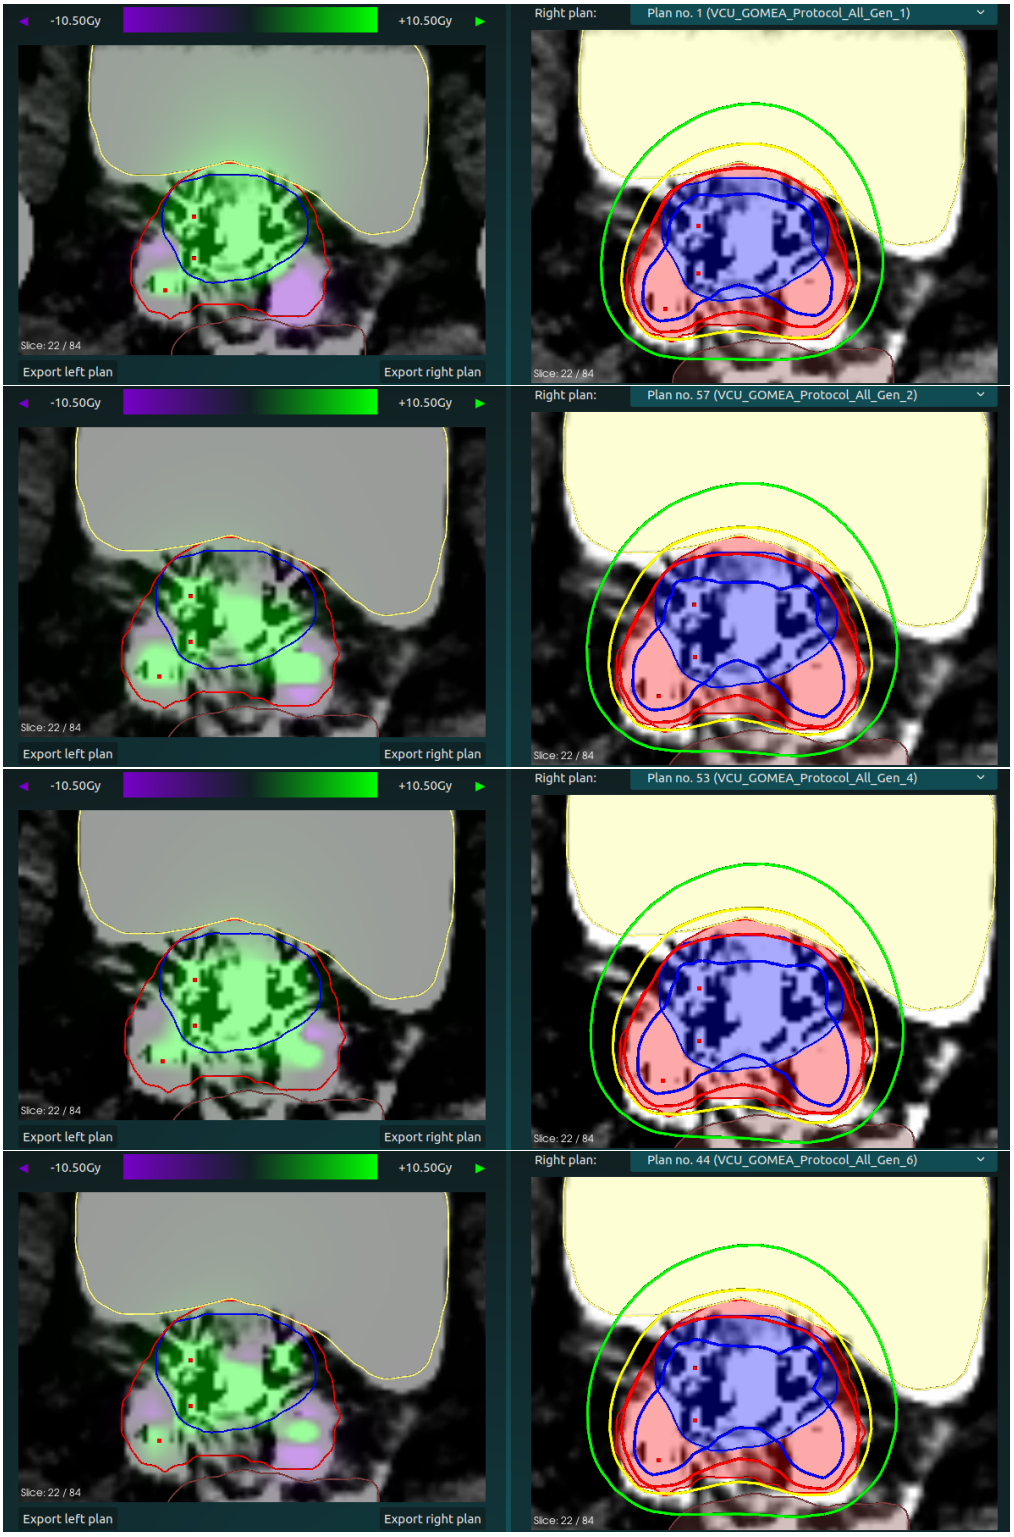

Figure 45: Validation Patient, CT Slice 22

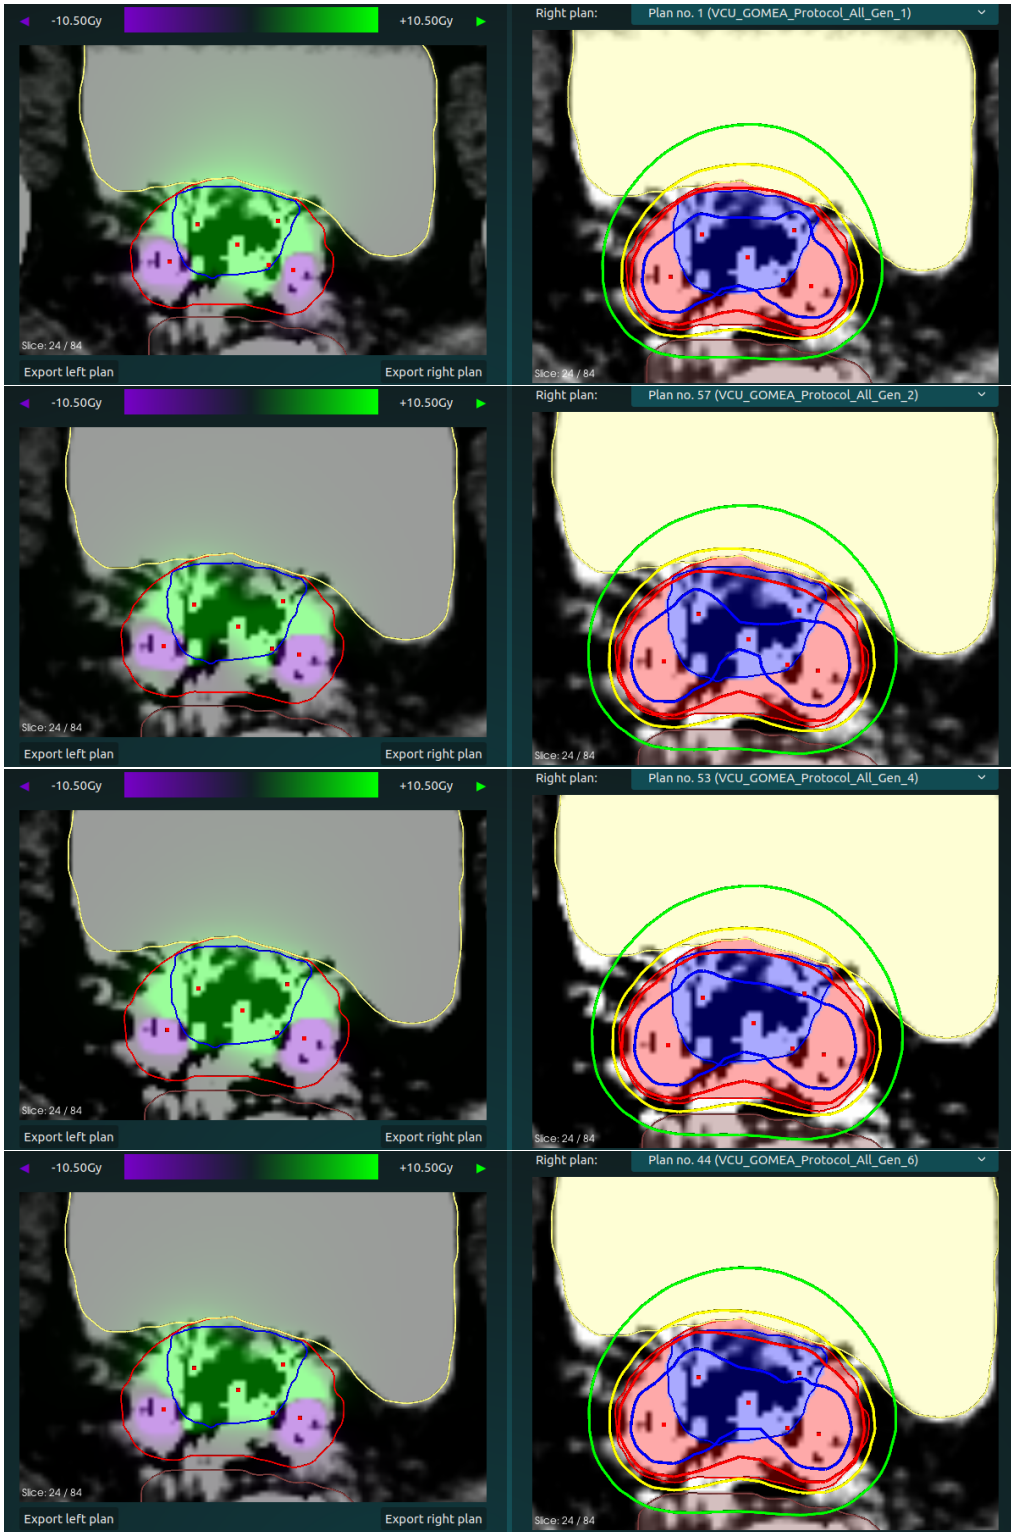

Figure 46: Validation Patient, CT Slice 24