

Bi-Level Optimization to Enhance Intensity Modulated Radiation Therapy Planning

Juan José MORENO^{1,*}, Savíns PUERTAS-MARTÍN^{1,2},
Juana L. REDONDO¹, Pilar M. ORTIGOSA¹, Anna ZAWADZKA³,
Paweł KUKOŁOWICZ³, Robert SZMURŁO⁴, Ignacy KALISZEWSKI⁵,
Janusz MIROFORIDIS⁵, Ester M. GARZÓN¹

¹ *Department of Informatics, University of Almería, Spain*

² *Information School, University of Sheffield, United Kingdom*

³ *Department of Medicine Physics, Maria Skłodowska Curie Memorial Cancer Centre and Institute of Oncology, Poland*

⁴ *Faculty of Electrical Engineering, Warsaw University of Technology, Poland*

⁵ *Systems Research Institute, Polish Academy of Sciences, Poland*

e-mail: juanjomoreno@ual.es, savinspm@ual.es, jfredondo@ual.es, ortigosa@ual.es, anna.zawadzka.fizyk@nio.gov.pl, Pawel.Kukolowicz@nio.gov.pl, robert.szmurlo@pw.edu.pl, ignacy.kaliszewski@ibspan.waw.pl, janusz.miroforidis@ibspan.waw.pl, gmartin@ual.es

Received: November 2023; accepted: June 2024

Abstract. Intensity Modulated Radiation Therapy is an effective cancer treatment. Models based on the Generalized Equivalent Uniform Dose (gEUD) provide radiation plans with excellent planning target volume coverage and low radiation for organs at risk. However, manual adjustment of the parameters involved in gEUD is required to ensure that the plans meet patient-specific physical restrictions. This paper proposes a radiotherapy planning methodology based on bi-level optimization. We evaluated the proposed scheme in a real patient and compared the resulting irradiation plans with those prepared by clinical planners in hospital devices. The results in terms of efficiency and effectiveness are promising.

Key words: Intensity Modulated Radiation Therapy (IMRT), Genetic Algorithms, Generalized Equivalent Uniform Dose (gEUD), Multi-Objective Optimization, Bi-Level Optimization.

1. Introduction

Intensity Modulated Radiation Therapy (IMRT) is an effective radiation therapy technique in cancer treatment. It requires personalized irradiation plans (RT plans) that specify 3D dose distributions which effectively destroy cancer cells while minimizing side effects on healthy tissues. Subsequently, linear accelerators equipped with multileaf collimators deliver radiation beams to patients from several fixed angles. Each beam is decomposed into a regular grid of (thousands of) beamlets with varying intensities that can be individually controlled. Therefore, every RT plan is defined by the specific intensities of all the

*Corresponding author.

beamlets over all beams. When preparing an RT plan, the goal is twofold. On the one hand, plans are sought that ensure depositing the prescribed doses in planning target volumes (PTVs), and, on the other hand, such plans should protect organs at risk (OARs). These two goals are contradictory and, as such, must be traded-off, which means that high quality RT plans, i.e. effective and at the same time causing as little harmful side effects as possible, require significant efforts from the RT planners (Breedveld *et al.*, 2019). This contradiction leads the way to multi-objective optimization in IMRT.

To obtain effective RT plans, various multi-objective optimization models with several radiation effect measures (or criteria, at least one measure for each PTV and OAR), physical or biological, have been proposed (Breedveld *et al.*, 2019; Cho, 2018; Ehr Gott *et al.*, 2010). Physical measures capture characteristics of dose distributions. Usually, in PTVs, physical measures take the form of sums of absolute or squared values of the difference between the delivered and prescribed doses. On the other hand, biological measures refer to (estimated) biological effects of radiation in organs (Ólafsson *et al.*, 2005). Doses delivered to OARs are controlled by imposing physical constraints on acceptable average or maximal radiation doses per voxel (a voxel being a 3D box constituting the irradiation planning mesh). In PTVs and OARs, the appropriateness of dose distributions is also assessed by the shapes of dose-volume histograms (DVHs). Favourable DVH shapes are enforced by appropriate dose-volume constraints, which as a rule rend the optimization problems nonconvex (Bortfeld, 1999). Given the multiplicity of beams, beamlets, and voxels, the resulting models are computationally complex. The most handleable are linear or linearized models (Romeijn *et al.*, 2006; Ólafsson and Wright, 2006). To handle nonconvex constraints, in Breedveld and Heijmen (2017), the interior point method has been adapted. Another approach to cope with nonconvexity is to resort to heuristics, or a combination of heuristic and exact methods (Yihua Lan *et al.*, 2006). Stochastic approaches are an alternative (Ahmad and Bergen, 2010; Moreno *et al.*, 2021). Another option is the convex formulation of dose-volume histograms constraints (Fu *et al.*, 2019; Romeijn *et al.*, 2004). Finally, another possibility is to establish two types of dose-volume constraints, namely hard and soft constraints, and define a solution of two stages, the first more general and the second focused on the subthreshold voxels (Mukherjee *et al.*, 2020).

Nowadays, capitalizing on the progress in oncology research, biologically-motivated measures of radiation effects play an important role in clinical practice. Despite the difficulties in their implementation for clinical use, they are recommended by the ICRU.¹ Reports as an advanced level of treatment reporting, incorporating evolving concepts. Measures of dose distribution effectiveness, such as the tumour control probability (TCP) and the normal tissue complication probability (NTCP) (Alber and Nüssli, 1999), the biologically effective dose (BED) (Saberian *et al.*, 2015), and the generalized equivalent uniform dose (gEUD) (Breedveld *et al.*, 2019; Niemierko, 1996; Wu *et al.*, 2002), are gaining popularity. The gEUD metric, based on the linear-quadratic cell survival model, has been reported to be the most relevant for radiotherapy (Niemierko, 1996; Wu *et al.*, 2002; Ólafsson *et al.*, 2005). Therefore, it is the focus of this work.

¹International Committee for Radiological Units.

For gEUD-based RT planning optimization, optimal solutions can be efficiently computed by gradient methods, as proposed in Choi and Deasy (2002). However, for every PTV and OAR several parameters need to be set by the RT planner. Usually, the planner starts from a set of parameters recommended in the literature and later adjusts them to the particular patient case by trial and error, to satisfy the imposed physical constraints. The quality of the final RT plan depends heavily on the knowledge and skills of the planner.

Anyway, gathering all clinical prescriptions in an optimization model to obtain automatic planning is a challenge. Recent studies propose the integration of various models to improve the quality of the computed optimal plans. For example, Huang *et al.* (2022) computes the plans by performing meta-optimization of treatment planning hyperparameters, and a meta-scoring is executed by constructing a tier list of the relevant criteria (e.g. dose homogeneity, conformity, spillage, and OAR sparing). Wang *et al.* (2023b) explores the use of Bayesian optimization methods to optimize the adjustable planning parameters included in both the dose objectives and their corresponding weights. So Bayesian optimization is used to solve a problem of a reduced dimension focused on the planning parameters. In Wang *et al.* (2023a), a new approach for automatic planning based on an integrated deep reinforcement learning-based framework for the IMRT technique has been proposed. In this approach, a hybrid objective function has been proposed that combine gEUD and physical dose constraints, and the trained artificial neural network can mimic the actions of the physician during optimization and adjust the plan parameters to obtain high-quality plans. In Won *et al.* (2021), a dose volume-gEUD-based optimization method to hippocampus-sparing whole-brain radiation therapy using volumetric modulated arc therapy has been applied and evaluated. The authors conclude that the dose volume-gEUD optimization plan showed a better dosimetric profile compared to the dose volume-based optimization one. However, in that work, the value of a parameter shaping the gEUD function for the hippocampus is selected by inspecting ten integer values in the range [1, 40] before the optimization step begins.

In our work, we try to fill this gap by automating the tuning of the gEUD parameters. We present a novel bi-level optimization model that integrates the gEUD metric-based model with the dose-volume constraints. In this work, we introduce two main contributions:

Firstly, we propose an RT planning methodology based on a bi-level optimization scheme. On the lower level, a gEUD-based objective function with all its parameters fixed, including the gEUD parameters, is optimized using an (exact) gradient algorithm. On the upper level, these parameters are optimized using an evolutionary algorithm. The output of this scheme is a collection of non-dominated RT plans, each representing a solution to the multi-objective optimization problem, with varying priorities.

Secondly, to facilitate the analysis of these RT plans, we have developed a decision tool. This tool assists the planner in selecting the final plan by providing comprehensive information and comparisons between different plans, helping to make decisions.

The rest of the paper is structured as follows: In Section 2, we present the proposed methodology for generating IMRT plans. Section 3 provides an explanation of PersEUD, the system we have developed and implemented for delivering IMRT RT plans. In Sec-

tion 4, we present the context and main results obtained from three RT plans using real patient data. Finally, Section 5 summarizes the conclusions drawn from our study.

2. Computational Methods and Theory: IMRT Planning Based on the gEUD and Physical Constraints

This section describes the proposed methodology for obtaining IMRT plans by combining gEUD and physical constraints based on oncology physicians' recommendations and individual patient anatomy.

To facilitate the reading, the notation used throughout the paper is summarized in the Table 1.

2.1. gEUD-Based IMRT Planning

RT planning can be viewed as an optimization process, in which the fluence maps, represented by vectors of non-negative numbers x , define the radiation intensities of individual beamlets. Deposition matrices D translate fluencies x to doses deposited in voxels, so that the doses in voxels are computed as product Dx . The RT planning goal is to compute fluencies x that deposit prescribed and homogeneous doses to PTVs and acceptable doses to OARs.

gEUD is a biology-motivated measure to evaluate radiation effects, based on the concept of the uniform radiation dose delivered to a patient organ, that causes the same effect as a nonuniform dose (Niemierko, 1996; Wu *et al.*, 2002). In the case of gEUD, radiation effects in a PTV or an OAR, both referred to as structure s , are evaluated by the following function that aggregates these effects over all voxels belonging to the structure s :

$$gEUD_s(x, a_s) = \left(\frac{1}{|M_s|} \sum_{j \in M_s} d_j(x)^{a_s} \right)^{\frac{1}{a_s}}, \quad (1)$$

Table 1
Notation and naming conventions.

Notation	Meaning
x	fluence map
D	deposition matrix (translates x to doses in voxels)
D_j	j -th row of D
$d(x) = Dx$	vector of doses deposited in voxels
$d_j(x) = D_j x$	dose deposited in voxel j
$T = \{t\}$	set of indices of PTVs
M_t	set of voxels in t -th PTV
EUD_t^0	prescribed uniform dose for t -th PTV
$R = \{r\}$	set of indices of OARs
M_r	set of voxels in r -th OAR
EUD_r^0	maximal uniform dose for r -th OAR
$P = \{p\} \subseteq R$	subset of high-priority protected OARs

where $|M_s|$ is the number of voxels of the structure s ; $d_j(x) = D_j x$ is the radiation dose deposited in voxel j of structure s by fluence map x and the parameter a_s represents the radiation effect on the structure. Its value can be taken from the literature or it can be adjusted individually by trial and error for each patient case.

According to Wu *et al.* (2002), clinically meaningful RT plans can be obtained by computing the maximum of the following function $F(x, \phi)$, built over the gEUD:

$$F(x, \phi) = \prod_{t \in T} \frac{1}{1 + \left(\frac{EUD_t^0}{gEUD_t(x, a_t)}\right)^{n_t}} \cdot \prod_{r \in R} \frac{1}{1 + \left(\frac{gEUD_r(x, a_r)}{EUD_r^0}\right)^{n_r}}, \quad (2)$$

where EUD_t^0 is the prescribed dose for t -th PTV, EUD_r^0 is the maximum uniform dose at r -th OAR; n_r and n_t express the importance of the prescriptions for the corresponding structure; ϕ represents the set of parameters involved in the F definition, i.e., ϕ is an instance of parameters n_t, n_r, a_t, a_r and EUD_r^0 with $t \in T, r \in R$.

As is, the objective function $F(x, \phi)$ only controls underdosage inside PTVs. However, overdosage control in those volumes is also important. For this purpose it is common to define complementary structures, introduced as ‘‘virtual PTVs’’ in Wu *et al.* (2002), that are treated as OARs. In this work, for each PTV we have defined a virtual PTV. To lighten optimization costs, we have interrelated the respective parameters, namely, for each PTV t , the corresponding virtual PTV r has $EUD_r^0 = EUD_t^0 + 1$, $a_r = -a_t$ and $n_r = n_t$.

In Choi and Deasy (2002), the convexity of function (1) was studied. It was shown that for a range of parameters, function (1) is a convex function of fluence maps (x). In Romeijn *et al.* (2004), the convexity of (2) was analysed and the conclusion was that it is convex for the range of values of the parameters a_s considered in practice. This means that IMRT RT plans can be efficiently obtained by gradient methods, as suggested in Wu *et al.* (2002).

RT planning based on maximization of function $F(x, \phi)$ goes as follows. At the start, the values of $n_t, n_r, a_t, a_r, t \in T, r \in R$, are selected according to suggestions from the literature. Next, in successive planning cycles, the parameters a_p and $EUD_p^0, p \in P$, are manually tuned by the RT planner. This is a tiresome and time-consuming process, requiring high-expertise planners. In this work, we propose to select the gEUD parameters automatically by multi-objective optimization. In the next section, we present this idea in more detail.

2.2. Multi-Objective Optimization for gEUD-Based IMRT Planning

Automated parameter tuning can result in several, even many, RT plans generated in one planning session. We present a multi-objective optimization model which serves that purpose. We also outline a novel approach to solve this model by a hybrid optimization scheme, coupling exact and heuristic (evolutionary) optimization methods.

2.2.1. The Multi-Objective Optimization Model

Here we present a multi-objective optimization model capable of handling physical constraints imposed on RT plans. Such constraints are individually selected in the optimization process based on the individual patient's anatomy.

For structure s we consider the following dose distribution statistical measures:

$$\begin{aligned}
 D_s^{\min}(x) &= \min_{j \in M_s} d_j(x), \\
 \overline{D_s(x)} &= \frac{1}{|M_s|} \sum_{j \in M_s} d_j(x), \\
 D_s^{\max}(x) &= \max_{j \in M_s} d_j(x).
 \end{aligned} \tag{3}$$

With these measures we define constraints in structures. Four constraints are defined for each PTV:

$$\begin{aligned}
 D_t^{\min}(x) &\geq LB_t, \\
 \overline{D_t(x)} &\geq \overline{LB_t}, \\
 \overline{D_t(x)} &\leq \overline{UB_t}, \\
 D_t^{\max}(x) &\leq UB_t,
 \end{aligned} \tag{4}$$

where, for given t , LB_t and UB_t are the lower and upper bound for the dose in any voxel of the structure, $\overline{LB_t}$, $\overline{UB_t}$ are lower and upper bound for the average dose in the structure.

Doses in parallel² OARs are constrained by upper bounds on the average dose in the structure:

$$\overline{D_r(x)} \leq \overline{UB_r}, \tag{5}$$

whereas doses in serial³ OARs are constrained by upper bounds on the maximal dose in individual voxels of the structure:

$$D_r^{\max}(x) \leq UB_r. \tag{6}$$

Constraint violations in structures, PTVs and OARs, are captured by the following constraint violation functions:

²An organ is called *parallel* if its functionality is preserved despite partial radiation damage, e.g. the salivary gland.

³An organ is called *serial* if it loses its functionality completely if any part of it is damaged, e.g. the spinal cord.

$$\begin{aligned}
 C_s^{\min}(x) &= \begin{cases} LB_s - D_s^{\min}(x), & \text{if } LB_s \text{ is defined in } s \text{ and } LB_s > D_s^{\min}(x), \\ 0, & \text{otherwise;} \end{cases} \\
 \overline{C_s^{\min}}(x) &= \begin{cases} \overline{LB_s} - \overline{D_s(x)}, & \text{if } \overline{LB_s} \text{ is defined in } s \text{ and } \overline{LB_s} > \overline{D_s(x)}, \\ 0, & \text{otherwise;} \end{cases} \\
 \overline{C_s^{\max}}(x) &= \begin{cases} \overline{D_s(x)} - \overline{UB_s}, & \text{if } \overline{UB_s} \text{ is defined in } s \text{ and } \overline{UB_s} < \overline{D_s(x)}, \\ 0, & \text{otherwise;} \end{cases} \\
 C_s^{\max}(x) &= \begin{cases} D_s^{\max}(x) - UB_s, & \text{if } UB_s \text{ is defined in } s \text{ and } UB_s < D_s^{\max}(x), \\ 0, & \text{otherwise.} \end{cases}
 \end{aligned} \tag{7}$$

Since obviously constraint violations should be minimized, we formulate the following objective functions. The first function, denoted f_0 , aggregates constraint violations over all structures:

$$f_0(x) = \sum_{s \in S} C_s^{\min}(x) + \overline{C_s^{\min}}(x) + \overline{C_s^{\max}}(x) + C_s^{\max}(x). \tag{8}$$

The priorities to protect the distinguished subsets of OARs, indexed by p , $p \in P \subseteq R$, are represented in the model by separate objective functions being the averages of the deposited doses, $\overline{D_p(x)}$:

$$f_p(x) = \overline{D_p(x)}, \quad p \in P, \tag{9}$$

or

$$f_p(x) = D_p^{\max}(x), \quad p \in P. \tag{10}$$

Clearly, low values of individual functions $f_0, f_1, \dots, f_{|P|}$ signal acceptable plans.

The multi-objective optimization problem consists of $|P| + 1$ objective functions $f_0, \dots, f_{|P|}$, all of them to be minimized:

$$\min_X (f_0, f_1, \dots, f_{|P|}), \tag{11}$$

where X is the set of technically feasible fluencies.

This model applies for any method of fluence map optimization, capable to handle functions (8) and (9).

The classical approach to use model (11) for RT planning, with no reference to the gEUD, would be to produce a set of fluence maps x for which $f_0(x), \dots, f_{|P|}(x)$ are mutually nondominated or, with the use of exact optimization methods, even nondominated in X . That can be achieved with a form of scalarization of the multi-objective problem (11) (Ehrgott *et al.*, 2010; Kaliszewski *et al.*, 2016).

2.2.2. Combining the Multi-Objective Optimization Model and the gEUD-Based Optimization

We propose the following iterative scheme for automated tuning of gEUD parameters, i.e. elements of ϕ , $\phi \in \Phi$, where Φ is the set of admissible parameters ϕ . In each iteration:

1. Each collection of gEUD parameters ϕ is evaluated with respect to values of $|P| + 1$ criteria $f_0(x^*(\phi)), \dots, f_{|P|}(x^*(\phi))$, where $x^*(\phi) = \operatorname{argmax}_X F(x, \phi)$ is derived by an exact optimizer.
2. A multi-objective evolutionary optimization algorithm searches set Φ for collections of parameters ϕ , $\phi \in \Phi$, such that tuples $f_0(x^*(\phi)), \dots, f_{|P|}(x^*(\phi))$ are mutually nondominated and they dominate analogous tuples derived in the previous iteration.

The process stops when the differences between tuples in successive iterations become insignificant. The final collection of tuples $(f_0(x^*(\phi)), \dots, f_{|P|}(x^*(\phi)))$ represents an approximation of the Pareto efficient set of RT plans $x^*(\phi)$ (the set of all nondominated RT plans at the final iteration).

A distinctive feature of this scheme is that it hybridizes exact and heuristic optimization methods (Jourdan *et al.*, 2009).

3. An IMRT Planning System with Automated gEUD Parameter Tuning: PersEUD

We have developed and implemented a system, named PersEUD, to deliver IMRT RT plans based on the gEUD, automated parameter ϕ tuning, and the optimization hybrid scheme described in Section 2.2.2. It is composed of four modules:

1. **EUDGD**, to deliver the optimal solutions $x^*(\phi)$ maximizing function (2) for a given collection of parameters ϕ ,
2. **EvalMO**, to compute values of functions $f_0, f_1, \dots, f_{|P|}$ for $x^*(\phi)$,
3. **EvolTuning**, a multi-objective evolutionary algorithm, to provide collections of parameters ϕ , nondominated in terms of functions $f_0, f_1, \dots, f_{|P|}$,
4. **Decision Tool**, to control the number of collections of parameters ϕ presented to RT planners.

The flow diagram of PersEUD is shown in Fig. 1. One run of the EvolTuning module produces K collections of parameters ϕ , mutually nondominated in terms of functions $f_0, f_1, \dots, f_{|P|}$. For each ϕ^k , $k \in K$, the module EUDGD computes $x^*(\phi^k)$ that maximizes function (2) and next, the module EvalMO computes values of functions $f_0, f_1, \dots, f_{|P|}$. These values are used by the EvolTuning module to identify mutually nondominated parameters ϕ directing in turn the search for new K collections of ϕ .

It is worth mentioning the versatility of PersEUD, as it can easily accommodate alternative criteria, because criteria $f_0, f_1, \dots, f_{|P|}$ are only involved in EvalMO module. The next four subsections provide more detailed descriptions of the modules of PersEUD.

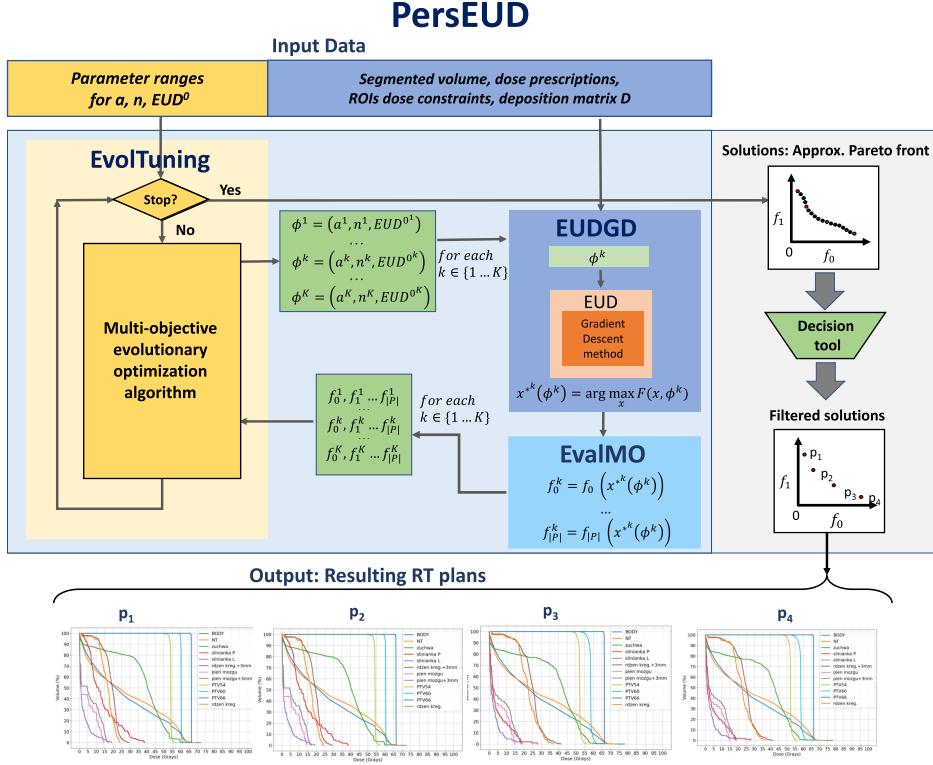


Fig. 1. PersEUD flow diagram.

3.1. The EUDGD Module

To maximize function (2), the Gradient Descent method has been custom-coded and the resulting algorithm (GD) is the backbone of the EUDGD module.

The EUDGD module receives as inputs for each patient case: the patient data, the deposition matrix, the beamlet geometry, the ROI segmentation, parameters ϕ and the number of steps. It returns the optimal fluence $x^*(\phi)$. To generate feasible fluence maps, the EUDGD module needs additional configuration parameters. In the current software release, these parameters are set to default values, but they can be customized. The default GD parameter values reported in this work are: 20000 descent steps, $2e^{-7}$ step size, and 0.3 as the maximum beamlet intensity limit. Lastly, a $3 \times$ smoothing kernel has been applied to all beamlets after every descent step to facilitate the delivery of fluencies defined by RT plans by the radiation equipment.

The number of descent steps required by EUDGD to provide acceptable fluence maps and the size of the auxiliary data structures result in high computational and memory demands. To maximize function (2) in reasonable times, we have applied High Performance Computing software and hardware techniques, as described in Moreno *et al.* (2022).

3.2. The EvalMO Module

For fluence maps $x^*(\phi^k)$ delivered by the EUDGD module, the EvalMO module computes the values of functions $f_0(x^*(\phi^k)), \dots, f_{|P|}(x^*(\phi^k))$ and sends them to EvolTuning module.

3.3. The EvolTuning Module

The EvolTuning consists of a multi-objective evolutionary optimization algorithm that derives parameters ϕ . It is initialized with collections of $n_s, a_s, EUD_s^0, s \in T \cup R$. In order to generate clinically acceptable plans and reduce the search space of the evolutionary process, feasibility ranges for each kind of parameter are to be set accordingly.

For the current implementation, the MOEA/D algorithm has been selected and seeded with parameter values recommended in Zhang and Li (2007), with the exception of the number of generations (epochs) that is set to 50 and the cardinality of populations set to 150.

EvolTuning uses the evaluations received from EvalMO to generate new collections of (potentially improved) parameters ϕ . The process stops when the limit of generations is reached. As a result, 150 mutually nondominated collections ϕ of parameters of function (2), each collection yielding radiation plan in the form of fluence map $x^*(x)$, 150 plans in all, constitute an approximation of the set of Pareto efficient collections ϕ (in terms of functions $f_0, \dots, f_{|P|}$).

3.4. Decision Tool

To avoid overwhelming the RT planner with so many plans, the module Decision Tool reduces the set of mutually nondominated collections ϕ . It performs this task as follows:

First, plans that minimize objective functions $f_0, f_1, \dots, f_{|P|}$ separately are selected, and they define the extreme points of the respective Pareto front. Next, a limited number of plans, fairly distributed between those extreme points, are added to form the reduced approximation of the Pareto front. Observe that even if f_0 for some plan takes a positive but small value (meaning that there are some constraint violations), it can be still a clinically viable option as long as it offers plausible values of $f_1, \dots, f_{|P|}$ at the price of a small constraint violation.

4. Results

As the proof of concept of our proposed methodology, we have prepared three RT plans for real patient data specified in Section 4.2.

The EUDGD module uses a radiation dose deposition model, yielding the deposition matrix D , developed by researchers at the Faculty of Electrical Engineering of the Warsaw University of Technology (Starzyński *et al.*, 2015). The final plan evaluations have been carried out using the commercially available treatment planning system Eclipse (v 15.6,

Varian Medical Company). By this, the dose deposition model, as well as the proposed methodology, were tested in a clinical regime.

4.1. Implementation and Experimentation Platforms

The gEUD-based RT planning system, described in the previous section, consists of multiple modules connected by a messaging system based on the ZeroMQ library (ZeroMQ, 2023).

The EvolTuning module is implemented in Java and it takes advantage of the well-known jMetal framework (Durillo and Nebro, 2011). This framework provides well-crafted implementations of multiple algorithms, allowing the user to select the best one for a given problem. Although in this experimentation we have used the MOEA/D algorithm, our modular system allows us to swap the evolutionary algorithm with minimal changes.

For the EUDGD and EvalMO modules, we have developed two versions: a CUDA C (NVIDIA, 2023) version for NVIDIA-based GPUs, and an OpenMP C (OpenMP, 2023) version for multicore CPUs. These two versions can be used interchangeably or even in parallel to fully exploit all the available resources of the given platform.

Finally, the Decision Tool is a Python-based script that is called after the EvolTuning process terminates. All the figures shown in this section are generated using Python tools implemented in the Decision Tool module.

Our experiments have been run on the High-Performance Cluster of the SAL (Supercomputación–Algoritmos) research group, located at the University of Almeria. Two kinds of computing nodes have been used. The first node contains two AMD EPYC 7302 (32 CPU cores), 512 GB of DDR4 RAM, and two NVIDIA Tesla V100 (32 GB). The second node contains two Intel Xeon E5-2620v3 (12 CPU cores), 64 GB of DDR3 RAM, and two NVIDIA Kepler K80 (12 GB).

4.2. Patient Data and Plan Evaluation Metrics

For the experimentation, we have used a Head and Neck IMRT real patient case treated with nine radiation beams. In this case, three PTVs with different prescribing radiation dose deposition levels 66 Gy, 60 Gy and 54 Gy (Gray, symbol Gy, is a unit of radiation dose) are defined. The most important is the PTV with the highest prescribed dose (PTV66) because the highest concentration of tumour cells is there. The other two PTVs are treated prophylactically, so the dose distribution homogeneity in them is less critical. In addition, five OARs are singled out: the spinal cord +3mm, the brainstem +3mm, the left salivary gland, the right salivary gland, and the mandible. Also, the normal tissue is defined as a region of the patient body outside all OARs and all PTVs. The dose deposition model contains 30265 beamlets interacting with 94647 voxels representing the irradiated part of the patient body.

Physical restrictions imposed on RT plans by oncology physicians are converted to the bounds defined in the model presented in Subsection 2.2.1. Table 2 presents the respective bound values for organs considered in the experiments. Those values are repeated also in Table 4 and Table 6. The bound values for PTVs are set by the following rules:

Table 2

Plan 1. Dose bounds, actual doses, and Dx% metrics. Values that exceed their constraint are marked in bold.

Region of Interest	Dose bounds				Dose Act.	Dx%	
	LB	\overline{LB}	\overline{UB}	UB		Bound	Act.
	Gy						
Normal tissue	–	–	–	74.25	71.52	–	–
Mandible	–	–	–	70.00	68.28	–	–
Salivary gland R.	–	–	26.00	–	14.26	–	–
Salivary gland L.	–	–	26.00	–	12.56	–	–
Spinal cord +3mm	–	–	–	50.00	41.43	–	–
Brainstem +3mm	–	–	–	60.00	28.67	–	–
PTV54	48.60	52.92	55.08	59.40	53.47	–	–
D98% for PTV54	–	–	–	–	–	51.30	48.00
D2% for PTV54	–	–	–	–	–	57.78	56.22
PTV60	54.00	58.80	61.20	66.00	60.04	–	–
D98% for PTV60	–	–	–	–	–	57.00	52.06
D2% for PTV60	–	–	–	–	–	64.20	64.70
PTV66	59.40	64.67	67.32	72.60	66.00	–	–
D98% for PTV66	–	–	–	–	–	62.70	63.74
D2% for PTV66	–	–	–	–	–	70.62	67.32

$$LB_t = 0.90 \times \text{prescribed dose for } PTV_t,$$

$$\overline{LB}_t = 0.98 \times \text{prescribed dose for } PTV_t,$$

$$\overline{UB}_t = 1.02 \times \text{prescribed dose for } PTV_t,$$

$$UB_t = 1.10 \times \text{prescribed dose for } PTV_t, t \in T.$$

Explicit constraints on DVH shapes are not a part of model (11) as they are implicitly controlled by gEUD. On the other hand, the DVH shape for PTVs ideally should look like the sequence of three-line segments: horizontal (at 100% level), vertical (at the prescribed deposited dose value), and horizontal again (at 0% level) (cf. Fig. 2, Fig. 4, Fig. 6). DVHs are primary tool to verify ex-post the homogeneity of dose distributions in PTVs. As such, they are always visually inspected by RT planners and oncology physicians for a holistic evaluation. A proxy measure of PTV homogeneity is the dose-volume metric Dx%. For a given PTV, Dx% is the minimal dose deposited in x% of the most irradiated PTV voxels.

In our case, metrics D98% and D2% are used, namely plans are to satisfy

$$D98\% \geq 95\% \times \text{prescribed dose for } PTV_t, \text{ and}$$

$D2\% \leq 107\% \times \text{prescribed dose for } PTV_t$ (i.e. in the latter case: the minimal dose deposited in 2% of the most irradiated PTV voxels should be less or equal than $107\% \times EUD_t^0$).

By this, we have the following levels of Dx%

for PTV66 Gy: $D98\% \geq 62.70$ Gy, $D2\% \leq 70.62$ Gy,

for PTV60 Gy: $D98\% \geq 57.00$ Gy, $D2\% \leq 64.20$ Gy,

for PTV54 Gy: $D98\% \geq 51.30$ Gy, $D2\% \leq 57.78$ Gy.

D2% and D98% represent two points on the respective DVHs, thus they give a rough characterization of DVHs.

4.3. Numerical Experiments

The planning system PersEUD is run three times, each time with a different aim. The aim of the first run is to derive a number (150 by default) of mutually nondominated plans that compromise constraint violations versus the sum of the average of radiation doses deposited in the left and the right salivary glands (a bi-objective optimization problem). In the second run, the aim is to derive plans that compromise constraint violations versus doses deposited in the spinal cord + 3mm (a bi-objective optimization problem), whereas, in the third run, the aim is to derive plans that compromise constraint violations versus doses deposited in the left and the right salivary gland versus doses deposited in the spinal cord + 3mm (a tri-criteria optimization problem).

The derived plans are evaluated by medical physicist experts (below: the Experts) who, in their daily work, make treatment plans in a commercial RT planning system in a series of try-and-correct interactions. Once they are satisfied with the result, they submit the plan to the oncology physician responsible for the case for approval. If the plan is disapproved, the process is repeated. In complex cases, two or three cycles may be needed to secure the final physician approval.

In the whole process, they deal with one RT plan at a time. The three runs of PersEUD produce 450 plans, fully automatic. For proof-of-concept analyses, the Experts selected one plan from those produced in each PersEUD run and pre-selected by the Decision Tool module.

4.4. *PersEUD Run 1: Compromising Constraint Violations Versus the Sum of the Average of Radiation Doses Deposited in Both Salivary Glands*

To fulfill the aim of the first PersEUD run, the second objective f_1 is the sum of the average radiation doses deposited in both glands, $|P| = 1$, and f_0 is defined by formula (8).

From the resulting RT plans, Experts select one that does not violate any constraint ($f_0(x) = 0$). This plan is denoted as Plan 1 and it is presented in Table 2. Table 3 presents the parameters of gEUD measures and of the function $F(x, \phi)$ for this plan, Fig. 2 presents its dose-volume histograms, and Fig. 3, the dose distributions for this plan on two exemplary cross sections of the patient irradiated part (head and neck).

As seen in Table 2, this plan provides adequate dose distributions in PTVs, and, as intended, preserves both salivary glands, and it does this surprisingly well. The average doses are 14.26 Gy for the right and 12.56 Gy for the left salivary gland. These values are well below the imposed bound of 26 Gy. All other OARs with no protection priority are also below their acceptable maximal doses. The most notable is the maximum dose in the spinal cord + 3mm, equal to 41.43 Gy, while the bound is 50 Gy, and the maximum dose in the brainstem + 3mm, equal to 28.67 Gy, while the bound is 60.00 Gy.

The actual average dose deposited in PTV66 (the most important PTV) is 66 Gy, as prescribed, and $D98\% \geq 62.70$ Gy and $D2\% \leq 70.62$ Gy.

The actual average dose deposited in PTV60 is 60.04 Gy and that is almost the perfect match as this value is within $\pm 2\%$ range from the target value. However, $D98\%$ is significantly unmet (52.06 Gy vs 57.00 Gy expected) due to the protection of the salivary gland. $D2\%$ is also exceeded but by a clinically insignificant value.

Table 3

Plan 1. Parameters of gEUD (a_s) and of $F(x, \phi)$ (n_s). Black: Parameters ϕ suggested in the literature. Blue: Parameter search ranges. Green: Optimal parameters derived by the EvoITuning module.

Region of Int.	EUD_S^0	a_s	n_s
Normal tissue	74.25	40.00	5.00
Mandible	70.00	10.00	5.00
Salivary gland R.	[0.5, 26] : 4.37	[1, 100] : 1.01	[1, 100] : 100.00
Salivary gland L.	[0.5, 26] : 3.93	[1, 100] : 1.19	[1, 100] : 18.04
Spinal cord +3mm	50.00	10.00	5.00
Brainstem +3mm	60.00	10.00	5.00
PTV54	54.00	[-100, -1] : -96.04	[1, 100] : 34.41
PTV60	60.00	[-100, -1] : -62.81	[1, 100] : 66.77
PTV66	66.00	[-100, -1] : -90.32	[1, 100] : 92.63




Fig. 2. Plan 1. Dose-volume histograms.

The actual average dose deposited in PTV54 is 53.47 Gy and that is also almost the perfect match as this value is within $\pm 2\%$ range from the target value. Moreover, $D2\% \leq 57.78$ Gy, but $D98\% \not\geq 51.30$.

As mentioned earlier, in regions treated prophylactically (in the considered case PTV 60 and PTV54) slight violations of $D98\%$ and $D2\%$ related constraints are acceptable.

4.5. PersEUD Run 2: Compromising Constraint Violations Versus Doses Deposited in the Spinal Cord +3mm

To fulfill the aim of the second PersEUD run, the second objective f_1 is the average radiation dose deposited in the spinal cord +3mm, $|P| = 1$, and f_0 is defined by formula (8).




Fig. 3. Plan 1. Dose distributions for $Z = 75$ (left) and $Z = 91$ (right). Brown: Normal tissue. Pink: Mandible. Green: Salivary gland R. Blue: Salivary gland L. Cyan: Spinal cord +3mm. Red: PTV54. Orange: PTV60.

Table 4

Plan 2. Dose bounds, actual doses, and Dx% metrics. Values that exceed their constraint are marked in bold.

Region of Interest	Dose bounds				Dose Act.	Dx%	
	LB	\overline{LB}	\overline{UB}	UB		Bound	Act.
	Gy						
Normal tissue	–	–	–	74.25	67.30	–	–
Mandible	–	–	–	70.00	71.13	–	–
Salivary gland R.	–	–	26.00	–	22.08	–	–
Salivary gland L.	–	–	26.00	–	22.19	–	–
Spinal cord +3mm	–	–	–	50.00	11.02	–	–
Brainstem +3mm	–	–	–	60.00	9.71	–	–
PTV54	48.60	52.92	55.08	59.40	54.30	–	–
D98% for PTV54	–	–	–	–	–	51.30	48.03
D2% for PTV54	–	–	–	–	–	57.78	58.81
PTV60	54.00	58.80	61.20	66.00	59.68	–	–
D98% for PTV60	–	–	–	–	–	57.00	58.81
D2% for PTV60	–	–	–	–	–	64.20	63.80
PTV66	59.40	64.67	67.32	72.60	66.00	–	–
D98% for PTV66	–	–	–	–	–	62.70	60.87
D2% for PTV66	–	–	–	–	–	70.62	71.66

From the resulting RT plans, the Experts select one that does not violate any constraint ($f_0(x) = 0$). This plan is denoted as Plan 2 and is presented in Table 4. Table 5 presents parameters of gEUD and of function $F(x, \phi)$ for this plan. Fig. 4 presents its dose-volume histograms, and Fig. 5, the dose distributions for this plan on two exemplary cross sections of the patient irradiated part (head and neck).

In Plan 2, the maximum dose deposited in the spinal cord +3mm is 11.02 Gy, noticeably lower than in the Plan 1. Although protection of the salivary glands in this plan is not a priority, the average dose is 22.19 Gy and 22.08 Gy in the left and right salivary gland,

Table 5

Plan 2. Parameters of gEUD (a_s) and of $F(x, \phi)$ (n_s). Black: Parameters ϕ suggested in the literature. Blue: Parameter search ranges. Green: Optimal parameters derived by the EvolTuning module.

Region of Int.	$EU D_s^0$	a_s	n_s
Normal tissue	74.25	40.00	5.00
Mandible	70.00	10.00	5.00
Salivary gland R.	26.00	1.00	5.00
Salivary gland L.	26.00	1.00	5.00
Spinal cord +3mm	[0.5, 50] : 0.50	[1, 50] : 49.30	[1, 100] : 10.25
Brainstem +3mm	60.00	10.00	5.00
PTV54	54.00	[-100, -1] : -67.50	[1, 100] : 17.76
PTV60	60.00	[-100, -1] : -49.83	[1, 100] : 25.35
PTV66	66.00	[-100, -1] : -7.02	[1, 100] : 14.07




Fig. 4. Plan 2. Dose-volume histograms.

respectively, which is far below the dose tolerance for these organs. This is at the price of a slight violation of the maximal dose allowed in the mandible (70 Gy) where the actual dose deposited is 71.13 Gy.

The actual average dose deposited in PTV66 is 66 Gy, as prescribed. However, $D98\% \not\geq 62.70$ Gy (violation by 1.83 Gy), and $D2\% \not\leq 70.62$ Gy (violation by 1.04 Gy).

The actual average dose deposited in PTV60 is 59.68 Gy and that is almost the perfect match as this value is within $\pm 2\%$ range from the target value. Moreover, $D98\% \geq 57.00$ Gy and $D2\% \leq 64.20$ Gy.

The actual average dose deposited in PTV54 is 54.30 Gy and that is also almost the perfect match as this value is within $\pm 2\%$ range from the target value. However, $D98\% \not\geq 51.30$ Gy, but $D2\% \leq 57.78$ Gy.




Fig. 5. Plan 2. Dose distributions for $Z = 75$ (left) and $Z = 91$ (right). Brown: Normal tissue. Pink: Mandible. Green: Salivary gland R. Blue: Salivary gland L. Cyan: Spinal cord +3mm. Red: PTV54. Orange: PTV60.

As mentioned earlier, in regions treated prophylactically (i.e. PTV 60 and PTV54) slight violations of D98% and D2% related constraints are acceptable.

4.6. *PersEUD Run 3: Compromising Constraint Violations Versus Doses Deposited in the Left and the Right Salivary Gland Versus Doses Deposited in the Spinal Cord +3mm*

To fulfill the aim of the third PersEUD run, the second objective function f_1 is the average of the radiation doses deposited in the left and the right salivary glands, the third objective function is the maximal radiation dose deposited in the spinal cord +3mm, $|P| = 2$, and f_0 is defined by formula (8).

From the resulting RT plans, the Experts select one that does not violate any constraint ($f(x) = 0$). This plan is denoted as Plan 3 and is presented in Table 6. Table 7 presents parameters of gEUD and of function $F(x, \phi)$ for this plan, Fig. 6 presents its dose-volume histograms, and Fig. 7, the dose distributions for this plan on two exemplary cross sections of the 3D model of the patient irradiated part (head and neck).

In Plan 3, the average dose deposited in the right and the left salivary gland is 13.94 Gy and 13.14 Gy, respectively. This is much lower than the maximal admissible value of 26 Gy, while the maximal dose deposited in the spinal cord +3mm is 24.06 Gy (the maximal acceptable value is 50 Gy). Although Plan 1 and Plan 2 offer slightly lower doses in respective priority protected OARs (the salivary glands in the first case, and the spinal cord +3mm in the second case), Plan 3 well balances the doses delivered to both salivary glands and to the spinal cord +3mm.

The actual average dose deposited in PTV66 is 66 Gy, as prescribed, and $D98\% \geq 62.70$ Gy, and $D2\% \leq 70.62$ Gy. The actual average dose deposited in PTV60 is 60.28 Gy, within $\pm 2\%$ range from the target value. Moreover, $D98\% \geq 57.00$ Gy and $D2\% \leq 64.20$ Gy. The actual average dose deposited in PTV54 is 55.34 Gy and that is also almost

Table 6

Plan 3. Dose bounds, actual doses, and Dx% metrics. Values that exceed their constraint are marked in bold.

Region of Interest	Dose bounds				Dose Act.	Dx%	
	LB	\overline{LB}	\overline{UB}	UB		Bound	Act.
	Gy						
Normal tissue	–	–	–	74.25	73.14	–	–
Mandible	–	–	–	70.00	69.83	–	–
Salivary gland R.	–	–	26.00	–	13.94	–	–
Salivary gland L.	–	–	26.00	–	13.14	–	–
Spinal cord +3mm	–	–	–	50.00	24.06	–	–
Brainstem +3mm	–	–	–	60.00	21.78	–	–
PTV54	48.60	52.92	55.08	59.40	55.34	–	–
D98% for PTV54	–	–	–	–	–	51.30	49.64
D2% for PTV54	–	–	–	–	–	57.78	59.13
PTV60	54.00	58.80	61.20	66.00	60.27	–	–
D98% for PTV60	–	–	–	–	–	57.00	52.30
D2% for PTV60	–	–	–	–	–	64.20	65.49
PTV66	59.40	64.67	67.32	72.60	66.00	–	–
D98% for PTV66	–	–	–	–	–	62.70	63.15
D2% for PTV66	–	–	–	–	–	70.62	67.63

Table 7

Plan 3. Parameters of gEUD (a_s) and of $F(x, \phi)$ (n_s). Black: Parameters ϕ suggested in the literature. Blue: Parameter search ranges. Green: Optimal parameters derived by the EvolTuning module.

Region of Int.	EUD_s^0	a_s	n_s
Normal tissue	74.25	40.00	5.00
Mandible	70.00	10.00	5.00
Salivary gland R.	[0.5, 26] : 4.76	[1, 100] : 1.01	[1, 100] : 18.25
Salivary gland L.	[0.5, 26] : 3.79	[1, 100] : 1.31	[1, 100] : 11.17
Spinal cord +3mm	[0.5, 50] : 1.80	[1, 50] : 1.33	[1, 100] : 12.79
Brainstem +3mm	60.00	10.00	5.00
PTV54	54.00	[−100, −1] : −65.55	[1, 100] : 54.11
PTV60	60.00	[−100, −1] : −87.12	[1, 100] : 57.66
PTV66	66.00	[−100, −1] : −33.27	[1, 100] : 18.62

the perfect match as this value is within $\pm 2.5\%$ range from the target value. However, $D98\% \not\geq 51.30$ Gy, but $D2\% \leq 57.78$ Gy.

As mentioned earlier, in regions treated prophylactically (i.e. PTV60 and PTV54) slight violations of $D98\%$ and $D2\%$ related constraints are acceptable.

4.7. Analysis of PersEUD Plans in a Clinical Setting

Table 8 provides a comparison of the three plans. At a glance, the Experts disapproved Plan 2 due to its underperformance in PTV66, the most important PTV in the considered case. The explanation for the disapproval is that such an underperformance cannot be outweighed by the excellent performance in the spinal cord +3mm and the brainstem




Fig. 6. Plan 3. Dose-volume histograms.




Fig. 7. Plan 3. Dose distributions for $Z = 75$ (left) and $Z = 91$ (right). Brown: Normal tissue. Pink: Mandible. Green: Salivary gland R. Blue: Salivary gland L. Cyan: Spinal cord +3mm. Red: PTV54. Orange: PTV60.

+3mm. From the remaining two, the Experts prefer Plan 3 because it offers low radiation doses deposited in both salivary glands, the spinal cord +3mm, and brainstem +3mm.

To verify the clinical viability of Plan 3 (below: P3), we compare it with two plans prepared by the Expert in Eclipse™ Treatment Planning system (below: Eclipse) from Varian Medical Systems, namely:

R1 *Starting plan*: an RT plan generated automatically by Eclipse at the start of the planning process.

Table 8
A comparison table for the three plans. Values that exceed their corresponding constraint are marked in bold.

Region of Interest	Bound type	Bound	Plan 1	Plan 2	Plan 3
		Gy			
Normal tissue	Average	74.25	71.52	67.30	73.14
Mandible	Maximum	70.00	68.28	71.13	69.83
Salivary gland R.	Average	26.00	14.26	22.08	13.94
Salivary gland L.	Average	26.00	12.56	22.19	13.14
Spinal cord +3mm	Maximum	50.00	41.43	11.02	24.06
Brainstem +3mm	Maximum	60.00	28.67	9.71	21.79
PTV54	Average	54.00	53.47	54.30	55.34
	D98%	51.30	48.00	48.03	49.64
	D2%	57.78	56.22	58.81	59.13
PTV60	Average	60.00	60.04	59.68	60.27
	D98%	57.00	52.06	58.81	52.30
	D2%	64.20	64.76	63.80	65.49
PTV66	Average	66.00	66.00	66.00	66.00
	D98%	62.70	63.74	60.87	63.15
	D2%	70.62	67.32	71.66	67.63




Fig. 8. DVHs of Plan 3 (solid) compared with R1, the starting plan (dashed).

R2 *Blind-expert plan*: an RT plan prepared by the Expert without seeing P3.

All plans have to satisfy the same constraints (specified in Table 6). The comparisons are by DVH and, for illustration, by one representative cross section of the patient model.

The first comparison (Fig. 8) is between P3 and R1. As displayed, P3 reduces the doses delivered to all OARs, with comparable or better dose distributions in PTVs.

When comparing P3 with R2 (Fig. 9), R2 provides a better dose distribution in PTVs,




Fig. 9. DVHs of plan 3 (solid) compared with the plan R2, the blind-expert plan (dashed).




Fig. 10. DVHs of the guided plan R3 (solid) compared with R2 (dashed).

however at the expense of higher dose deposited in OARs, especially in both salivary glands and the spinal cord +3mm. Notably, in P3 the doses deposited in the salivary glands are 20% lower than in R2. And in P3 the dose deposited in PTV66 is within the specified bounds.

Next, in Fig. 11, we compare with *Guided-expert plan* (R3), a plan prepared by the Expert using P3 as the Eclipse starting plan, with R2 (Fig. 10). R3 provides a similar to

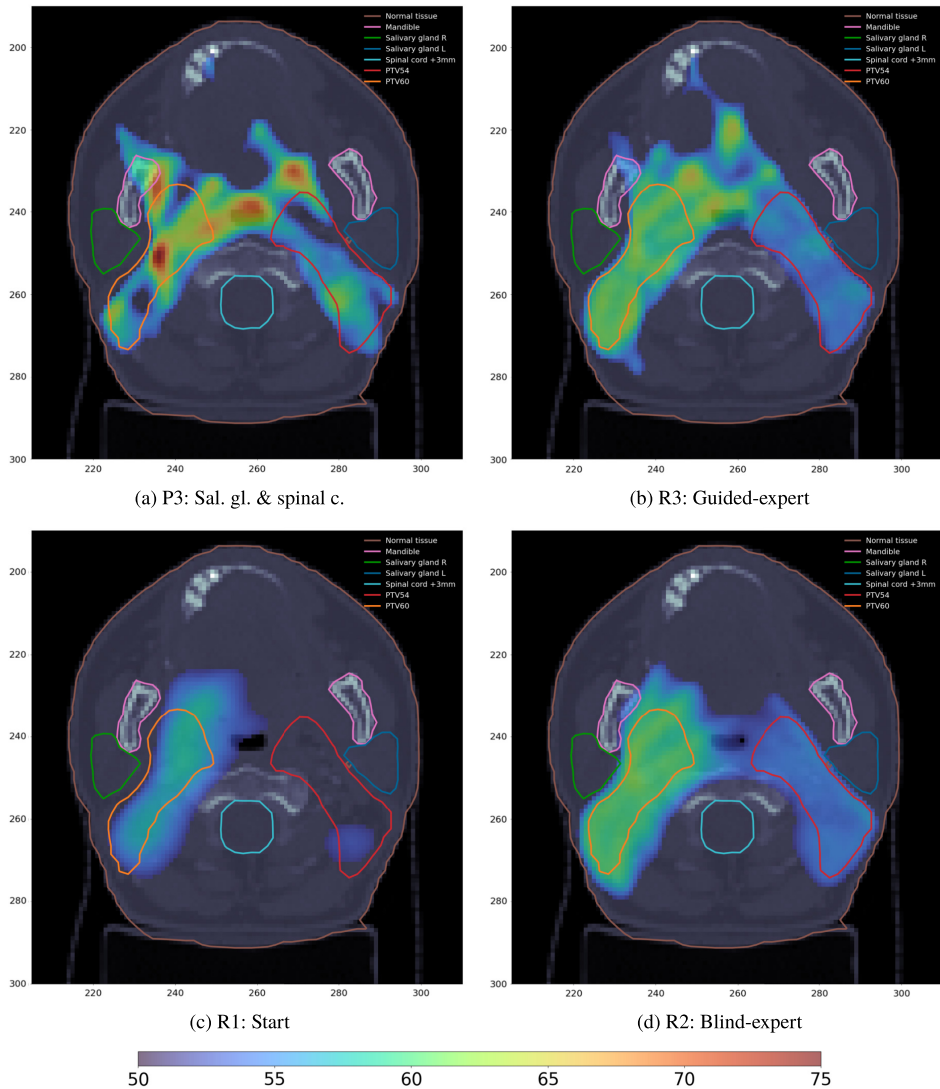


Fig. 11. Dose deposition of the four compared plans for $Z = 83$. Only doses higher than 51.3 Gy are represented.

R2 dose distribution in PTVs and a significant reduction in the doses deposited in OARs. This shows that plans produced by the PersEUD system can be used to enhance RT plans produced by commercial systems currently in use.

As preliminary conclusions from our work, it can be assumed that using plans from PersEUD as starting points in RT planning systems routinely used in an oncology hospital should improve RT plan quality and reduce RT plan preparation times. This claim has to be validated in more clinical cases.

5. Conclusions

The gEUD-based RT planning optimization has the potential to deliver better RT plans than planning routines based on physical criteria. This is because it directly refers to the biological phenomena in the irradiated cells, whereas the physical criteria measure radiation effects indirectly. Despite its potential, the wide usage of biological optimization techniques in current clinical practices is hampered by the problematic translation between input parameters and resulting dose distributions (RT plans), requiring time-consuming manual model tuning. To help with this situation, we propose a system for automated parameter tuning by optimization. To our best knowledge, it is the first system of this sort.

The distinctive feature of the proposed system is that it consists of two optimization levels. Because of the number of unknown parameters and the complex nature of the gEUD-based function (2), this function cannot be optimized directly. Instead, the bi-level optimization scheme has been adopted. On the upper level, the space of the parameters is searched for promising collections of parameters, and the search is guided by values of objective functions of the multi-objective optimization model proposed. Any reasonable heuristic can perform this search, and we have employed evolutionary multi-objective optimization to this aim. On the lower level, the gEUD-based function (function (2)) optimization is performed by a custom-implemented exact optimization method.

Starting with a solution provided by PersEUD, rather than starting from scratch, can significantly reduce the time it takes an expert to develop a radiotherapy treatment plan. This improvement underlines the efficiency and effectiveness of the planning process by providing a solid starting point. More importantly, this approach can be used independently of the commercial treatment planning system used by the expert to refine the final solution. However, to comprehensively validate its efficacy, further testing with a broader spectrum of clinical cases is needed.

The bi-level optimization scheme we apply here to RT planning optimization is clearly also applicable to any optimization problem in which the search is needed not only for optimal values of model variables but also for favourable (at best: optimal) model parameters.

Funding

This work has been supported by Grant PID2021-123278OB-I00 funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe”; and by projects PDC2022-133370-I00 and TED2021-132020B-I00 funded by MCIN/AEI/10.13039/501100011033 and by European Union Next GenerationEU/PRTR. Savíns Puertas Martín is a fellow of the “Margarita Salas” grant (RR_A_2021_21), financed by the European Union (NextGenerationEU).

References

- Ahmad, S.U., Bergen, S.W.A. (2010). A genetic algorithm approach to the inverse problem of treatment planning for intensity-modulated radiotherapy. *Biomedical Signal Processing and Control*, 5(3), 189–195. <https://doi.org/10.1016/j.bspc.2010.03.001>.

- Alber, M., Nüssli, F. (1999). An objective function for radiation treatment optimization based on local biological measure. *Physics in Medicine & Biology*, 44(2), 479–493. <https://doi.org/10.1088/0031-9155/44/2/014>.
- Bortfeld, T. (1999). Optimized planning using physical objectives and constraints. *Seminars in Radiation Oncology*, 9(1), 20–34. [https://doi.org/10.1016/S1053-4296\(99\)80052-6](https://doi.org/10.1016/S1053-4296(99)80052-6).
- Breedveld, S., Heijmen, B. (2017). Data for TROTS – The Radiotherapy Optimisation Test Set. *Data in Brief*, 12, 143–149. <https://doi.org/10.1016/j.dib.2017.03.037>.
- Breedveld, S., Craft, D., van Haveren, R., Heijmen, B. (2019). Multi-criteria optimization and decision-making in radiotherapy. *European Journal of Operational Research*, 277(1), 1–19. <https://doi.org/10.1016/j.ejor.2018.08.019>.
- Cho, B. (2018). Intensity-modulated radiation therapy: a review with a physics perspective. *Radiation Oncology Journal*, 36(1), 1–10. <https://doi.org/10.3857/roj.2018.00122>.
- Choi, B., Deasy, J.O. (2002). The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Physics in Medicine and Biology*, 47(20), 3579–3589. <https://doi.org/10.1088/0031-9155/47/20/302>.
- Durillo, J.J., Nebro, A.J. (2011). jMetal: a Java framework for multi-objective optimization. *Advances in Engineering Software*, 42, 760–771. <https://doi.org/10.1016/j.advengsoft.2011.05.014>.
- Ehrgott, M., Guler, C., Hamacher, H.W., Shao, L. (2010). Mathematical optimization in intensity modulated radiation therapy. *Annals of Operations Research*, 175(1), 309–365. <https://doi.org/10.1007/s10479-009-0659-4>.
- Fu, A., Ungun, B., Xing, L., Boyd, S. (2019). A convex optimization approach to radiation treatment planning with dose constraints. *Optimization and Engineering*, 20, 277–300. <https://doi.org/10.1007/s11081-018-9409-2>.
- Huang, C., Nomura, Y., Yang, Y., Xing, L. (2022). Meta-optimization for fully automated radiation therapy treatment planning. *Physics in Medicine & Biology*, 67(5). <https://doi.org/10.1088/1361-6560/ac5672>.
- Jourdan, L., Basseur, M., Talbi, E.-G. (2009). Hybridizing exact methods and metaheuristics: a taxonomy. *European Journal of Operational Research*, 199(3), 620–629. <https://doi.org/10.1016/j.ejor.2007.07.035>.
- Kaliszewski, I., Miroforidis, J., Podkopaev, D. (2016). *Multiple Criteria Decision Making by Multiobjective Optimization – A Toolbox*. Springer. <https://doi.org/10.1007/978-3-319-32756-3>.
- Moreno, J.J., Miroforidis, J., Filatovas, E., Kaliszewski, I., Garzón, E.M. (2021). Parallel radiation dose computations with GENOCOP III on GPUs. *Journal of Supercomputing*, 77, 66–76. <https://doi.org/10.1007/s11227-020-03254-6>.
- Moreno, J.J., Miroforidis, J., Kaliszewski, I., Garzón, E.M. (2022). Parallel EUD models for accelerated IMRT planning on modern HPC platforms. In: *14th International Conference on Parallel Processing and Applied Mathematics (PPAM 2022)*. https://doi.org/10.1007/978-3-031-30445-3_12.
- Mukherjee, S., Hong, L., Deasy, J., M., Z. (2020). Integrating soft and hard dose-volume constraints into hierarchical constrained IMRT optimization. *Medical Physics*, 47(2), 414–421. <https://doi.org/10.1002/mp.13908>.
- Niemierko, A. (1996). Reporting and analyzing dose distributions: a concept of equivalent uniform dose. *Medical Physics*, 24(1), 103–110. <https://doi.org/10.1118/1.598063>.
- NVIDIA (2023). CUDA Toolkit Documentation v11.7.1. <https://docs.nvidia.com/cuda/index.html>. Last Accessed on 2023/11/17.
- Ólafsson, A., Wright, S.J. (2006). Linear programming formulations and algorithms for radiotherapy treatment planning. *Optimization Methods and Software*, 21(2), 201–231. <https://doi.org/10.1080/10556780500134725>.
- Ólafsson, A., Jeraj, R., Wright, S.J. (2005). Optimization of intensity-modulated radiation therapy with biological objectives. *Physics in Medicine and Biology*, 50(22), 5357–5379. <https://doi.org/10.1088/0031-9155/50/22/010>.
- OpenMP (2023). The OpenMP API specification for parallel programming. <https://www.openmp.org/>. Last Accessed on 2023/11/17.
- Romeijn, H., Dempsey, J., Li, J. (2004). A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology*, 49(10), 1991–2013. <https://doi.org/10.1088/0031-9155/49/10/011>.
- Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A. (2006). A new linear programming approach to radiation therapy treatment planning problems. *Operations Research*, 54(2), 201–216. <https://doi.org/10.1287/opre.1050.0261>.
- Saberian, F., Ghate, A., Kim, M. (2015). Optimal fractionation in radiotherapy with multiple normal tissues. *Mathematical Medicine and Biology: A Journal of the IMA*, 33(2), 211–252. <https://doi.org/10.1093/imammb/dqv015>.

- Starzyński, J., Szmurlo, R., Chaber, B., Krawczyk, Z. (2015). Open access system for radiotherapy planning. In: *2015 16th International Conference on Computational Problems of Electrical Engineering (CPEE)*, pp. 204–206. <https://doi.org/10.1109/CPEE.2015.7333376>.
- Wang, H., Bai, X., Wang, Y., Lu, Y., Wang, B. (2023a). An integrated solution of deep reinforcement learning for automatic IMRT treatment planning in non-small-cell lung cancer. *Frontiers in Oncology*, 13. <https://doi.org/10.3389/fonc.2023.1124458>.
- Wang, Q., Wang, R., Liu, J., Jiang, F., Yue, H., Du, Y., Wu, H. (2023b). High-dimensional automated radiation therapy treatment planning via Bayesian optimization. *Medical Physics*, 50(6), 3773–3787. <https://doi.org/10.1002/mp.16289>.
- Won, Y.J., Lee, E., Cho, S.J., Kim, K.S. (2021). Generalized equivalent uniform dose-based biological optimization in hippocampus-sparing whole-brain radiation therapy. *Journal of the Korean Physical Society*, 79, 1171–1179. <https://doi.org/10.1007/s40042-021-00321-w>.
- Wu, Q., Mohan, R., Niemierko, A., Schmidt-Ullrich, R. (2002). Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology Biology Physics*, 52, 224–235. [https://doi.org/10.1016/S0360-3016\(01\)02585-8](https://doi.org/10.1016/S0360-3016(01)02585-8).
- Yihua Lan, Y., Li, C., Ren, H., Zhang, Y., Min, Z. (2006). Fluence map optimization (FMO) with dose–volume constraints in IMRT using the geometric distance sorting method. *Physics in Medicine and Biology*, 57(20), 201–231. <https://doi.org/10.1088/0031-9155/57/20/6407>.
- ZeroMQ (2023). ØMQ – The Guide. <https://zguide.zeromq.org/>. Last Accessed on 2023/11/17.
- Zhang, Q., Li, H. (2007). MOEA/D: a multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6), 712–731. <https://doi.org/10.1109/TEVC.2007.892759>.

J.J. Moreno is a doctoral researcher in the Department of Informatics at the University of Almería (UAL), Spain. He obtained his BSc and MSc in computer science from UAL in 2015 and 2018, respectively. Since 2017, he has been actively engaged in research as a member of the Supercomputing–Algorithms research group at UAL. His areas of research expertise encompass High-Performance Computing (HPC), radiotherapy optimization, and electron tomography image processing.

S. Puertas-Martín is a post-doctoral researcher at the Department of Informatics of the University of Almería, Spain. He obtained his PhD in computer science at the University of Almería in 2020. He is a member of the Supercomputing–Algorithms Research Group at the University of Almería in Spain and the Chemoinformatics Research Group at the University of Sheffield in the United Kingdom. His research interests are drug discovery, global optimization and high-performance computing.

J.L. Redondo holds the position of full professor at the Department of Informatics at the University of Almería, Spain. She earned her PhD in computer science from the same university and is an esteemed member of the Supercomputing–Algorithms Research Group. Her research focuses on High-Performance Computing, global optimization, and their diverse applications.

P.M. Ortigosa is a full professor of architecture and computer technology of the University of Almería, Spain. She received MSc degrees in physics and electronic engineering from the University of Granada in 1994 and 1996, respectively, and a PhD in computer science from the University of Málaga in 1999. She is a member of the Supercomputing–Algorithms Research Group of the University of Almería. Her research focuses on high-performance computing, metaheuristic global optimization, computational intelligence,

deep learning, and the application to several real problems. Recently she has been working on the Internet of Things.

A. Zawadzka, a distinguished medical physicist, currently serves in the Medical Physics Department at the Maria Skłodowska-Curie National Research Institute of Oncology. In addition to her role, she holds the position of head of the Treatment Planning Laboratory within the same institution. Her academic journey led to the successful completion of a PhD in medical physics from the Maria Skłodowska-Curie National Research Institute of Oncology in 2018. A. Zawadzka is actively engaged as a lecturer at Warsaw University and contributes as a regional consultant in the field of medical physics. Her scientific pursuits are concentrated on treatment planning, dosimetry, and the quality control of radiotherapy treatments.

P. Kukołowicz serves as a medical physicist at the Medical Physics Department of the Maria Skłodowska-Curie National Research Institute of Oncology, an institution with a rich history dating back to 1934. As a prominent figure, he holds the position of the head of this esteemed institution. Notably, he has been the president of the Polish Society of Medical Physics from 2011 to 2014 and then again from 2014 to 2018. Since 2018, he is a consultant in medical physics for the Ministry of Health. P. Kukołowicz is also recognized for his role as an advisor to 14 PhD candidates in the field of medical physics. His scientific pursuits primarily revolve around treatment planning, dosimetry, and the quality control of radiotherapy treatments.

R. Szmurło holds the position of an assistant professor at the Warsaw University of Technology, Poland. He has actively contributed to over 10 research projects, funded by entities such as the Polish Ministry of Science, commercial enterprises, and European Union grants, all related to computer simulation methods. The subjects of the projects were related to modelling applications of electromagnetic fields in medical treatment, modelling electric impulse power supply systems, artificial intelligence in medicine for Radiotherapy planning, methods of modelling information systems, among others.

I. Kaliszewski is a full professor in the Systems Research Institute of the Polish Academy of Sciences and in Warsaw School of Information Technology. His scientific interests are optimization, multiple criteria decision making, computer-aided decision making, and also identification, quantification, and management of risk in business organizations.

J. Miroforidis is an associate professor at the Systems Research Institute of the Polish Academy of Sciences, Poland. He obtained his PhD in computer science from the same institution in 2010. His scientific interests are multi-objective optimization, multiple criteria decision making, and computer-aided decision making.

E.M. Garzón is a full professor at the Department of Informatics of the University of Almería, Spain. She obtained her PhD in computer science from the University of Almería in 2000. She is the head of the Supercomputing-Algorithms Research Group at the same institution. Her research activity is centred on high-performance computing (HPC) addressed to extend applications of scientific computation. Her works have been related to different fields and disciplines.