

Trading robustness: a scenario-free approach to robust Multi-Criteria Optimization for Treatment Planning

Remo Cristoforetti^{a,b,c}, Philipp Süß^d, Tobias Becher^{a,b,c}, Niklas Wahl^{a,b}

^aDepartment of Medical Physics in Radiation Oncology, German Cancer Research Center – DKFZ, Im Neuenheimer Feld 280, 69120 Heidelberg, Germany

^bHeidelberg Institute for Radiation Oncology – HIRO, Im Neuenheimer Feld 280, 69120 Heidelberg, Germany

^cFaculty of Physics and Astronomy, Heidelberg University, Heidelberg, Germany

^dFraunhofer Institute for Industrial Mathematics (ITWM), Kaiserslautern, Germany

Version 1, typeset October 15, 2025

Corresponding author(s):

Remo Cristoforetti: remo.cristoforetti@dkfz-heidelberg.de

Philipp Süß:

Tobias Becher:

Niklas Wahl: n.wahl@dkfz-heidelberg.de

Abstract

Background: Treatment planning in radiotherapy is inherently a multi-criteria optimization (MCO) problem, as it requires balancing competing clinical goals. Traditionally, the treatment’s robustness is not formulated as a part of this decision making problem, but dealt with separately through margins or robust optimization.

Purpose: This work facilitates integration of robustness into multi-criteria optimization using a recently proposed efficient “scenario-free” (s-f) robust optimization approach: Utilizing variance reduction objectives, whose computation is independent of the number of chosen error scenarios, robustness can become part of the multi-criteria decision making process at minimal computational overhead.

Methods: The s-f approach relies on the fast evaluation of the expected dose distribution and mean variance during optimization independent of the scenario number. This is achieved by precomputation of expected dose influence and total variance influence matrices, which can then be used for repeated solving of subproblems in the two explored MCO approaches: Lexicographic Ordering (LO) and full Pareto Front (PF) approximation. Different prioritization strategies within the LO approach are used to assess the impact of variance reduction on the optimization outcome. A 3-objective PF approximation, including a variance reduction objective, is generated to visualize and analyze trade-offs between the competing objectives.

The robust optimization is performed including 100 scenarios modeling setup and range errors, as well as organ motion, on 3D- and 4DCT lung cancer patient datasets. Robustness analysis is performed to assess and explore the efficacy of all optimization strategies.

Results: The s-f approach enabled robust optimization in MCO with computational times comparable to nominal MCO. Both MCO strategies highlighted the interplay between dosimetric and variance reduction objectives. The LO approach showed how prioritization affects plan quality and robustness, while the PF analysis revealed a clear trade-off between robustness and organ-at-risk sparing.

Conclusions: The proposed s-f robust optimization approach allowed the efficient application of robust MCO by significantly reducing the required computational time. The reported analysis highlighted the conflicting trade-off nature of plan robustness and dosimetric quality, demonstrating how robust MCO supports a more informed and flexible decision-making process in treatment planning.

I Introduction

The primary objective of radiotherapy treatment planning is to achieve an optimal balance between clinical goals, including adequate coverage of the target volume and minimal exposure of the surrounding organs at risk (OARs). As clinicians must navigate trade-offs between these competing clinical goals, the treatment planning problem inherently represents a multi-criteria decision making problem. In treatment planning, this problem has been traditionally approached by representing the trade-offs as a weighted sum of objectives to be optimized with scalar optimization techniques. In practice, this often results in a tedious iterative tuning process to find clinically acceptable plans whose quality further depends on the planners' experience^{1, 2}. To systematize the navigation of the planning trade-offs, more sophisticated multi-criteria optimization (MCO) approaches have been introduced into clinical workflows³⁻⁹. These methods offer advantages over traditional weighted-sum approaches by enabling consistent and automated navigation of a priori goals or by allowing a fast and interactive a posteriori navigation of available solutions^{5, 7}. Common examples for a priori methods are hierarchical or lexicographic optimization^{3, 10}, while the most common a posteriori method relies on Pareto surface approximation^{4, 8}.

Beyond the challenges of decision-making, the treatment plan quality can be significantly affected by uncertainties such as anatomical variations, patient setup or range errors¹¹⁻¹⁴. According to the specific radiation modality and treatment site, robust optimization techniques can effectively reduce the required margin size and the integral dose to the patient^{15, 16}.

However, traditional robust optimization approaches involve evaluating multiple error scenarios during both dose calculation and optimization which makes them computationally expensive. At the same time, standard MCO methods, typically require the sequential generation of multiple plans, exacerbating the computational complexity of the problem and scaling the optimization time and memory footprint to unfeasibility.

Conventional robust optimization algorithms rely on well defined mathematical metrics

such as max value operators, for min-max optimization approaches or expected value operators for probabilistic planning^{15, 16}. Despite their rigorous formulation, predicting the impact of these operators on the final plan remains challenging, and their effectiveness closely relates to the accuracy of the uncertainty model employed. These operators influence the problem formulation in binary manner, i.e once applied, they are enforced fully on all solutions.

Two fundamental limitations thus hinder the integration of traditional robust optimization with Multi Criteria Optimization (MCO). The first is a practical limitation posed by the computational cost of a full scenario based approach. The second is a conceptual limitation posed by the rigid application of robustness operators, which prevents the direct and flexible control over how robust optimization influences the final solution.

Several approaches tackled the first practical limitation in both optimization and robustness analysis. Examples include Polynomial Chaos Expansion^{17, 18} and Analytical Probabilistic Models (APM)^{19, 20}, which utilize analytical methods to propagate uncertainty through dose calculation. These techniques efficiently estimate the first moments of the dose for each voxel, and eliminate the explicit calculation of all error scenarios during optimization. However, they often still entail significant computational costs associated with constructing the analytical model itself.

A solution to overcome the second limitation, and enable robust MCO, was proposed in literature²¹ where traditional robust operators are applied to selected objectives. A combination of robust and non robust objectives builds the Pareto front approximation allowing the control of plan robustness through trade-off between these objectives. This approach offers a viable solution to the conceptual limitations of robust optimization, but still bares the computational burden of explicit individual scenario evaluation during optimization.

As an alternative approach, the APM framework itself suggests incorporating variance reduction objectives alongside standard dosimetric ones in the optimization problem formulation. In this case robustness is quantified through variance as a well defined metric, and supports a convex optimizable objective function.

Variance can be estimated during optimization through the explicit evaluation of all error scenarios, nonetheless bearing again the same computational cost. Alternatively, the recently proposed *scenario-free* approach to robust optimization enables the natural and efficient integration of mean variance reduction objectives²². This efficiency arises from precalculation of probabilistic quantities over the set of error scenarios, avoiding explicit evaluation of each scenario during optimization.

Including variance reduction objectives in MCO using the scenario-free approach can potentially address both limitations. It enables explicit flexible control over dose variance at minimal computational cost, effectively introducing robustness as a criterion in the decision making process.

In this work we propose and investigate the application of the scenario-free approach to two MCO techniques, namely a Lexicographic Ordering (LO) and a Pareto front approximation and navigation approach. The LO strategy is a so-called (*2pec*) approach^{7, 10} while the Pareto front approximation is obtained through a sandwiching algorithm^{23–25}. Several

optimization configurations are explored to investigate the impact of variance reduction as a decision making criterion for the different strategies.

II Materials and Methods

II.A Multi Criteria Optimization

For the purpose of this work, the MCO problem can be stated as

$$\begin{aligned} \min_{\mathbf{x}} \quad & \vec{F}(\mathbf{x}) = (F_1(\mathbf{x}), \dots, F_n(\mathbf{x})) \\ \text{s.t.} \quad & \mathbf{x} > 0 \\ & c_k(\mathbf{x}) \leq 0 \end{aligned} \quad , \quad (1)$$

where $\mathbf{x}$ is the set of optimization variables, usually representing the fluence weights and $\vec{F}(\mathbf{x})$ are the competing cost functions in vectorized form: Each $F_i(\mathbf{x})$ component represents an objective in the treatment plan and constitutes a dimension in the objective function space. The problem also includes an arbitrary number of $c_k(\mathbf{x})$ constraints that limit the feasibility space. Solutions for which no F_i can improve without degrading at least one other F_j are called Pareto efficient, or Pareto optimal.

Solution of problem (1) through standard optimization algorithms requires a scalarization of the vectorized cost function $\vec{F}(\mathbf{x})$. The standard approach consists in performing a linear combination of the objectives and minimize $F(\mathbf{x}) = \sum_i w_i F_i(\mathbf{x})$ upon assignment of relative importance weights w_i . The weight vector $\vec{w}$ is assigned a priori based on the clinical relevance of the associated objective. The minimal value of the cost function obtained this way corresponds to a single point in the objective space, thus to a single optimal plan, and its position on the Pareto surface is eventually determined by the weight vector itself. In practice, a clinically relevant plan is achieved heuristically tuning the importance vector $\vec{w}$.

II.A.1 A posteriori approach: Pareto Front navigation

This “trial-and-error” process of tuning weights, waiting for the next solution, and repeating with an adapted set of weights $\vec{w}$, may practically be replaced by an automated strategy to generate plans with different sets of weights $\vec{w}$. The method of choosing $\vec{w}$ can be formalized, for example using sandwiching algorithms, to adequately cover the Pareto front. Combined with a navigation technique to interpolate between generated points, the clinical decision is then performed a-posteriori. The planner can explore the tradeoffs between objectives and their impact on the plan metrics by intuitive graphical user interfaces.

The accuracy of the surface approximation so of the interpolated plans, is proportional to the number of calculated points. Given the computational cost of optimizing a single plan, the set of $\vec{w}$ vectors should be appropriately chosen to achieve sufficient accuracy with minimal number of optimizations.

II.A.2 A priori approach: Lexicographic Ordering

A second approach to the scalarization of problem (1) consists in applying a lexicographic ordering strategy. Contrary to the Pareto front approximation and navigation, the clinical decision is performed a-priori by means of a well defined hierarchical ordering of the objectives which are addressed sequentially. Each objective is optimized individually and constraining all the other objectives to their previously achieved value. Formally, problem 1 is scalarized as:

$$\begin{aligned} \min_{\mathbf{x}} \quad & F_i(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x} > 0 \\ & c_k(\mathbf{x}) \leq 0 \\ & F_j(\mathbf{x}) \leq \epsilon_j \end{aligned} \quad , \quad (2)$$

for all $j < i$ and $\epsilon_j = F_j(\mathbf{x}^*)$ is the previously achieved value, i.e. $\mathbf{x}^*$ is the solution achieved for $j = i$. This way objective F_j does not deteriorate upon optimization of objective F_i .

In particular, the (*2pec*) approach applied in this work performs the optimization in two steps, or phases. The first phase optimizes each objective up to a prescribed clinical goal b_i assessing its reachability. During this phase ϵ_j in 2 is defined as:

$$\epsilon_j = \begin{cases} b_j & F_j(\mathbf{x}^*)\delta < b_j \\ F_j(\mathbf{x}^*)\delta & F_j(\mathbf{x}^*)\delta \geq b_j \end{cases}$$

where δ is a slack variable introduced for numerical stability. The second phase optimizes each objective to their best and tunes the optimality of the solution. In this phase, all other objectives are constrained, the index j spans $\{1, \dots, n\} \setminus i$ and $\epsilon_j = F_j(\mathbf{x}^*)\delta$.

Both techniques allow to explore the interaction between the optimization objectives $F_i(\mathbf{x})$ and assess their impact on the optimized plan.

II.B The scenario-free approximation to robust optimization

Traditional robust optimization approaches may be expressed as applying a robustness operator to the objective function $\vec{F}(\mathbf{x})$. Such operators usually rely on a set of individual error scenarios that need to be explicitly evaluated in each iteration during optimization. This implies significant computational overhead in the application of such robust optimization techniques as the optimization time for each problem to be solved scales with the number of included scenarios. The complexity of the problem scales thus with the number of scenarios, which is exacerbated in the case of MCO approaches requiring optimization of multiple subproblems.

Furthermore, the application of a robustness operator fully enforces attention to the modeled uncertainty. The effect of robust optimization on the dose distribution strongly depends on the nature of the operator used and whether the resulting plan is overly optimistic or overly conservative directly derives from this choice. In an MCO context, the optimization of robust objectives modifies the shape of the Pareto surface, but not the dimensionality of

the objective space itself. An optimization trajectory in this space is thus limited to only access robust solutions when robust optimization is applied, and only non-robust solutions when it is not, preventing the exploration of intermediate solutions.

As an alternative approach to robust optimization, the scenario-free approach²² enables efficient evaluation of the expected dose distribution $\mathbb{E}[\mathbf{d}]$ and mean variance $\mathbb{V}_m[\mathbf{d}(\mathbf{x})]$ during optimization. These quantities are estimated from precomputed probabilistic quantities through:

$$\mathbb{E}[\mathbf{d}] = \mathbb{E}[\mathbf{D}]\mathbf{x}, \quad \text{with: } \mathbb{E}[\mathbf{D}] = \sum_s p_s \mathbf{D}^s \quad (3)$$

$$\mathbb{V}_m[\mathbf{d}] = \frac{1}{N} \mathbf{x}^T \mathbf{\Omega} \mathbf{x}, \quad \text{with: } \mathbf{\Omega} = \mathbb{E}[\mathbf{D}^T \mathbf{D}] - \mathbb{E}[\mathbf{D}]^T \mathbb{E}[\mathbf{D}] \quad (4)$$

where $\mathbb{E}[\mathbf{D}]$ and $\mathbf{\Omega}$ are the expected-dose- and a variance- influence matrix, respectively. Both matrices are generated during dose calculation from the full set of error scenarios. This way no individual scenario information is retained during optimization. p_s is a probability weight associated to each scenario s . It is determined by the probability distribution modeling the uncertainty source, together with the adopted sampling method.

With the scenario-free approach, conventional dosimetric objectives are evaluated on the expected dose distribution. The optimization cost function is formulated as a composite of the dosimetric objective and the mean variance term:

$$\mathcal{F}(x) = f(\mathbb{E}[\mathbf{d}(\mathbf{x})]) + \mathbb{V}_m[\mathbf{d}(\mathbf{x})] \quad (5)$$

This combination is motivated by the formal equivalence to the traditional expected-value robust optimization approach when the least-squares objective function is considered.

However, the variance reduction term can be treated as a separate independent variance reduction objective. This enables its integration into the clinical decision-making process, with its full potential most effectively explored through an MCO approach.

Additionally, accumulating multiple variance-influence matrices on selected sets of error scenarios allows to separate the variance contribution of different sources of uncertainty. In an MCO framework this allows tradeoffs not only of variance against dosimetric goals, but also among variance goals for different uncertainty models.

II.C MCO Robust Optimization Strategies

Integration of the scenario-free robust optimization approach with MCO allows for several optimization strategies to be designed. With the a priori LO approach, the importance of variance reduction is controlled through a predefined priority rank. The a posteriori Pareto front approximation instead, provides full flexibility to evaluate the impact of variance reduction on the plan.

II.C.1 LO strategies

The initial exploration of trade-offs with the LO approach is performed including only setup and range uncertainties. Three prioritization strategies are defined to investigate the role of variance reduction in the optimization process. A fourth strategy is introduced to address target motion uncertainty through a 4D variance reduction objective.

For the first strategy, both dosimetric and variance reduction objectives are combined at every priority level. This enables direct comparison with a non-robust, margin-based approach, in which the PTV is generated by expanding the CTV with an isotropic margin and only the nominal planning CT is considered. Such comparison allows the validation of the proposed method and the estimation of relative optimization times of the scenario-free robust LO.

The second strategy prioritizes dosimetric goals over variance reduction objectives, ensuring that the dosimetric quality of the plan is not compromised by the introduction of variance reduction. In this configuration, robustness represents an additional feature of the plan that is achieved on top of, rather than instead of, dosimetric quality.

The third strategy places target variance reduction above all other objectives, explicitly favoring robustness for target coverage.

For the fourth strategy two variance reduction steps follow the dosimetric optimization. The two objectives address two different sources of uncertainty, namely target motion/anatomical changes and setup and range errors. The separation is achieved by accumulating the Ω variance-influence matrices on separate sets of error scenarios, each encoding the desired uncertainty source.

II.C.2 Pareto Front Approximation

A robust Pareto front approximation is generated considering three combined objectives: two dosimetric objectives for OARs sparing and one variance reduction objective for the target. This allows for a three-dimensional representation of the Pareto front to be visualized for descriptive purposes. The Pareto surface is obtained using the total expected- and variance-influence matrices accounting for the combined setup, range, and breathing motion uncertainties. Multiple points are sampled on the Pareto surface using the sandwiching algorithm to create a convex-hull approximation.

II.D Treatment planning application

II.D.1 Objective functions and clinical goals

Table 1 presents a collection of objective functions considered in this work. The selection of such objectives for each ROI was guided by the clinical interpretability of the goal values applied for the LO approach, as well as the intention to demonstrate that the method's effectiveness is not limited to a specific set of objectives.

The variance reduction objective was formulated as minimization of the variance term in eq. (5). The equivalent uniform dose (EUD) objective depends on the n parameter as

defined in table 1. This parameter encodes the volume effect of the specific organ and for parallel OARs, values of $n \approx 1$ are usually applied^{26, 27}. An higher value of n increases the penalty and the attention of the optimizer for higher doses. In this work a value of $n = 3.5$ was chosen. This choice allows to maintain a clear clinical interpretation of the goal value associated to this objective while inducing an stronger reduction of higher doses in the lung. This choice additionally ensures the convexity of the EUD function²⁸.

Table 1: Objective function definitions, relative parameters and prescriptions. The objective functions are applied in different configurations according to the specific optimization strategy applied. The $\hat{\cdot}$ operator applied to d^i represents the *mean* operator.

Objective Type	Functional form	Parameters	Goal
Max Dose	$d_{max} + \log(\sum_i e^{\frac{d_i - d_{max}}{\epsilon d_{max}}})$	$\epsilon = 1e - 3$	60 Gy
EUD	$\left[\sqrt[n]{\frac{1}{N} \sum_i d_i^n} - d_{ref} \right]^2$	$d_{ref} = 0 \text{ Gy}, n = 3.5$	$[40 \text{ Gy}]^2$
Mean Dose	$ \hat{d}_i - d_{ref} $	$d_{ref} = 0 \text{ Gy}$	15 Gy
Squared Deviation	$\frac{1}{N} \sum_i (d_i - d_{ref})^2$	$d_{ref} = 60 \text{ Gy}$	$[1 \% \cdot d_{ref} \text{ Gy}]^2$
Variance reduction	$\mathbb{V}_m[d_i] - v_{pres}$	$v_{pres} = 0 \text{ Gy}$	1 Gy^2
EUD + Sq. Overdosing	$\sqrt[n]{\frac{1}{N} \sum_i d_i^n} - d_{ref}^{EUD} + \frac{1}{N} \sum_i (d_i - d_{ref}^o)^2 \Theta(d_i - d_{ref}^o)$	$d_{ref}^{EUD} = 0$ $d_{ref}^o = 20 \text{ Gy}$	

Different trade-offs between such objectives are explored using the aforementioned robust optimization strategies.

Table 2 collects the priority lists designed for the strategies (Strategy 1-4). Variance reduction for the target structure is prioritized differently for each strategy described in section II.C.1, balancing its relative importance against the dosimetric objectives for the OARs. For all strategies, an additional minimum dose constraint is applied to the target at 95 % of the prescribed dose.

For the Pareto Front approximation, multiple objective functions for the same structure are combined and treated as single objectives (table 1). For the OARs, this includes a least-squares deviation on the EUD plus a squared overdose objective on dose, forming a single objective. A variance minimization term is applied on the target. Additional minimum and maximum dose constraints are applied to the target to ensure coverage and avoid hot spots. The minimum and maximum target doses are set to 50 Gy and 55 Gy.

II.D.2 Patient datasets

In this work, a lung cancer patient dataset was selected from the cancer imaging archive²⁹. The dataset consists of a 4D-CT scan with 10 phases representing the breathing motion. This choice allows for the separate investigation of the different considered sources of uncertainty, particularly of setup and range uncertainties as well as anatomical motion.

For the investigation of setup and range errors, only the nominal planning CT, corresponding to the deep breath-hold position, was considered. Setup errors were modeled as

rigid shifts of the isocenter in all three spatial directions, sampled from normal distributions with zero mean and a 2.25 mm standard deviation. Range errors were modeled as relative and absolute deviations and sampled from normal distributions with zero mean and standard deviations of 3.5 % and 1 mm, respectively.

100 combined error scenarios were generated to compute the probabilistic quantities. These quantities allow to estimate an expected dose distribution and a variance reduction objective referred to in the following as 3D-expected-dose and 3D-variance. This set of probabilistic quantities was used to investigate the LO optimization Strategies (1-3).

To address the uncertainty arising from breathing motion, a second Ω matrix was computed including only the 4D-CT phases. The resulting variance reduction objective $\mathbb{V}_m^{4D}[\mathbf{d}]$ quantifies thus the uncertainty associated with breathing motion.

An expected dose influence matrix was also accumulated including the combined dataset of 3D error scenarios and 4D-CT phases. This matrix allows to estimate the expected dose of the combined dataset. The LO strategy 4 was applied considering this latter expected dose influence matrix for the application of the dosimetric objectives. The variance reduction objectives were instead applied through the $\mathbb{V}_m^{4D}[\mathbf{d}]$ and $\mathbb{V}_m^{3D}[\mathbf{d}]$ operators.

A third dataset was generated by combining setup and range uncertainties across all 4D-CT phases. In this case, 10 setup and range errors were sampled for each breathing phase, resulting in a total of 100 scenarios. The 10 setup and range error scenarios are the same for each of the 10 CT phases, assuming thus a full correlation between 3D scenarios across phases. Given the nature of the uncertainty sources this is closer to a realistic treatment setup where the setup and range uncertainty reflect the interfractional positioning variability and systematic over- and undershoot. These scenarios were combined into a single expected dose influence and variance influence matrix, capturing the combined variation induced by all three sources.

All treatment plans were generated using the same irradiation setup, consisting of a three-field proton plan with gantry angles fixed at 45°, 90°, and 135°. A constant RBE value of 1.1 was assumed for all optimizations, and no rescaling of the dose distribution was applied. Additionally, the PTV margin for the non robust plan was set to 4 mm.

II.E Robustness Analysis

Robustness analysis was conducted for each optimized plan using the complete set of scenarios employed during the pre-calculation of the respective probabilistic quantities. In particular, the 100 3D scenario pool was used to evaluate the robustness of LO strategies (1-4). For the forth LO strategy, the robustness analysis was additionally performed on the 10 4D-CT phases dataset to investigate the impact of breathing motion alone, and on the combined total dataset of 3D and 4D scenarios, i.e the total 110 scenarios. The robustness analysis for the Pareto surface was conducted for all sampled solutions on the set of correlated 3D and 4D scenarios.

Plan quality and robustness were evaluated through multiple metrics, including dose distributions, dose-volume histograms (DVHs), standard deviation distributions, and stan-

dard deviation volume histograms (SDVHs). An SDVH provides a relationship between the standard deviation of the delivered dose and volume. DVH points are also reported for a selection of plans. Unless otherwise indicated, the reported values refer to the median value of the DVH metric on the set of error scenarios and the relative inter quartile range (IQR).

II.F Algorithm implementation

Both the scenario-free approach and the LO approach were implemented in the open-source treatment planning toolkit matRad³⁰. matRad served as the development environment for individual scenario dose calculations, accumulation of probabilistic quantities, and the final robustness analysis performed on all obtained solutions.

In contrast, the sandwiching algorithm for Pareto front approximation and the corresponding navigation tool were developed by the Fraunhofer ITWM^{31, 32} and incorporate the same probabilistic quantities computed in matRad.

Table 2: Lexicographic ordering strategies applied to investigate the impact of variance reduction on the optimized plan. The strategies differ from each other in the prioritization grade of the variance reduction objective. The fourth and fifth column report the quantities on which the given objective functions are applied at the corresponding step. The minimum dose constraint was applied to all the strategies.

Constraints					
Structure		Type	Parameters		
Target		Min Dose	95 % · Prescription		
Objectives					
	Priority	Objective Type	Structure	Rob	ITV
Strategy 1	1	Max Dose	Heart	$\mathbb{E}[\boldsymbol{d}]$	$\boldsymbol{d}$
	2	EUD	Lung		
	3	Mean Dose	Heart		
	4	Squared Deviation	Target		
Strategy 2	1	Max Dose	Heart	$\mathbb{E}[\boldsymbol{d}]$	
	2	EUD	Lung		
	3	Mean Dose	Heart		
	4	Squared Deviation	Target	$\mathbb{V}_m[\boldsymbol{d}]$	
	5	Variance	Target		
Strategy 3	1	Variance	Target	$\mathbb{V}_m[\boldsymbol{d}]$	
	2	Max Dose	Heart		
	3	EUD	Lung		
	4	Mean Dose	Heart	$\mathbb{E}[\boldsymbol{d}]$	
	5	Squared Deviation	Target		
Strategy 4	1	Max Dose	Heart	$\mathbb{E}[\boldsymbol{d}]$	
	2	EUD	Lung		
	3	Mean Dose	Heart		
	4	Squared Deviation	Target	$\mathbb{V}_m^{4D}[\boldsymbol{d}]$	
	5	Variance	Target		
	6	Variance	Target		

III Results

III.A LO approach

All the reported results for the LO approach refer to the second phase of the *2pec* algorithm.

III.A.1 Strategy 1: Algorithm validation

The first set of results focuses on validating the scenario-free robust LO approach. Figure 1, Strategy 1 presents a comparative analysis between the margin-based and the robust plans at the last step of the optimization sequence as reported in table 2. The robust plan integrates both dosimetric and variance reduction objectives.

The DVHs include the target, left lung and heart structures. The reported solid line DVHs are computed from the expected dose distribution, while DVH bands span the 5th to 95th percentiles of the distribution of DVHs computed across all scenarios. A representative axial slice of the expected dose distribution is also reported for both algorithms.

The robust plan achieves a markedly improved robustness within the target, as demonstrated by the narrower DVH band relative to the nominal plan. In addition to improved target uniformity, the robust plan also results in better sparing of OARs. Notably, the integral dose to the lung is lower in the robust plan. The V_{20} and V_{30} values for the lung in the nominal plan are (45.9%) and (27.3%), respectively, whereas the corresponding values for the robust plan are (44.6%) and (26.2%), indicating a reduction in high- and mid-dose volumes within the lung. All the measured IQRs for these DVH points were below 3.5%.

Finally, the relative optimization times for all steps range between 1.5 and 1.6. This indicates that the s-f approach is approximately 1.5 times slower than the nominal optimization. Such overhead is introduced by the additional computation of the variance operator. Variance estimation is performed at each optimization step as a constraint value on the variance is always applied during the second phase of the *2pec* approach.

III.A.2 Strategy 2: Prioritization of dosimetric objectives

The second prioritization strategy focuses initially on dosimetric objectives alone. To highlight the impact of variance reduction, only the final dosimetric step (Step 4) and the first variance reduction step (Step 5) are reported for the robust plan in fig. 1. The robust plan at Step 4 in this analysis differs from the nominal plan presented for Strategy 1 (fig. 1) only in that dosimetric objectives are applied to the expected dose distribution rather than to the nominal one. In Strategy 2, (fig. 1) the DVHs and expected dose distributions are shown on the left, while the standard deviation (SD) distribution and SDVHs are presented on the right.

At Step 4, the DVH plot still displays a relatively wide DVH band, particularly for the target, indicating that substantial uncertainty remains at this stage. This is further corroborated by the SDVH plot, where the SDVH for the nominal plan—corresponding to Strategy 1, Step 4 in fig. 1—is reported with a dashed line for reference. The SDVH for




Figure 1: Comparison between optimized plans for 3 different LO optimization strategies. (Strategy 1, Step 4) Comparison between the nominal plan (left) and the scenario-free robust plan (right) for the last optimization step of the Strategy 1. (Strategy 2) Comparison between DVHs and expected dose distribution (left) and SDVHs and standard deviation maps (right) for the last dosimetric (Step 4) and variance reduction (Step 5) steps of Strategy 2. (Strategy 3) DVHs, expected dose distribution, SDVHs and standard deviation distribution of the last step of the optimization Strategy 3. For the DVH plots, the solid line corresponds to the DVH computed for the expected dose distribution, while the dotted lines correspond to the 25-75 percentiles. The colored DVH band spans the 5-95 percentiles over the scenarios distribution. For the SDVH of Strategy 2, solid lines for the target correspond to the SDVH for the corresponding step, strategy and plan, the dotted lines serve as a comparison reference for the nominal plan (Step 4) and the previous step (Step 5). All colorbar values are reported in Gy.

the robust plan is shown with the solid green line and differs only slightly from the nominal counterpart.

Variance minimization for the target structure is introduced in Step 5. The DVH band for the target becomes narrower, and the SDVH curve shifts to the left, indicating reduced standard deviation. For comparison, the SDVH for Step 4 of the same strategy is retained as a dashed line in the SDVH plot. No major changes are observed in the dose distribution between Steps 4 and 5. The SD distribution shows instead reduction of high standard deviation values and a more homogeneous pattern within the target.

III.A.3 Strategy 3: Prioritization of Variance reduction objectives

The third prioritization strategy places variance reduction within the target volume as the primary objective. Figure 1, Strategy 3 presents the DVH and SDVH analyses for the robust plan at the last steps of the optimization process. A highly robust dose distribution is achieved for the target as early as the first step of optimization and minimal variance is maintained throughout the remaining optimization steps.

The dose distribution reported for this strategy shows a lower target conformity and higher OARs doses compared to the results presented in section (III.A.1) and (III.A.2). At the same time the SD distributions and SDVH curves highlight how the standard deviation is highly reduced in the target compared to the previous approaches.

III.A.4 Strategy 4: 4D robustness

The fourth optimization strategy was designed to address setup and range (3D) uncertainties separately from breathing motion (4D) uncertainties. Figure 2 illustrates the robustness analysis associated with this strategy, focusing on the variance reduction achieved in Steps 5 and 6. In Step 5, 4D variance reduction is performed, and the robustness analysis was conducted on both the 3D-only and the 4D-only scenario datasets. For both robustness analyses, the figure also reports the SDVH for the target structure at the previous dosimetric-only step (Step 4) for direct comparison.

Both analyses present a reduction in variance for the target structure. The 4D analysis highlights a considerable reduction in target uncertainty introduced by the 4D-variance reduction objective. Concurrently, the 3D robustness analysis for the same Step also presents a contained reduction in target variance.

Figure 2 also presents the DVH and SDVH analysis for the final optimization step (Step 6), in which the 3D variance reduction is performed. At this stage, the robustness analysis was performed on the combined dataset of 3D and 4D scenarios, thereby accounting for the total uncertainty arising from both sources. For comparison, the SDVH plot also includes the curves obtained for the target, on the same dataset, in Steps 4 and 5, corresponding to dosimetric-only optimization and 4D-only variance reduction, respectively. Across these steps, the SDVH curve progressively shifts to the left, reflecting the reduction in uncertainty.




Figure 2: Robustness analyses performed for the Step 5 and Step 6 of the LO Strategy 4. Top left: SDVH analysis and SD distribution computed on the 3D-only set of error scenarios. The green dashed line corresponds to the target SDVH obtained at Step 4 for the same strategy and 3D error scenario dataset. Top right: SDVH and SD analysis for the same step but performed for the 4D scenario dataset only including the nominal CT-phases. Bottom: DVH and expected-dose distribution (left) and SDVH and SD distribution (right) performed for Step 6 on the complete combined set of 3D and 4D error scenarios. All the reported values are in Gy.

III.B Pareto Front approximation

A three-dimensional Pareto front approximation was constructed using the objective selection outlined in table 1. The resulting approximation, illustrated in fig. 3(b), represents the trade-offs among competing objectives within a normalized objective space $[0,1]$. The plans were sampled through the sandwiching algorithm, for a total of 15 solutions. Each solution is displayed as a black dot on the three-dimensional Pareto surface. To facilitate comparison among objectives, all bi-objective projections of the Pareto front are also shown (fig. 3(a,c,d)). In addition to the surface representation, fig. 3 reports a DVH and SDVH analysis for 25 solutions sampled on the Pareto surface. The solutions were sampled through convex interpolation of the surface along the target variance reduction axis. The solutions are color coded according to the objective function value for the target variance, from the lowest variance (yellow) to the highest (blue). A strong correlation between the objective values of the different solutions is observed, as the solutions with the lowest variance are the solutions with the worst DVHs for both the OARS.

Figure 3(h) further highlights the correlation among solutions through a scatter plot for two selected DVH and SDVH metrics, namely V_5 for the heart and SD_{50} for the target. SD_{50} corresponds to the standard deviation value received by 50% of the target volume. To further highlight the correlation among the DVH metrics, the points are color coded according to the value of a third DVH metric, namely D_{20} for the lung.

Figure 4 reports the robustness analysis for two selected plans, chosen to highlight the extremes of the target variance dimension. The first plan, referred to as Solution 1, corresponds to the solution with the lowest variance reduction objective for the target. Its




Figure 3: Representation of the Pareto front (a-d) for the considered case and correlation analysis for interpolated solutions on the surface (e-h). Black dots on the complete surface (b) and the three projections (a,c,d) correspond to points obtained through the sandwiching algorithm. The yellow and blue circles in (a-d) highlight Solution 1 and 2 respectively. SDVHs for the target (e), DVHs for the heart (f) and lung (g) for solutions obtained through interpolation of the Pareto surface. DVH and SDVH lines with the same color correspond to the same solution. Scatter plot (h) relating SD_{50} for the target and V_5 for the heart for the same set of sampled solutions. The set is color coded according to the D_{20} metric for the lung of the corresponding solution.

DVH band is tightly clustered around the expected dose distribution, indicating both high robustness and improved homogeneity within the target.

In contrast, Solution 2 represents the solution with the highest variance in the target. This plan exhibits a wider DVH band and reduced target homogeneity but achieves larger sparing of the OARs, with lower DVH curves for both the heart and lungs compared to Solution 1.

The SDVH curves provide additional insight into these differences. For the target, the SDVH in Solution 1 is significantly lower than in Solution 2, confirming its superior robustness. However, Solution 2 shows a lower SDVH for the heart, in accordance with the observed decrease in its DVH curve. The SD distribution for Solution 1 also reflects a more homogeneous and generally lower standard deviation within the target structure.

Both solutions are highlighted by yellow (Solution 1) and blue (Solution 2) circles on the Pareto surface in fig. 3(a-d).

From a multi-objective optimization perspective, Solution 1 is the second-worst plan in terms of heart sparing (Heart objective = 0.99) and fourth-worst in terms of lung sparing (Lung objective = 0.94).

When the Pareto front is projected onto the plane defined by the heart and lung objectives (fig. 3(d)), Solution 2 lies closest to the origin, i.e., it is the point with the lowest Euclidean norm. In other words, the solution with the best compromise between heart and lung objectives is the solution with the worst variance for the target.




Figure 4: Comparison between DVHs and dose distribution (left), SD distribution and SDVHs (right) for Solution 1 (top) and Solution 2 (bottom). For the DVH plots, the solid line represents the DVH computed for the expected dose distribution, the dotted lines correspond to the 25-75 % percentiles of the DVH distribution over the error scenarios, and the colored band corresponds to the 5-95 % percentiles. All reported values are in Gy.

These trade-offs can be explored interactively through interpolation of the Pareto surface, allowing for dynamic assessment of alternative solutions. This functionality is enabled by a graphical decision-support tool developed by the Fraunhofer ITWM . The tool includes real-time visual updates of objective projections and DVH comparisons.

IV Discussion

The scenario-free (s-f) approach enables the efficient evaluation of both the expected dose distribution and the mean variance term during optimization. While these quantities could also be computed using traditional robust optimization methods and evaluating all the included scenarios, such approach is typically hindered by high computational demands and memory limitations. In MCO workflows, where multiple sequential optimizations are performed, these computational costs accumulate, making the use of traditional robust optimization techniques impractical. In contrast, the s-f approach offers a fast alternative for achieving robust optimization within MCO, requiring optimization times comparable to those of nominal, non-robust approaches.

Traditional approaches allow only a binary application of robust optimization, meaning that the robustness component is either fully included in the problem formulation or entirely excluded. The choice of the robustness operator or margin dictates the nature of the optimized plan, which may become too optimistic or too conservative depending on the settings^{15, 33, 34}. Because of this rigidity, the influence of robust optimization on the final dose distribution can only be adjusted by adapting the uncertainty source³⁵⁻³⁷, or, when possible, by adopting compromising strategies³⁸. However, the selection of an uncertainty model should be guided by evidence and clinical practice^{39, 40} whenever sufficient data is available.

In contrast, the scenario-free approach incorporates variance reduction as a criterion in the decision making process. This enables explicit control of plan robustness through variance reduction, which is not bounded to any a-priori choice of robustness operator. Moreover, if the underlying uncertainty model does not accurately represent the real, undetermined, probability distribution of the error source, the scenario-free approach still provides a consistent solution to practically manage the impact of such uncertainty on the final dose distribution.

At the same time, the explicit incorporation of variance reduction terms into the cost function eq. (5) can strongly influence both the definition of optimization objectives and the optimization trajectory itself. Determining appropriate prescription values and the relative weighting of variance reduction objectives is not straightforward and poses a conceptual challenge.

Thus, the contribution of this work is twofold. First, it demonstrates that robust MCO can be achieved with limited overhead compared to nominal MCO through the s-f approach, owing to its computational efficiency and flexibility. Second, by embedding variance reduction objectives within an MCO framework, this study enables systematic exploration of their effects on treatment planning, offering valuable insights into the behavior and impact of the s-f method.

To evaluate the performance and implications of the approach, two MCO strategies were applied and analyzed in detail.

IV.A Prioritization strategies for the LO approach

The lexicographic ordering framework enables the explicit assignment of priority levels to planning objectives, offering a clear and clinically interpretable structure. LO also relies on the assumption that all objectives can be ranked. The ranking priority may be clear for traditional dosimetric goals, when the DM relies on extensive planning experience, but is less clear for variance reduction objectives. In this work, we explored four prioritization strategies to examine how different objective orderings affect the resulting treatment plans.

The first strategy combined dosimetric and variance reduction objectives at each optimization step. The resulting robust plan demonstrated a reduced sensitivity to uncertainty compared to the margin-based nominal plan and achieved improved sparing of organs at risk. Additionally, optimization times remained comparable to those of nominal optimization, supporting the feasibility of the s-f approach in clinical workflows. A direct comparison between the scenario-free robust and a traditional robust optimization algorithm in a multi-criteria optimization context is challenging due to the computational demand of the latter.

This first approach however, served as a validation of the proposed method. It reflects a more conventional clinical application of LO, where the order of objectives is inherited from the nominal optimization and the robust objectives are obtained by replacement of their nominal counterpart. In this configuration, prioritization is primarily driven by structure- and objective-based hierarchy rather than by a deliberate choice between dosimetric and robustness objectives.

Alternatively, the s-f robust LO approach allows for the systematic exploration of alternative prioritization strategies that shift the balance between dosimetric and variance reduction goals. To investigate this, two contrasting strategies were considered. The first strategy (Strategy 2) prioritized dosimetric objectives and variance reduction is only performed at last, without compromising the achieved dosimetric goals. The second strategy (Strategy 3) reverses this priority by placing variance reduction at the top of the objective hierarchy.

For the first strategy (Strategy 2), a comparison between the robust and nominal plan at the end of the dosimetric phase effectively illustrates the difference between optimizing for expected versus nominal dose distributions. In the examined case, no substantial differences between the nominal and expected dose optimization were observed. It should be noted that all the presented results pertain to the second phase of the $2pec$ approach, in which all objectives are constrained to their goal values or previously optimized outcomes. Consequently, the optimization of dosimetric objectives for the robust plan is anyway conducted under loose variance constraints, set here to an optimistic value of 1 Gy^2 .

The subsequent variance reduction step effectively reduces the variance within the targeted structure thus proving the efficacy of the variance reduction objective in limiting the impact of uncertainty while still maintaining the achieved dosimetric plan quality.

For the second strategy (Strategy 3), variance reduction alone applied at the first step would push the dose to the target toward zero. This is prevented by the hard constraint posed on the minimum target dose at 95 % of the prescription.

This prioritization strategy yields an exceptionally robust dose distribution for the target, which remains stable across all subsequent optimization steps. However, this robustness comes at the expense of target conformity and reduced sparing of nearby organs at risk, highlighting the trade-offs inherent in such prioritization.

Together, these two opposing strategies exemplify the flexibility afforded by the s-f LO framework. They demonstrate how variance reduction objectives can be treated on equal footing with dosimetric goals, allowing the prioritization configuration to be fully customized for each individual patient case. This opens the possibility for a broader spectrum of optimization strategies, tailored to clinical intent and case-specific trade-offs between robustness and plan quality.

The analysis in fig. 2 illustrates how the scenario-free approach combined with lexicographic optimization allows the exploration of robustness against different sources of uncertainty.

This analysis demonstrates how the 4D variance reduction applied with Step 5 effectively reduces the uncertainty generated by breathing motion when it is assessed on the 4D scenario set. 4D variance minimization reduces the standard deviation for 3D scenarios as well, demonstrating how the two sources of uncertainty are not completely independent. Correlation among the uncertainty sources translates into correlation between the objectives and further demonstrates how variance reduction interacts as any other objective in the treatment plan. For this patient case and the magnitude of the chosen setup and range

errors, the uncertainty on the dose distribution introduced by breathing motion was lower compared to the uncertainty introduced by setup and range errors. This is reflected both by the standard deviation distributions and by the SDVH curves reported in fig. 2.

For the analysis of Step 6, the uncertainty is estimated across all 3D and 4D scenarios and both variance reduction steps contribute to a partial reduction of the overall uncertainty.

Regardless of the specific prioritization strategy, each step in the LO sequence allows for a controlled relaxation of the previously optimized objectives. This is managed through the introduction of the slack variable δ ,⁷, which ensures numerical stability during the optimization process. No significant degradation in the quality of any previously optimized objective was observed across the analyzed cases, for any of the applied Strategies.

While the LO approach provides a structured and clinically interpretable prioritization of objectives, it effectively traces only a single path through the objective space and along the Pareto surface. In contrast, a full Pareto front approximation enables the exploration of the entire trade-off landscape between competing objectives, offering an extended view of the available planning solutions.

IV.B Pareto front approximation

This analysis focused on a three-dimensional objective space, allowing full visualization of the Pareto front and direct exploration of how variance reduction objectives affect the overall plan quality. The representation clearly demonstrates how robustness, especially in the target, trades off against sparing of OARs.

Two extreme solutions along the target variance axis were highlighted. These exhibited significant differences in DVHs, DVH bands, SDVHs, and dose and standard deviation distributions. The most robust solution for the target (Solution 1) was among the least favorable for heart and lung sparing, while the best compromise plan for the OARs (Solution 2) presented the poorest target robustness. This inverse correlation reflects the optimizer’s behavior: robust target coverage requires expanding the high-dose region, thus reducing conformity and increasing dose to surrounding tissue; conversely, maximizing OAR sparing relies on a tighter dose distribution and steeper dose gradients, reducing thus robustness.

Correlation among objectives translates to correlation among DVH metrics as well. The analysis reported in fig. 3(e-h) for the interpolated solutions highlights how variability in the SDVH curves for the target translates to variability in the DVH curves for the two organs at risk. For this specific case, variability among DVH curves for the lung was lower than variability among DVH curves for the heart. The scatter plot reported in fig. 3(h) further enhances how metrics for target variance and organ sparing anti-correlate. When sampling the solutions along the variance reduction direction, intrinsic correlation among the OARs DVH metrics is also observed, as both of them improve when the variance objective worsen.

All sampled solutions lie on the Pareto front, guaranteeing Pareto efficiency and enabling the decision-maker to explore meaningful trade-offs between competing objectives. Combining this approach with the computational speed of the s-f framework allows efficient construction of the Pareto front approximation, effectively bringing robust Pareto front ex-

ploration within the reach of clinical use. The interactive decision-support tool further enhances this process by enabling real-time exploration of the Pareto surface, DVHs, and dose distributions, supporting informed clinical choices.

Although this analysis focused on three objectives only, realistic clinical scenarios involve higher-dimensional spaces. The extension of the approach to include additional dimensions poses no conceptual nor practical challenge. The sandwiching algorithm can handle the additional dimensions by sampling more solutions on the Pareto Front. Thanks to the s-f approach’s efficiency, the sampling of additional plans remains feasible within practical time constraints.

IV.C Summary and Outlook

The LO approach encodes precise clinical priorities directly into the problem formulation, offering a structured and guided optimization pathway. However, this inherently constrains the exploration of the objective space. In contrast, Pareto front approximation facilitates a broader exploration, allowing trade-offs among competing objectives to be fully visualized and evaluated. Depending on the clinical context, both strategies can support the decision-making process: some scenarios benefit from the clarity and structure of LO, while others may require the flexibility and comprehensiveness of Pareto-based exploration.

Hybrid strategies that combine the strengths of both approaches are also possible. For instance, a lexicographic, clinically informed ordering could be applied to the dosimetric objectives, followed by a Pareto front approximation performed solely in the variance reduction space. This would allow robustness to be balanced across different structures, or different sources of uncertainty, without compromising initial dosimetric goals. Additionally, the dual nature of the problem – addressing both dose and variance – suggests the potential for a bi-level optimization formulation. Recent work demonstrated the effectiveness of bi-level approaches for dose and LET optimization, particularly when supported by an efficient approximation of the domination cone⁴¹. A similar formulation could be explored in the context of dose and variance optimization as well.

V Conclusions

This work demonstrates the novel application of a scenario-free robust optimization approach to multi-criteria optimization. Successful integration into two MCO strategies — a lexicographic ordering approach and a full Pareto front approximation — was achieved. These approaches allowed to explore the impact of variance reduction objectives on the overall plan quality.

This approach effectively incorporates robustness as a criterion into the decision-making process, making it an integral, optimizable feature of the treatment plan. Notably, the observed optimization times remain comparable to those of a nominal MCO approach.

VI Acknowledgement

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), projects no. 443188743 and 457509854, is acknowledged.

VII Conflict of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

1. Kierkels RGJ, Visser R, Bijl HP, et al. Multicriteria Optimization Enables Less Experienced Planners to Efficiently Produce High Quality Treatment Plans in Head and Neck Cancer Radiotherapy. *Radiation Oncology (London, England)*. 2015;10:87.
2. Bohsung J, Gillis S, Arrans R, et al. IMRT Treatment Planning:- a Comparative Inter-System and Inter-Centre Planning Exercise of the ESTRO QUASIMODO Group. *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology*. 2005;76(3):354–361.
3. Zarepisheh M, Hong L, Zhou Y, et al. Automated and Clinically Optimal Treatment Planning for Cancer Radiotherapy. *INFORMS journal on applied analytics*. 2022;52(1):69–89.
4. Hong TS, Craft DL, Carlsson F, and Bortfeld TR. MULTICRITERIA OPTIMIZATION IN IMRT TREATMENT PLANNING FOR LOCALLY ADVANCED CANCER OF THE PANCREATIC HEAD. *International journal of radiation oncology, biology, physics*. 2008;72(4):1208–1214.
5. Breedveld S, Craft D, van Haveren R, and Heijmen B. Multi-Criteria Optimization and Decision-Making in Radiotherapy. *European Journal of Operational Research*. 2019;277(1):1–19.
6. Teichert K, Currie G, Küfer KH, et al. Targeted Multi-Criteria Optimisation in IMRT Planning Supplemented by Knowledge Based Model Creation. *Operations Research for Health Care*. 2019;23:100185.
7. Breedveld S, Storchi PRM, and Heijmen BJM. The Equivalence of Multi-Criteria Methods for Radiotherapy Plan Optimization. *Physics in Medicine & Biology*. 2009;54(23):7199.
8. Craft DL, Hong TS, Shih HA, and Bortfeld TR. Improved Planning Time and Plan Quality Through Multicriteria Optimization for Intensity-Modulated Radiotherapy. *International Journal of Radiation Oncology*Biophysics*. 2012;82(1):e83–e90.
9. Thieke C, Küfer KH, Monz M, et al. A New Concept for Interactive Radiotherapy Planning with Multicriteria Optimization: First Clinical Evaluation. *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology*. 2007;85(2):292–298.

10. Breedveld S, Storchi PRM, Voet PWJ, and Heijmen BJM. iCycle: Integrated, Multi-criterial Beam Angle, and Profile Optimization for Generation of Coplanar and Non-coplanar IMRT Plans. *Medical Physics*. 2012;39(2):951–963.
11. Lomax AJ. Intensity Modulated Proton Therapy and Its Sensitivity to Treatment Uncertainties 1: The Potential Effects of Computational Uncertainties. *Physics in Medicine & Biology*. 2008;53(4):1027.
12. Lomax AJ. Intensity Modulated Proton Therapy and Its Sensitivity to Treatment Uncertainties 2: The Potential Effects of Inter-Fraction and Inter-Field Motions. *Physics in Medicine and Biology*. 2008;53(4):1043–1056.
13. Liebl J, Paganetti H, Zhu M, and Winey BA. The Influence of Patient Positioning Uncertainties in Proton Radiotherapy on Proton Range and Dose Distributions. *Medical Physics*. 2014;41(9):091711.
14. Kinoshita N, Shimizu M, Motegi K, et al. Quantification of Uncertainties in Reference and Relative Dose Measurements, Dose Calculations, and Patient Setup in Modern External Beam Radiotherapy. *Radiological Physics and Technology*. 2025;18(1):58–77.
15. Fredriksson A. A Characterization of Robust Radiation Therapy Treatment Planning Methods—from Expected Value to Worst Case Optimization. *Medical Physics*. 2012;39(8):5169–5181.
16. Unkelbach J, Alber M, Bangert M, et al. Robust Radiotherapy Planning. *Physics in Medicine and Biology*. 2018;63(22):22TR02.
17. Perkó Z, van der Voort SR, van de Water S, Hartman CMH, Hoogeman M, and Lathouwers D. Fast and Accurate Sensitivity Analysis of IMPT Treatment Plans Using Polynomial Chaos Expansion. *Physics in Medicine and Biology*. 2016;61(12):4646–4664.
18. Hoogeman MS, Water S van de, Voort S van der, Perko Z, Heijmen B, and Lathouwers D. Fast and Accurate Sensitivity Analysis of Random and Systematic Errors Using Polynomial Chaos Expansion in IMPT. *International Journal of Radiation Oncology, Biology, Physics*. 2015;93(3):S149.
19. Bangert M, Hennig P, and Oelfke U. Analytical Probabilistic Modeling for Radiation Therapy Treatment Planning. *Physics in Medicine and Biology*. 2013;58(16):5401–5419.
20. Wahl N, Hennig P, Wieser HP, and Bangert M. Analytical Probabilistic Modeling of Dose-Volume Histograms. *Medical Physics*. 2020;47(10):5260–5273.
21. Chen W, Unkelbach J, Trofimov A, et al. Including Robustness in Multi-Criteria Optimization for Intensity-Modulated Proton Therapy. *Physics in Medicine and Biology*. 2012;57(3):591–608.
22. Cristoforetti R, Hardt JJ, and Wahl N. Scenario-Free Robust Optimization Algorithm for IMRT and IMPT Treatment Planning. *Medical Physics*. 2025;52(7):e17905.
23. Rennen G, van Dam E, and den Hertog D. Enhancement of Sandwich Algorithms for Approximating Higher Dimensional Convex Pareto Sets. 2009. URL: <https://papers.ssrn.com/abstract=1427721> (visited on 06/26/2023). Pre-published.

24. Craft DL, Halabi TF, Shih HA, and Bortfeld TR. Approximating Convex Pareto Surfaces in Multiobjective Radiotherapy Planning. *Medical Physics*. 2006;33(9):10.
25. Bokrantz R and Forsgren A. An Algorithm for Approximating Convex Pareto Surfaces Based on Dual Techniques. *INFORMS Journal on Computing*. 2013;25(2):377–393.
26. Matuszak M, Matrosic C, Jarema D, et al. Priority Driven Plan Optimization in Locally Advanced Lung Patients Based on Perfusion SPECT Imaging. *Advances in Radiation Oncology*. 2016;1.
27. Allen Li X, Alber M, Deasy JO, et al. The Use and QA of Biologically Related Models for Treatment Planning: Short Report of the TG-166 of the Therapy Physics Committee of the AAPM: TG-166 Report. *Medical Physics*. 2012;39(3):1386–1409.
28. Choi B and Deasy JO. The Generalized Equivalent Uniform Dose Function as a Basis for Intensity-Modulated Treatment Planning. *Physics in Medicine & Biology*. 2002;47(20):3579.
29. Hugo GD, Weiss E, Sleeman WC, Balik S, and Keall PJ. Data from 4D Lung Imaging of NSCLC Patients (Version 2). 2016.
30. Wieser HP, Cisternas E, Wahl N, et al. Development of the Open-Source Dose Calculation and Optimization Toolkit matRad. *Medical Physics*. 2017;44(6):2556–2568.
31. Lammel I, Küfer KH, and Süß P. Efficient Approximation Quality Computation for Sandwiching Algorithms for Convex Multicriteria Optimization. *Journal of Optimization Theory and Applications*. 2025;204(3):41.
32. Collicott C, Bonacker E, Lammel I, Teichert K, Walczak M, and Süß P. Interactive Navigation of Multiple Convex Patches. *Journal of Multi-Criteria Decision Analysis*. 2021;28(5–6):311–321.
33. Fredriksson A and Bokrantz R. A Critical Evaluation of Worst Case Optimization Methods for Robust Intensity-Modulated Proton Therapy Planning. *Medical Physics*. 8Part1 2014;41:081701.
34. Zaghian M, Cao W, Liu W, et al. Comparison of Linear and Nonlinear Programming Approaches for "Worst Case Dose" and "Minmax" Robust Optimization of Intensity-Modulated Proton Therapy Dose Distributions. *Journal of Applied Clinical Medical Physics*. 2017;18(2):15–25.
35. Badiu V, Trier Taasti V, Defraene G, van Elmpt W, and Sterpin E. Balancing Robustness and Adaptation Rate for Proton Therapy of Lung Cancer Patients. *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology*. 2024;196:110290.
36. Tarp IS, Taasti VT, Jensen MF, Vestergaard A, and Jensen K. Benefit of Range Uncertainty Reduction in Robust Optimisation for Proton Therapy of Brain, Head-and-Neck and Breast Cancer Patients. *Physics and Imaging in Radiation Oncology*. 2024;31:100632.

-
37. Van de Water S, van Dam I, Schaart DR, Al-Mamgani A, Heijmen BJM, and Hoogeman MS. The Price of Robustness; Impact of Worst-Case Optimization on Organ-at-Risk Dose and Complication Probability in Intensity-Modulated Proton Therapy for Oropharyngeal Cancer Patients. *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology*. 2016;120(1):56–62.
 38. Sevilla AC, Cabal G, Wahl N, Puerta ME, and Rivera JC. A Robust Optimization Model for Intensity-Modulated Radiotherapy: Cheap-Minimax. *Medical Physics*. 2025;52(5):3360–3376.
 39. Van Herk M. Errors and Margins in Radiotherapy. *Seminars in Radiation Oncology*. High-Precision Radiation Therapy of Moving Targets. 2004;14(1):52–64.
 40. Fracchiolla F, Widesott L, Righetto R, et al. Development and Clinical Application of a Probabilistic Robustness Evaluation Tool for Pencil Beam Scanning Proton Therapy Treatments. *Physica Medica*. 2025;131:104938.
 41. Schubert M. *Reoptimization of Multi-Criteria Proton Therapy Plans Including Linear Energy Transfer and Relative Biological Effectiveness*. Fraunhofer Verlag, 2025.