

Direct incorporation of patient-specific efficacy and toxicity estimates in radiation therapy plan optimization

Running title: RT planning using personalized models

Daniel F Polan PhD¹, Marina Epelman PhD², Victor W Wu PhD², Yilun Sun PhD^{1,3}, Markus Varsta DSc⁴, Daniel R Owen PhD¹, David Jarema BS¹, Charles K Matrosic PhD¹, Shruti Jolly MD¹, Caitlin A Schonewolf MD¹, Matthew J Schipper PhD^{1,3}, and Martha M Matuszak PhD¹

¹Department of Radiation Oncology, University of Michigan, Ann Arbor, MI, 48109, USA

²Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, 48109, USA

³Department of Biostatistics, University of Michigan, Ann Arbor, MI, 48109, USA

⁴Varian Medical Systems, Palo Alto, CA, 94304, USA

Corresponding Author: Daniel Polan, polandan@med.umich.edu

Abstract

Purpose: Current radiation therapy (RT) treatment planning relies mainly on pre-defined dose-based objectives and constraints to develop plans that aim to control disease while limiting damage to normal tissues during treatment. These objectives and constraints are generally population-based, in that they are developed from the aggregate response of a broad patient population to radiation. However, correlations of new biologic markers and patient-specific factors to treatment efficacy and toxicity provide the opportunity to further stratify patient populations and develop a more individualized approach to RT planning. We introduce a novel intensity-modulated radiation therapy (IMRT) optimization strategy that directly incorporates patient-specific dose response models into the planning process. In this strategy, we integrate the concept of utility-based planning where the optimization objective is to maximize the predicted value of overall treatment utility, defined by the probability of efficacy (e.g., local control) minus the weighted sum of toxicity probabilities. To demonstrate the feasibility of the approach, we apply the strategy to treatment planning for non-small cell lung cancer (NSCLC) patients.

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as [doi: 10.1111/mp.15940](https://doi.org/10.1111/mp.15940).

This article is protected by copyright. All rights reserved.

Methods and materials: We developed a prioritized approach to patient-specific IMRT planning. Using a commercial treatment planning system (TPS), we calculate dose based on an influence matrix of beamlet-dose contributions to regions-of-interest. Then, outside of the TPS, we hierarchically solve two optimization problems to generate optimal beamlet weights that can then be imported back to the TPS. The first optimization problem maximizes a patient's overall plan utility subject to typical clinical dose constraints. In this process, we facilitate direct optimization of efficacy and toxicity trade-off based on individualized dose-response models. After optimal utility is determined, we solve a secondary optimization problem that minimizes a conventional dose-based objective subject to the same clinical dose constraints as the first stage but with the addition of a constraint to maintain the optimal utility from the first optimization solution. We tested this method by retrospectively generating plans for five previously treated NSCLC patients and comparing the prioritized utility plans to conventional plans optimized with only dose metric objectives. To define a plan utility function for each patient, we utilized previously published correlations of dose to local control and grade 3–5 toxicities that include patient age, stage, microRNA levels, and cytokine levels, among other clinical factors.

Results: The proposed optimization approach successfully generated RT plans for five NSCLC patients that improve overall plan utility based on personalized efficacy and toxicity models while accounting for clinical dose constraints. Prioritized utility plans demonstrated the largest average improvement in local control (16.6%) when compared to plans generated with conventional planning objectives. However, for some patients the utility-based plans resulted in similar local control estimates with decreased estimated toxicity.

Conclusion: The proposed optimization approach, where the maximization of a patient's RT plan utility is prioritized over the minimization of standardized dose metrics, has the potential to improve treatment outcomes by directly accounting for variability within a patient population. The implementation of the utility-based objective function offers an intuitive, humanized approach to biological optimization in which planning trade-offs are explicitly optimized.

Key words: IMRT, radiotherapy outcomes, hierarchical optimization

1. Introduction

Advances in radiation therapy (RT) planning and delivery techniques have led to improved outcomes for cancer patients, both in terms of increased control rates and decreased treatment toxicity.¹⁻³ However, balancing disease control with adverse therapy effects continues to be a challenge in RT and across many other cancer treatment modalities. In external beam RT, this trade-off is inherent because delivery of radiation to the tumor often requires the use of beam trajectories that traverse adjacent normal tissue structures prior to and after converging on the target.⁴ Consequently, providing a sufficient therapeutic dose to the tumor may result in radiation-induced normal tissue toxicities, which can negatively impact a patient's quality-of-life during and after treatment.¹⁻⁵ Modern RT techniques, such as intensity-modulated radiation therapy (IMRT), have made it possible to deliver highly non-uniform dose distributions with steep dose gradients between target locations and organs-at-risk (OARs). While this has enabled improved OAR sparing, optimizing an RT plan to

best balance the individualized potential for disease progression and treatment toxicities remains hindered by conventional dose-metric-based planning techniques.

Conventional RT treatment planning attempts to address these trade-offs by optimizing radiation delivery based on predefined dose-based objectives that have either been previously correlated with measures of control and toxicity or determined from expert consensus and institutional standards. In inverse IMRT treatment planning, these surrogates of biological response are often incorporated into a single objective function through individually weighted linear or quadratic penalties, and the fluence map is optimized to minimize the sum of these penalties.⁶ During this process, treatment guidelines and acceptable clinical trade-offs are implicitly translated and incorporated into the optimization objective by adjusting the relative importance of each dose-based penalty, relying heavily on the planner's experience and expertise to guide the process.⁷ Although a substantial amount of research has focused on weight selection to improve plan quality,⁸⁻¹⁶ this optimization approach is limited by the inability to directly include individualized, quantitative estimates of the biological outcomes in the optimization process.

Recently, advancements in optimization approaches and availability of commercial products for biological optimization have facilitated the move beyond conventional dosimetric planning, allowing for the inclusion of tumor control probability (TCP) and normal tissue toxicity probability (NTCP) models in the plan optimization process.¹⁷ Although implementations of available biological optimization methods differ, they are generally restricted in the ability to use both estimated biological response objectives and conventional dose-based objectives to drive the optimization process.¹⁷⁻¹⁹ Instead, dose-based objectives may only be regarded as hard constraints in the optimization process, potentially limiting the ability to decrease dose to OARs with undefined biological response objectives.¹⁷⁻¹⁹ Additionally, these systems continue to rely on aggregated response models and do not account for patient-to-patient variability in tissue radiosensitivity. Recent studies have focused on addressing this through the inclusion of personalized models in the planning process.²⁰⁻²³ While these studies have demonstrated the ability to optimize treatments based on patient-specific outcome predictions, they remain limited to three-dimensional conformal radiotherapy (3DCRT) optimizing the beam angles and monitoring units²⁰⁻²² or prescription.²³ Additionally, these methods do not explicitly incorporate currently accepted clinical dose constraints which may impede clinical acceptance or diminish the realized benefit of the planning strategies.

To overcome these limitations, we propose an IMRT optimization method, termed Prioritized Utility Optimization (PUO), that augments the traditional dosimetric inverse treatment planning process by directly incorporating quantitative estimates of personalized biological response while maintaining current dose-based planning objectives. This builds upon previous work that introduced the concept of utility-based planning, in which individualized optimal dose was selected by maximizing plan utility defined by a weighted combination of predicted efficacy and toxicity probabilities.²⁴⁻²⁵ These studies demonstrated the ability of utility-based planning to increase efficacy while maintaining or reducing toxicity levels in a population but were limited to scaling of fixed dose distributions. In the proposed IMRT optimization approach, we translate this methodology into a constrained hierarchical optimization problem that prioritizes maximization of plan utility over minimization of typical dose metrics and is subject to clinical hard constraints throughout. Through this approach, we

add the ability to redistribute dose and aim to further exploit predictive models of efficacy and toxicity based on biological markers, clinical factors, and patient demographics, to improve the personalization of planning trade-offs. To the authors' knowledge, this is the first study to propose a prioritized approach to biologically based treatment planning. This approach attempts to overcome the current limitations of biological optimization techniques by improving user-interpretability of the biological objective and optimization process, directly integrating tradeoff exploration between predictive outcome models, and allowing for the inclusion of dose-based constraints and objectives. Overall, the methodology outlined in this study aims to provide patient-specific outcome-based treatment planning without compromising currently accepted dose-based treatment planning standards. In this study, we introduce and evaluate the feasibility of the proposed method, PUO, by applying it to cohort of previously treated non-small cell lung cancer (NSCLC) patients.

2. Materials and Methods

2.A. Optimization Overview

Our proposed approach utilizes a two-stage optimization process to first maximize plan utility and then minimize dose-based metrics. Figure 1 provides the mathematical formulation of the optimization problems and Table I provides a full description of the notation used in defining the problems. In Stage 1, the objective is to maximize plan utility defined by the predicted efficacy probability based on target dose, minus the weighted sum of predicted toxicity probabilities across multiple OARs based on respective OAR doses. The efficacy and toxicity models are assumed to be monotonic functions of a dose metric, with increasing dose to the structure never decreasing the predicted probability of the event occurring to that structure. The weighting parameter for each toxicity, ϑ , represents the undesirability of toxicity relative to the efficacy measure and can be adjusted based upon physician and/or patient preference, or tuned to result in an acceptable rate of predicted toxicity across a patient population.²⁵

The maximization of plan utility in Stage 1 is subject to inviolable clinical criteria for the plan, guaranteeing that any solution meets dose-based hard constraints (1b–1f). These constraints are based on the following three rationales. First, for any individual patient, increasing predicted efficacy can only occur by increasing dose to the target. Since a single target dose metric is used to predict efficacy, maximization of utility through increasing efficacy may result in unacceptable target dose heterogeneity. This has previously been noted as a potential issue in biological optimization methods and is addressed through constraint (1b).¹⁷ Second, it is possible that optimal solutions exist at dose values above or below levels that physicians would be comfortable prescribing. Hard constraints (1c) and (1d) are used to avoid these solutions and avoid extrapolating model predictions to dose values

beyond those in the dataset used to fit the models. Lastly, it is unlikely that every dose-limiting OAR will have a toxicity model corresponding to every dose metric typically used to limit OAR doses during planning. Therefore, constraints (1d–1f) are included to ensure that any solution remains within traditional OAR dose limits. To avoid non-convex functions of dose, dose-volume constraints (DVC), which are typically represented by value-at-risk (VaR) metrics (e.g., V_{20Gy} , $D_{0.1cc}$), are instead represented by upper conditional value-at risk (CVaR⁺) metrics in constraint (1f). CVaR⁺ represents a convex DVC that captures the *mean* upper-tail dose of a structure's dose-volume histogram and has previously been used to formulate linear RT optimization problems.²⁶⁻²⁸

In Stage 2, the objective is to minimize dose-based metrics similar to traditional plan optimization approaches. Stage 1 (biological objective) and Stage 2 (dose-based objective) are intentionally separated to eliminate the need to weight outcomes directly against dose metrics and maintain user interpretability of the utility objective and outcome weighting factors in Stage 1. Although the utility objective is prioritized, the inclusion of dose-based constraints in Stage 1 prevents the solution from compromising physical dose requirements. However, Stage 1 effectively only optimizes the OAR dose-metrics used in toxicity models for the utility function. Therefore, the plan resulting from the Stage-1 optimization solution may not be Pareto optimal across all clinically relevant OARs and corresponding dose metrics, and dose metrics not used in the toxicity models can still potentially be reduced. Therefore, Stage 2 can be considered a solution “clean-up” or refinement step, in which all OAR doses are reduced while attempting to maintain the optimal plan utility from Stage 1. For computational efficiency, the objective function is a sum of linearized equivalent uniform dose (ℓEUD). ℓEUD is a convex piecewise-linear (PWL) approximation of generalized equivalent uniform dose and is calculated as a linear weighted combination of the mean and maximum doses for a structure.²⁹ In this formulation, the structure-specific parameter α defines the importance of max dose and mean dose, ranging from 0 to 1. Constraints (2a–2f) are simply carried over from constraints (1a–1f) in Stage 1. However, a new constraint, (2g), is added to preserve the optimal plan utility from Stage 1. This constraint allows for relaxation of the Stage-1 plan utility value to allow sufficient search space for the Stage-2 optimization to reduce OAR doses. The relaxation parameter, η , determines the allowable plan utility degradation between Stage 1 and Stage 2. Although a value of 1 is ideal (maintaining the optimal utility found in Stage 1), in practice a value slightly less than 1 may be required to provide the optimization algorithm sufficient search space to improve dose metrics in Stage 2.

2.B. Technical Implementation

Optimization first requires calculation of an influence matrix, A , which represents the dose contribution of a discretized beam to a set of voxels, or points, within a patient volume. To facilitate this process, we built a plugin using a research version of the Eclipse Scripting Application Program Interface (ESAPI v15.5) (Varian Medical System Inc., Palo Alto, CA) to allow integration of our process with our commercial Eclipse treatment planning system (TPS). Integration with the TPS provided access to patient images, structures, beam configuration, and dose calculation as needed throughout the planning process. Point clouds were generated using pseudorandom point sampling

of pre-defined structures within a patient's body with adjustable structure-specific point spacing. Multi-field IMRT plans were configured using pre-defined beam angle templates that were initially fit to the target structure(s) and then expanded to create a beam that could be divided into uniform beamlet sizes. To calculate the influence matrix, we used an ESAPI method that separates a beam into beamlets based on a user-defined size parameter and calculates a full-scatter influence matrix using the Analytical Anisotropic Algorithm (AAA v15.5.11) (Varian Medical System Inc., Palo Alto, CA). To reduce the overall size and complexity (i.e., number of variables and constraints) of the optimization problems, this clinically accurate influence matrix calculation was separated into primary and secondary dose contributions. Secondary dose contributions, consisting of low dose scatter components, were removed from the influence matrix and summed for each point separately across each beam. These contributions were then incorporated into the optimization dose estimation by multiplying the total low dose scatter component by the corresponding average beam fluence. The decision variables in the optimization problem were beamlet fluence intensities, x , that were forced to be non-negative with constraints (1a) and (2a).

Maximization of the utility function (f_1) requires that the efficacy model is concave relative to dose, and the toxicity models are convex relative to dose, within the relevant prediction range. Although predictive models of efficacy and toxicity could take on many non-linear forms, they are typically represented by sigmoidal functions of dose (e.g., NTCP and TCP). Methods exist for transforming these types of partially convex and concave functions into purely convex or concave functions for the purposes of optimization and are useful in determining Pareto efficient solutions.³⁰⁻³² However, these transforms obscure the relative importance of absolute changes in efficacy and toxicity in a weighted-sum approach and do not preserve the interpretability of the utility function and weights (ϑ 's). Therefore, PWL relaxations of these models were implemented for the Stage-1 objective function and for constraint (2g) in Stage 2, reducing the problems to linear programs which can be efficiently solved with off-the-shelf commercial optimization solvers. PWL functions are routinely used to approximate nonlinear optimization problems with more computationally manageable linear ones while preserving continuity and concavity.³³ For the utility function f_1 , PWL approximations extended across the convex or concave envelope of the dose response functions with extensions beyond the envelope acceptable within the maximum allowable error of 0.5% (absolute). To create the linear segments within the PWL approximations, we sampled breakpoints on the original nonlinear dose-response functions starting with 10 segments per function. If the initial approximation was found to have error exceeding the allowable threshold at any point within the approximation, additional segments were added until the error was below the threshold. Additionally, since many predictive models are based on biological dose, rather than physical dose, we implemented the ability for models to be based on equivalent dose in 2 Gy fractions (EQD₂) for OARs using PWL approximations (starting with 4 segments), and linear scaling for target volumes with approximation errors less than 1 Gy (EQD₂) over the relevant dose range. To eliminate approximation errors during plan evaluation, plan utility and individual components of the utility function were recalculated after optimization using the original functions. The resulting linear optimization problems were solved using a third-party commercial optimization solver Gurobi v9.0.3 (Gurobi Optimization, LLC, Beaverton, OR) through the .NET interface on a workstation with dual Intel Xeon E5-2620 v4 8-core processors and 64 GB of memory. Interior point, or barrier, methods

have recently demonstrated good performance across linear and non-linear optimization problems in radiation therapy, including prioritized approaches.^{34,35} Therefore, for our linear problems, we tested both barrier and simplex algorithms implemented in Gurobi to determine which method was more suitable for this set of optimization problems.

Following optimization, the optimal fluence map can be directly imported into the TPS using ESAPI, allowing for leaf sequencing, final dose calculation, and plan visualization. Since our optimization process does not include a fluence smoothing stage or intermediate full dose calculation, dose discrepancies after leaf sequencing and final dose calculation are expected but are not the focus of this initial feasibility study. Therefore, quantitative plan comparisons for this study are limited to optimization solutions prior to these steps when possible.

2.C.1. Retrospective planning study

For this initial demonstration of the proposed optimization strategy, five previously treated NSCLC patients were selected for retrospective treatment planning and plan comparison. These patients were selected to be representative of the variation in patient geometry, target size, and model inputs from a cohort of 125 stage II-III NSCLC patients previously treated on institution review board (IRB) approved prospective studies. Table II provides details for the five patients. All patients were treated with definitive standard or dose-escalated 3DCRT with or without sequential or concurrent chemotherapy. These studies included data collection previously used to generate predictive models of local regional progression-free survival at 2 years (LRPFS_{2y}),³⁶ grade ≥ 3 cardiac event within 2 years (CE_{3+,2y}),³⁷ grade ≥ 3 radiation esophagitis (RE₃₊),³⁸ grade ≥ 3 radiation-induced lung toxicity (RILT₃₊)³⁹ based on patient demographics, clinical factors, and biomarkers. The RE₃₊ and RILT₃₊ models were based on toxicities reported and graded during regularly scheduled follow-up evaluations for up to 2 years following RT.^{38,39} Although the referenced LRPFS_{2y} and CE_{3+,2y} models were developed using longer term retrospective review of patient data, for this study we limited these models to time endpoints of 2 years following RT to match the timeframe for the RE₃₊ and RILT₃₊ models, which is also similar to the reported median overall survival time for this patient cohort.^{36,37} We used these previously published models, further described in Table III, to generate personalized predictive dose response curves, shown in Fig. 2. These patient-specific predictions are utilized in Stage 1 of the PUO process for optimization and calculation of overall plan utility, as well as constraint (2g) of Stage 2. For this initial study, θ values were uniformly set to 1, representing equal trade-off between efficacy and each toxicity in the utility function. To avoid use of a non-convex DVC in the optimization process, max dose was used in place of D_{2cc} for the esophagus model, leading to slight over-estimation of RE₃₊ during optimization (typically $< 0.5\%$, absolute). Although a CVaR⁺ metric could alternatively be used to estimate D_{2cc} , it also represents an approximation and results in similar over-estimation error. Therefore, to reduce computational complexity, max dose was used during optimization. However, final utility values for plan evaluation were calculated using D_{2cc} for the RE₃₊ prediction.

IMRT plans using the PUO approach were generated for the five patients. These 30-fraction plans consisted of 8 non-opposing 6MV photon treatment fields placed at standardized beam angles based on target laterality. Beam angle templates were determined by a dosimetrist with clinical experience

in IMRT planning for lung cases and were reviewed for appropriateness based on the planning geometry in these five cases. For the influence matrix calculation, beamlets were set to a size of 0.5 x 0.5 cm except for patient 5 which required 1.0 x 1.0 cm beamlets due to the large planning target volume (PTV) volume and computational limitations. Volume dose was calculated at a resolution of 2.5 x 2.5 x 3.0 mm³ and the influence matrix was sampled with 2 mm to 6 mm isotropic point spacing depending on the volume and importance of the structure.

Structure constraints used in the optimization are shown in Table IV. Mean and max OAR dose constraints were based on our institution's standard-of-care dose objectives for conventionally fractionated lung RT. The PTV mean EQD₂ limits (50–100 Gy, $\alpha/\beta = 10$) were selected to reflect the range of prescription doses from the original patient cohort corresponding to the development of the response models. For this cohort treated on IRB approved prospective studies, the median EQD₂ prescription dose was 70 Gy (range: 47–96 Gy, $\alpha/\beta = 10$).³⁷ Heterogeneity in the PTV dose was constrained to be similar to the heterogeneity considered clinically acceptable at our institution. CVaR⁺ constraints were added as convex representations of the VaR-based DVCs used clinically at our institution, including lung-GTV $V_{20\text{Gy}} < 35\%$, heart $V_{30\text{Gy}} < 50\%$, and heart $V_{50\text{Gy}} < 25\%$. These constraints were tuned based on a linear fit of CVaR⁺ to the corresponding VaR metric based on a retrospective analysis of IMRT plans from a separate cohort of 30 previously treated NSCLC patients. Additionally, a CVaR⁺ constraint for the esophagus was added as a conservative representation of $D_{2\text{cc}} < 68$ Gy. VaR dose metrics were checked after plan optimization to ensure that plans remained within the original VaR clinical limits. ξ EU weighting factors were set to 0.5 for all structures with equal weighting between structures for the Stage-2 objective. Based on empirical testing, η was set to 0.999 (0.1% overall utility loss acceptable), constraining Stage 2 to retain nearly the same overall plan utility as Stage 1 while still providing the optimization algorithm sufficient search space.

2.C.2 Planning Comparison

The proposed PUO method was compared to the clinically delivered 3DCRT plans, and retrospectively generated dose-only optimization (DOO) IMRT and volumetric modulated arc therapy (VMAT) plans. DOO plans were generated to represent plans with the same beam angles and hard constraints as the PUO plans but optimized with only a typical dose-metric-based objective. These plans were created using a similar workflow as the PUO plans, but only used Stage 2 of the PUO process and excluded the utility constraint (2g). An additional constraint for minimum target dose of 60 Gy (30 fractions) was added to make plans that followed current clinical planning prescription guidelines. VMAT plans with a 60 Gy (30 fractions) target dose goal were generated in Eclipse by a dosimetrist with prior experience in clinical lung treatment planning. These VMAT plans were used as a modern, clinically deliverable benchmark for the DOO IMRT plans generated with our external optimization process and fixed beam angles. To provide a uniform comparison between planning methods, VMAT plans were normalized to the $D_{95\%}$ of the corresponding DOO plans (normalization values < 102%). 3DCRT, VMAT, DOO, and PUO plans were compared by calculating the final plan utility. Dose-volume histograms (DVHs) of DOO and PUO plans directly from the

optimal solutions were also compared. Non-tumor integral dose (NTID) was calculated for the DOO and PUO plans using the methods outlined by D'Souza and Rosen assuming uniform normal tissue density and voxel size.⁴⁰ Variable prescription doses in PUO plans have the potential to influence NTID, so differences should be considered in the context of potential dose escalation and de-escalation relative to DOO plans. To provide a comparison of intermediate dose spill, R50% (defined as the ratio of 50% prescription isodose volume to PTV) was calculated for the DOO and PUO plans. Additionally, results from Stage 1 of the optimization process without clinical dosimetric OAR hard constraints (1c–1f), labeled $PUO_{s1,mod}$, were compared to the final PUO utility results to determine what effect, if any, the inclusion of OAR hard constraints had on the maximum achievable plan utility. $PUO_{s1,mod}$ plans are expected to show similar or increased plan utility compared to the unmodified PUO algorithm because they ignore acceptable OAR dose limits. Utility results are compared in terms of absolute utility difference. Therefore, given θ values of 1, any change in plan utility directly corresponds to a combination of changes in the predicted absolute probabilities of efficacy and toxicity. For example, if the predicted probability of efficacy increases by 5% (absolute, e.g., 55% to 60%) with predicted toxicity remaining the same, then the overall utility improvement would be 0.05. Alternatively, if the same increase in efficacy occurs, but one of the predicted toxicity probabilities increases by 7% (absolute), then the overall utility change would be $0.05 - 0.07 = -0.02$.

3. Results

For all five NSCLC cases, our approach successfully generated optimal beamlet weights that maximize utility while remaining within dose-based constraints. Total (Stage 1 and Stage 2) solution times ranged from 16 to 71 minutes, with the number of beamlets and points in the optimization problems ranging from 1144 to 2610 and 41888 to 96523, respectively, requiring 0.6–1.9 GB when encoded in protocol buffers format. Separation of the primary and secondary dose contributions within the influence matrix reduced the number of non-zero matrix elements by 83–92%, which greatly reduced solver time while still providing an accurate estimates of plan dose and predicted outcomes. On average, 78% of total solver time was spent on Stage 2, of which 67% was spent in the crossover solver phase to produce a basic solution. For our linear problems, Gurobi's barrier method was found to be faster than primal and dual simplex methods but required modification of the numerical focus to the maximum setting and decreasing the barrier convergence tolerance to reduce crossover solver time.

Plan utility values and target $D_{95\%}$ values for the various planning methods are shown in Table II with a further breakdown of the utility metrics shown in Fig 3. The PUO method resulted in average absolute utility improvements of 0.21 (range: 0.09–0.43), 0.17 (0.05–0.31), and 0.16 (0.07–0.30) when compared to 3DCRT, VMAT, and DOO plans, respectively. DOO IMRT plans were shown to have similar utility to VMAT plans, demonstrating that the DOO plans provide a clinically reasonable comparison for PUO plans. Analysis of individual components of the utility function showed that

when compared to DOO, PUO improved absolute predicted LRPFS_{2y} by, on average, 16.6% (range: 2.7–32.6%). Improvements in predicted efficacy were met with smaller changes in predicted grade 3–5 toxicities, with average changes of 0.1% (-0.5–0.5%), -0.9% (-4.0–1.7%), and 1.0% (-0.3–3.5%) for predicted CE_{3+,2y}, RE₃₊, and RILT₃₊, respectively. PUO resulted in increased target dose, based on analysis of PTV D_{95%}, for three of the five patients, with one patient's PUO plan retaining a 60 Gy D_{95%} and another's decreasing target dose slightly. NTID was on average 11.6% (-4.4–44%) higher in the PUO plans compared to the DOO plans, with the largest increases occurring for patient 1 (44%) and patient 2 (42%), corresponding to increases in prescription doses in the PUO plans of 38.3% and 42.2%, respectively, relative to the DOO plans. NTID and R50% metrics were found to be similar or better in the PUO plans. Comparison of the PUO_{S1,mod} utility metrics to the final PUO plan metrics showed that the incorporation of clinically relevant hard constraints reduced the potential maximum plan utility by, on average, 0.03 (range: 0.00–0.14). Patient 3 had the largest decrease, resulting mainly from the constraints on maximum esophagus and cord doses.

For patients 1 and 2, PUO led to large utility improvements over DOO. These improvements were mainly driven by increases in predicted LRPFS_{2y}, as visualized in Fig. 3, with the predicted probability of RILT₃₊ moderately increasing and probabilities of RE₃₊ and RP₃₊ remaining similar. The PUO plan for patient 3 had a moderate increase in utility compared to the DOO plan caused by increasing predicted LRPFS_{2y} with similar overall predicted toxicity risk. For patients 4 and 5, slight utility improvements in PUO plans were driven by improved predicted LRPFS_{2y} and decreased probabilities of RE₃₊ and RILT₃₊. For Patient 5, PUO planning also slightly decreased the predicted probability of CE_{3+,2y}.

Absolute and relative DVHs in Figs. 4 and 5 show comparisons of the DOO and PUO plans for patients 2 and 4, respectively. For patient 2, a large increase in target dose was noted with non-proportional increases in esophagus and cord doses. For patient 4, the PUO method slightly improved target coverage and decreased mean lung dose and maximum esophagus dose. Mean cardiac dose remained similar between the two methods, but the PUO plan increased maximum cord dose by 24.8 Gy. Variable increases in maximum cord dose were present in all PUO plans, since a cord toxicity risk model was not included in the Stage-1 optimization model. However, all maximum cord doses remained below 45 Gy as enforced by the hard constraint.

4. Discussion

The current clinical standard for RT planning, employing objective functions based upon standardized dose metrics, does not adequately account for patient variability in the optimization process. In many cases, these objectives are based on historical rates of control or toxicity in a patient population without incorporating personalized factors related to the radiosensitivity of tumors and OARs. Additionally, most cost functions require user-assigned weights for the dose penalties, relying on the expertise of planners and physicians to subjectively determine the appropriate weights for a patient. Through this method, allowable planning trade-offs, based on treatment outcomes rather than dose metrics, could be inconsistently decided upon since they are

not directly prioritized or optimized. The proposed method presented in this study aims to augment this standard approach by facilitating a direct exploration of planning trade-offs based on patient-specific predictive models of efficacy and toxicity. Through the implementation of a utility-based objective, we believe our strategy offers a more intuitive, tunable approach to balance a patient's potential therapeutic benefit and risk. The proposed method integrates the ability for non-uniform weighting of predicted toxicities in an interpretable manner with direct relation to the predicted probabilities of treatment outcomes. Therefore, toxicity weights (ϑ 's) can be elicited directly by determining the relative harm associated with a given toxicity compared to a given efficacy metric. These values could be assigned independently by a clinician prior to plan creation, potentially incorporating patient input, or based on expert consensus opinion determined through a method similar to Hobbs et al.⁴¹

We demonstrated the feasibility of this method in a small cohort of NSCLC patients, which is a population of interest given the evident trade-offs between local control and lung, heart, and esophageal toxicities in RT planning. Additionally, failure of the Radiation Therapy Oncology Group (RTOG) 0617 phase III randomized trial to demonstrate improved outcomes at higher doses suggests that utilization of dose escalation in the treatment of NSCLC may only be beneficial for a subset of patients.^{23,37,42,43} Although 60 Gy in 2 Gy fractions has become the standard-of-care for patients with stage III NSCLC, individualized dose escalation remains of interest to improve the relatively poor outcomes for these patients.⁴⁴⁻⁴⁶ This includes isotoxic approaches to individualize dose escalation in which escalation is based on pre-specified, uniform normal tissue dose limits without consideration of underlying heterogeneity in the patient population.^{45,46} Our method provides a quantitative approach for determining which patients may benefit from dose escalation or redistribution based on patient-specific clinical factors and biomarkers while also accounting for patient geometry and OAR dose limits.

While this study illustrates the potential for the proposed planning strategy, it does have several limitations. Although we compared plan utility metrics between multiple planning methods, DVH comparison was limited to the DOO and PUO plans prior to leaf sequencing and final volumetric dose calculation due to dose-discrepancies between optimal plans from the optimizer and plans imported, sequenced, and calculated inside of our TPS. However, dose colorwash comparisons following these steps, noting possible discrepancies, for DOO and PUO are available in a supplemental document for Patient 2 (Figure S-1) and Patient 4 (Figure S-2). These discrepancies have been previously noted in other optimization methods and are normally addressed through updating the optimization process with intermediate full dose calculations.⁴⁷ However, in our hierarchical approach, this update process could invalidate hard constraints. Future work will look at addressing this discrepancy through the implementation of a fluence smoothing optimization stage, heterogeneity constraints, or a correction step similar to that presented in Zarepisheh et al.⁴⁸ Additionally, we found Gurobi's barrier solver to perform well for this optimization problem and did not test other methods or algorithms beyond the simplex method included in Gurobi. It is possible that first-order algorithms could improve performance; however, the barrier solver performance could likely be improved through additional tuning of the algorithm's parameters. Furthermore, in a clinical setting, the crossover solver phase (used to convert an interior-point solution to a basic

solution) can be eliminated to improve solve times by more than 50% while causing very little impact to the optimal fluence map.

Our study did not evaluate the effect of beam angle selection on the optimization process or achievable plan utility. Further investigation will be required to determine this impact of this and if utility could be improved through beam angle optimization. Alternatively, arc-based IMRT delivery methods, such as VMAT, could be used to alleviate the need for beam angle selection. However, VMAT optimization methods are currently inhibited by non-convexity and associated local minima, and therefore require additional development to facilitate implementation of constrained hierarchical optimization strategies.⁴⁹

The PUO approach facilitates the use of personalized dose-response models in the optimization process, but aggregated response models could also be utilized. In testing this method for NSCLC patients, we used previously developed models of response based upon patients treated with 3DCRT. Ideally, modeling would be based on a cohort that was treated with the same delivery technique as that being optimized since dose response could be, in part, correlated to the underlying dose distribution. This distribution, particularly within OARs, could systematically differ between varying treatment delivery techniques. While our approach optimizes a patient's tumor dose, rather than assigning a fixed dose goal prior to optimization, modeling for this aspect is challenging because it requires previously treated patients to receive a range of prescription doses to determine the relationship between treatment efficacy and dose. In the era of increasingly uniform prescription doses, such as the 60 Gy standard-of-care for NSCLC patients, modeling dose-based efficacy may not be fully possible. In the absence of an efficacy model, the PUO strategy could still be utilized to reduce total patient-specific toxicity burden at a fixed prescription dose.

Our initial results from this method demonstrate that plan utility for NSCLC patients has the potential to be improved, although drawing additional conclusions from the results is limited by the small sample size. For these cases, increased plan utility resulted in similar or increased doses to structures that were unrepresented in the utility function. Increases to unrepresented structures are anticipated in this prioritized approach, but all doses remained within currently acceptable limits. The impact of these hard constraints on achievable plan utility are variable but can be large depending on plan geometry as noted in the comparison of the PUO_{S1,mod} and unmodified PUO plans for patient 3. If the increases in doses are determined to be clinically unacceptable, stricter constraints and penalties can be implemented, but this must be balanced with the impact on achievable utility. Alternatively, a more beneficial approach would be to directly capture clinically acceptable tradeoffs by adding additional components representing these structures to the utility objective. By implementing this strategy, the PUO's two-stage approach allows for the progressive transition to more holistic biologically based treatment planning without completely deviating from physical dose-based planning for structures with poorly understood or unmodeled dose response.

In this study, we introduced and demonstrated the feasibility of PUO, but did not investigate the clinical impact or acceptability of planning tradeoffs, including the potentially large variability in prescription doses with fixed fractionation, captured using PUO compared to DOO. *In silico* and prospective clinical trials will be required to determine the clinical efficacy and effectiveness of PUO.

Additionally, we did not investigate the effects of potential error and uncertainty in the dose-response models or the selection of objective function parameters (e.g., ϑ 's) on the optimization process and resulting plans. Future studies will aim to address this through application of the PUO strategy in additional disease sites and on larger cohorts of patients.

5. Conclusions

We developed and studied an inverse IMRT planning strategy where patient-specific radiosensitivities of tumors and normal tissues are directly factored into the optimization objective. Through this approach, we aim to improve a patient's overall RT outcome by balancing potential therapeutic benefit with the associated risk in an interpretable and tunable manner. First, a patient's overall plan utility, based upon personalized models of biological response, is maximized subject to relevant clinical constraints. Then, through a hierarchical optimization technique, a typical dose-based objective function is minimized while retaining maximal or near-maximal plan utility.

We demonstrated the feasibility of the approach using a cohort of NSCLC patients with previously developed predictive models of treatment efficacy and toxicity based on demographics, clinical factors, and biomarkers. The proposed planning method generated plans conforming to clinical constraints with improved overall utility when compared to plans generated using typical dose-based objectives.

Acknowledgments

This work was funded, in part, by NIH P01CA059827 and Varian Medical Systems, Inc. We thank Philip Boonstra PhD for his assistance in toxicity modeling and the Varian API team, including Wayne Keranen and Seppo Tuomaala, for supporting and providing changes to the Eclipse API to facilitate this work.

Disclosure of Conflicts of Interest

Matthew Schipper has previously consulted with Innovative Analytics outside of the submitted work. Shruti Jolly was part of the advisory boards of AstraZeneca and Varian Medical Systems and also was a consultant for the Michigan Radiation Oncology Quality Consortium (MROQC, funded by Blue Cross Blue Shield of Michigan) during this work. Markus Varsta is an employee of Varian Medical Systems. Martha Matuszak received research funding from Varian Medical Systems and serves as the co-director of MROQC. All other authors have no conflicts of interest to report.

References

Kachnic LA, Tsai HK, Coen JJ, et al. Dose-painted intensity-modulated radiation therapy for anal cancer: a multi-institutional report of acute toxicity and response to therapy. *Int J Radiat Oncol Biol Phys*. 2012;82(1):153-158. doi:10.1016/j.ijrobp.2010.09.030

Nutting CM, Morden JP, Harrington KJ, et al. Parotid-sparing intensity modulated versus conventional radiotherapy in head and neck cancer (PARSPORT): a phase 3 multicentre randomised controlled trial. *Lancet Oncol*. 2011;12(2):127-136. doi:10.1016/S1470-2045(10)70290-4

Klopp AH, Yeung AR, Deshmukh S, et al. Patient-Reported Toxicity During Pelvic Intensity-Modulated Radiation Therapy: NRG Oncology-RTOG 1203 [published correction appears in *J Clin Oncol*. 2019 Mar 20;37(9):761] [published correction appears in *J Clin Oncol*. 2020 Apr 1;38(10):1118]. *J Clin Oncol*. 2018;36(24):2538-2544. doi:10.1200/JCO.2017.77.4273

Rosenthal DI, Chambers MS, Fuller CD, et al. Beam path toxicities to non-target structures during intensity-modulated radiation therapy for head and neck cancer. *Int J Radiat Oncol Biol Phys*. 2008;72(3):747-755. doi:10.1016/j.ijrobp.2008.01.012

Kim JH, Jenrow KA, Brown SL. Mechanisms of radiation-induced normal tissue toxicity and implications for future clinical trials. *Radiat Oncol J*. 2014;32(3):103-115. doi:10.3857/roj.2014.32.3.103

Kessler ML, Mcshan DL, Epelman MA, et al. Costlets: A Generalized Approach to Cost Functions for Automated Optimization of IMRT Treatment Plans. *Optim Eng* 6, 421–448 (2005). doi:10.1007/s11081-005-2066-2

Deasy JO, Alaly JR, Zakaryan K. Obstacles and advances in intensity-modulated radiation therapy treatment planning. *Front Radiat Ther Oncol*. 2007;40:42-58. doi:10.1159/000106027

Xing L, Li JG, Donaldson S, Le QT, Boyer AL. Optimization of importance factors in inverse planning. *Phys Med Biol*. 1999;44(10):2525-2536. doi:10.1088/0031-9155/44/10/311

Wu X, Zhu Y. An optimization method for importance factors and beam weights based on genetic algorithms for radiotherapy treatment planning. *Phys Med Biol*. 2001;46(4):1085-1099. doi:10.1088/0031-9155/46/4/313

Bedford JL, Webb S. Elimination of importance factors for clinically accurate selection of beam orientations, beam weights and wedge angles in conformal radiation therapy. *Med Phys*. 2003;30(7):1788-1804. doi:10.1118/1.1582471

Lee T, Hammad M, Chan TC, Craig T, Sharpe MB. Predicting objective function weights from patient anatomy in prostate IMRT treatment planning. *Med Phys*. 2013;40(12):121706. doi:10.1118/1.4828841.

Chan T, Craig T, Lee T, Sharpe M. Generalized Inverse Multiobjective Optimization with Application to Cancer Therapy. *Operations Research*, 2014;62(3), 680-695. doi:10.1287/opre.2014.1267

Lee T, Hammad M, Chan TC, Craig T, Sharpe MB. Predicting objective function weights from patient anatomy in prostate IMRT treatment planning. *Med Phys*. 2013;40(12):121706.

doi:10.1118/1.4828841

Boutilier JJ, Lee T, Craig T, Sharpe MB, Chan TC. Models for predicting objective function weights in prostate cancer IMRT. *Med Phys*. 2015;42(4):1586-1595. doi:10.1118/1.4914140

Babier A, Boutilier JJ, Sharpe MB, McNiven AL, Chan TCY. Inverse optimization of objective function weights for treatment planning using clinical dose-volume histograms. *Phys Med Biol*.

2018;63(10):105004. Published 2018 May 10. doi:10.1088/1361-6560/aabd14

Shen C, Nguyen D, Chen L, et al. Operating a treatment planning system using a deep-reinforcement learning-based virtual treatment planner for prostate cancer intensity-modulated radiation therapy treatment planning. *Med Phys*. 2020;47(6):2329-2336. doi:10.1002/mp.14114

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. *Med Phys*.

2012;39(3):1386-1409. doi:10.1118/1.3685447

Feng Z, Tao C, Zhu J, et al. An integrated strategy of biological and physical constraints in biological optimization for cervical carcinoma. *Radiat Oncol*. 2017;12(1):64. Published 2017 Apr 4.

doi:10.1186/s13014-017-0784-1

Kan MW, Leung LH, Yu PK. The use of biologically related model (Eclipse) for the intensity-modulated radiation therapy planning of nasopharyngeal carcinomas. *PLoS One*. 2014;9(11):e112229. Published

2014 Nov 5. doi:10.1371/journal.pone.0112229

Modiri A, Stick LB, Rice SR, et al. Individualized estimates of overall survival in radiation therapy plan optimization - A concept study. *Med Phys*. 2018;45(11):5332-5342. doi:10.1002/mp.13211

Stick LB, Vogelius IR, Modiri A, et al. Inverse radiotherapy planning based on bioeffect modelling for locally advanced left-sided breast cancer. *Radiother Oncol*. 2019;136:9-14.

doi:10.1016/j.radonc.2019.03.018

Rechner LA, Modiri A, Stick LB, et al. Biological optimization for mediastinal lymphoma radiotherapy - a preliminary study. *Acta Oncol*. 2020;59(8):879-887. doi:10.1080/0284186X.2020.1733654

Scott JG, Sedor G, Scarborough JA, et al. Personalizing Radiotherapy Prescription Dose Using Genomic Markers of Radiosensitivity and Normal Tissue Toxicity in NSCLC. *J Thorac Oncol*.

2021;16(3):428-438. doi:10.1016/j.jtho.2020.11.008

Schipper MJ, Taylor JM, TenHaken R, Matuzak MM, Kong FM, Lawrence TS. Personalized dose selection in radiation therapy using statistical models for toxicity and efficacy with dose and biomarkers as covariates. *Stat Med*. 2014;33(30):5330-5339. doi:10.1002/sim.6285

Li P, Taylor JMG, Kong S, Jolly S, Schipper MJ. A utility approach to individualized optimal dose selection using biomarkers. *Biom J*. 2020;62(2):386-397. doi:10.1002/bimj.201900030

- Romeijn HE, Ahuja RK, Dempsey JF, Kumar A, Li JG. A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Phys Med Biol*. 2003;48(21):3521-3542. doi:10.1088/0031-9155/48/21/005
- Kishimoto S, Yamashita M. A successive LP approach with C-VAR type constraints for IMRT optimization. *Operations Research for Health Care*. 2018;17, 55-64. doi:10.1016/j.orhc.2017.09.007
- Wu VW, Epelman MA, Pasupathy KS, Sir MY, Deufel CL. A new optimization algorithm for HDR brachytherapy that improves DVH-based planning: Truncated Conditional Value-at-Risk (TCVaR). *Biomed. Phys. Eng. Express*. 2020;6. doi:10.1088/2057-1976/abb4bc
- Thieke C, Bortfeld T, Küfer KH. Characterization of dose distributions through the max and mean dose concept. *Acta Oncol*. 2002;41(2):158-161. doi:10.1080/028418602753669535
- Romeijn HE, Dempsey JF, Li JG. A unifying framework for multi-criteria fluence map optimization models. *Phys Med Biol*. 2004;49(10):1991-2013. doi:10.1088/0031-9155/49/10/011
- Hoffmann AL, den Hertog D, Siem AY, Kaanders JH, Huizenga H. Convex reformulation of biologically-based multi-criteria intensity-modulated radiation therapy optimization including fractionation effects. *Phys Med Biol*. 2008;53(22):6345-6362. doi:10.1088/0031-9155/53/22/006
- Kierkels RG, Korevaar EW, Steenbakkens RJ, et al. Direct use of multivariable normal tissue complication probability models in treatment plan optimisation for individualised head and neck cancer radiotherapy produces clinically acceptable treatment plans. *Radiother Oncol*. 2014;112(3):430-436. doi:10.1016/j.radonc.2014.08.020
- Kontogiorgis S. Practical piecewise-linear approximation for monotropic optimization. *INFORMS Journal on Computing*. 2000;12(4):324-340. doi:10.1287/ijoc.12.4.324.11877
- Mukherjee S, Hong L, Deasy JO, Zarepisheh M. Integrating soft and hard dose-volume constraints into hierarchical constrained IMRT optimization. *Med Phys*. 2020;47(2):414-421. doi:10.1002/mp.13908
- Wu VW, Epelman MA, Wang H, et al. Optimizing global liver function in radiation therapy treatment planning. *Phys Med Biol*. 2016;61(17):6465-6484. doi:10.1088/0031-9155/61/17/6465
- Sun Y, Hawkins PG, Bi N, et al. Serum MicroRNA Signature Predicts Response to High-Dose Radiation Therapy in Locally Advanced Non-Small Cell Lung Cancer. *Int J Radiat Oncol Biol Phys*. 2018;100(1):107-114. doi:10.1016/j.ijrobp.2017.08.039
- Dess RT, Sun Y, Matuszak MM, et al. Cardiac Events After Radiation Therapy: Combined Analysis of Prospective Multicenter Trials for Locally Advanced Non-Small-Cell Lung Cancer. *J Clin Oncol*. 2017;35(13):1395-1402. doi:10.1200/JCO.2016.71.6142
- Soni PD, Boonstra PS, Schipper MJ, et al. Lower Incidence of Esophagitis in the Elderly Undergoing Definitive Radiation Therapy for Lung Cancer. *J Thorac Oncol*. 2017;12(3):539-546. doi:10.1016/j.jtho.2016.11.2227

Hawkins PG, Boonstra PS, Hobson ST, et al. Radiation-induced lung toxicity in non-small-cell lung cancer: Understanding the interactions of clinical factors and cytokines with the dose-toxicity relationship. *Radiother Oncol*. 2017;125(1):66-72. doi:10.1016/j.radonc.2017.09.005

D'Souza WD, Rosen II. Nontumor integral dose variation in conventional radiotherapy treatment planning. *Med Phys*. 2003;30(8):2065-2071. doi:10.1118/1.1591991

Hobbs BP, Thall PF, Lin SH. Bayesian Group Sequential Clinical Trial Design using Total Toxicity Burden and Progression-Free Survival. *J R Stat Soc Ser C Appl Stat*. 2016;65(2):273-297. doi:10.1111/rssc.12117

Bradley JD, Paulus R, Komaki R, et al. Standard-dose versus high-dose conformal radiotherapy with concurrent and consolidation carboplatin plus paclitaxel with or without cetuximab for patients with stage IIIA or IIIB non-small-cell lung cancer (RTOG 0617): a randomised, two-by-two factorial phase 3 study. *Lancet Oncol*. 2015;16(2):187-199. doi:10.1016/S1470-2045(14)71207-0

Bradley JD, Hu C, Komaki RR, et al. Long-Term Results of NRG Oncology RTOG 0617: Standard-Versus High-Dose Chemoradiotherapy With or Without Cetuximab for Unresectable Stage III Non-Small-Cell Lung Cancer. *J Clin Oncol*. 2020;38(7):706-714. doi:10.1200/JCO.19.01162

Thureau S, Dubray B, Modzelewski R, et al. FDG and FMISO PET-guided dose escalation with intensity-modulated radiotherapy in lung cancer. *Radiat Oncol*. 2018;13(1):208. Published 2018 Oct 23. doi:10.1186/s13014-018-1147-2

Haslett K, Franks K, Hanna GG, et al. Protocol for the isotoxic intensity modulated radiotherapy (IMRT) in stage III non-small cell lung cancer (NSCLC): a feasibility study [published correction appears in *BMJ Open*. 2016 Jul 18;6(7):e010457corr1]. *BMJ Open*. 2016;6(4):e010457. Published 2016 Apr 15. doi:10.1136/bmjopen-2015-010457

Haslett K, Bayman N, Franks K, et al. Isotoxic Intensity Modulated Radiation Therapy in Stage III Non-Small Cell Lung Cancer: A Feasibility Study. *Int J Radiat Oncol Biol Phys*. 2021;109(5):1341-1348. doi:10.1016/j.ijrobp.2020.11.040

Siebers JV, Lauterbach M, Tong S, Wu Q, Mohan R. Reducing dose calculation time for accurate iterative IMRT planning. *Med Phys*. 2002;29(2):231-237. doi:10.1118/1.1446112

Zarepisheh M, Hong L, Zhou Y, et al. Automated intensity modulated treatment planning: The expedited constrained hierarchical optimization (ECHO) system. *Med Phys*. 2019;46(7):2944-2954. doi:10.1002/mp.13572

Unkelbach J, Bortfeld T, Craft D, et al. Optimization approaches to volumetric modulated arc therapy planning. *Med Phys*. 2015;42(3):1367-1377. doi:10.1118/1.4908224

Figure Legends

Figure 1. Mathematical formulation of the optimization problems.

Figure 2. Predictive efficacy and toxicity models used to estimate overall plan utility: Local regional progression-free survival at 2 years (LRPFS_{2y}, top left), Grade ≥ 3 cardiac event within 2 years (CE_{3+,2y}, top right), Grade ≥ 3 radiation esophagitis (RE₃₊, bottom left), Grade ≥ 3 radiation-induced lung toxicity (RILT₃₊, bottom right). All cases included in the generation of the models are plotted with gray solid lines while the five cases used for plan comparison are plotted with colored dashed lines. White background areas represent the model ranges used in study, whereas the gray background represents ranges excluded by hard constraints during optimization. The CE₃₊ model only stratifies into two distinct groups since it is based on one binary covariate, baseline cardiac disease, with heart dose.

Figure 3. Utility results for the original clinically delivered 3DCRT plans and retrospective VMAT, Dose-only optimization (DOO), and Prioritized utility optimization (PUO) replans for the five patients used for plan comparison. Results include overall plan utility (left), Local regional progression-free survival at 2 years (LRPFS_{2y}, middle-top), Grade ≥ 3 cardiac event within 2 years (CE_{3+,2y}, right-top), Grade ≥ 3 radiation esophagitis (RE₃₊, middle-bottom), Grade ≥ 3 radiation-induced lung toxicity (RILT₃₊, right bottom). As an example, Patient 1 estimated probabilities for LRPFS_{2y}, CE_{3+,2y}, RE₃₊, and RILT₃₊ were 53.6%, 4.3%, 0.4%, and 3.9%, for the DOO plan and 86.2%, 4.8%, 2.1%, 7.4% for the PUO plan. Therefore, the overall plan utility improvement using PUO is calculated as $[0.862 - (0.048 + 0.021 + 0.074)] - [0.536 - (0.043 + 0.004 + 0.039)] = 0.719 - 0.450 = 0.269$.

Figure 4. Absolute (top) and relative (bottom) DVH comparisons between the Dose-only optimization (DOO) plan (dashed lines) and the Prioritized utility optimization (PUO) plan (solid lines) for Patient 2.

Figure 5. Absolute (top) and relative (bottom) DVH comparisons between the Dose-only optimization (DOO) plan (dashed lines) and the Prioritized utility optimization (PUO) plan (solid lines) for Patient 4.