Exploring the limits of LET-based target-RBE enhancement in proton therapy for lung cancer

Feasible gains and their therapeutic relevance

By

M.O.M. Koch

In partial fulfilment of the requirements for the degree of:

Master of Science
in Biomedical Engineering Track Medical Physics

at the Delft University of Technology,
to be defended publicly on Tuesday August 12, 2025 at 12:30 PM

Supervisors:	Prof. dr. M.S. Hoogeman	
	Dr. S.J.M. Habraken	
	Msc. J.I. De Jong	
Thesis committee:	Prof. dr. M.S. Hoogeman,	TU Delft
	Dr. ir. D. Lathouwers,	TU Delft
	Dr. S.J.M. Habraken	Holland PTC

Exploring the limits of LET-based target-RBE enhancement in proton therapy for lung cancer

Feasible gains and their therapeutic relevance

Bsc. M.O.M. Koch[a], Msc. J.I. De Jong[b], Dr. S.J.M. Habraken[c,d,b], Prof. Dr. M.S. Hoogeman[b,c]

[a]*Delft University of Technology, Department of Radiation Science and Technology, Delft, The Netherlands*
[b]*University Medical Center Rotterdam, Department of Radiotherapy, Rotterdam, The Netherlands*
[c]*Holland PTC, Delft, The Netherlands*
[d]*Leiden University Medical Center, Department of Radiotherapy, Leiden, The Netherlands*

Abstract

Background and purpose: The variable RBE of protons could be strategically harnessed by confining high RBE to the tumor by optimizing the LET distribution, which has been shown to strongly correlate with RBE and does not contain the uncertainties associated with variable RBE models. However, increasing target LET compromises target dose robustness against setup and range uncertainties, as LET rises sharply past the Bragg peak, where the dose gradient is steep. Previous studies lack thorough robustness analyses and rely on LET-based objectives in optimization, which are not available to all institutions. Lung cancer, in particular, remains underexplored in the context of target-LET optimization. This study therefore (i) investigates the elevation of target LET without relying on LET-based optimization objectives, (ii) determines the maximum feasible increase while maintaining adequate target dose robustness and adherence to clinical goals, and (iii) assesses whether such increases translate into meaningful improvements in tumor control probability in a cohort of recently treated lung cancer patients.
Materials and methods: A LET-painting technique using two sets of opposing beams with beam-specific minimum dose objectives for proximal target segments was employed to elevate target LET. To find the near-feasible target LET, high inhomogeneity was leveraged in combination with CTV-based planning, under the assumption of motion-mitigation. Probabilistic evaluation using Polynomial Chaos Expansion (PCE) was then performed to ensure precise adherence to adequate robustness. LET-optimized treatment plans, employing this recipe, were then compared to dose-optimized treatment plans for 10 lung cancer patients, based on dose-averaged LET (LET_d), $LET_d \times D$, Unkelbach RBE, and EUD-based TCP values.
Results: An increase of 50% in target LET_d was observed for the probabilistically robust LET-optimized treatment plans compared to the dose-optimized plans, raising the RBE-weighted dose by 6 Gy and mean target RBE from 1.1 to 1.16. These gains were associated with increased mean target dose, reduced target dose robustness and increased beam-entry doses. The elevated LET_d significantly improved tumor control probability, although a large part of the total gain in TCP was due to increased physical dose.
Conclusion: Clinically significant increases in TCP due to enhanced target LET_d are possible without the use of LET_d-based objectives, by leveraging high inhomogeneity and precisely adhering to adequate target dose robustness. It should be determined whether these trade-offs in addition to changes in OAR doses are clinically acceptable and should be directly compared to dose escalation strategies alone. Future technologies might help enhance benefits or improve trade-offs between inhomogeneity, robustness and target LET, making target RBE-enhancement through LET to improve treatment efficacy more clinically viable.

Keywords: Proton therapy, LET-optimization, RBE, Probabilistic evaluation, Polynomial Chaos Expansion, Lung cancer

1. Introduction

Proton therapy offers distinct advantages over conventional photon-based radiotherapy, due to the unique physical properties of protons. As charged particles, they interact continuously with tissue, gradually losing energy in the process, further increasing the likelihood of interaction. This process results in a sharp deposition of most of their energy at the end of their range, in what is known as the Bragg peak. This characteristic dose deposition results in high target-dose conformity and reduced radiation burden to organs at risk (OARs). In contrast, photons deposit a high entrance dose, followed by a more gradual dose deposition that can extend far beyond the target, ultimately irradiating a higher portion of healthy tissue.

In addition to these dosimetric advantages, protons also result in greater biological impact compared to photons, described by the Relative Biological Effectiveness (RBE). This increased RBE stems from a more dense spatial pattern of ionization produced by protons, which results in more complex DNA damage and therefore a higher likelihood of cell death [1]. As protons slow down, the density of ionizations increases further. This ionization density is closely related to the varying Linear Energy Transfer (LET), defined as the amount of energy deposited per unit distance traveled by the particle. Initially the LET of protons rises gradually, but increases sharply at the distal edge of the Bragg peak, where protons stop. This increase in LET leads

to a corresponding gradual increase in RBE, followed by a peak in RBE in the same region.

Despite this known variation in RBE due to varying LET, a constant RBE-value of 1.1, determined as a conservative estimate of in vivo studies in the 1970s, is still the recommended standard in clinical proton therapy today [2]. This RBE value is deliberately conservative with respect to target dose, to minimize the risk of RBE overestimation and subsequent underdosing of the tumor [2]. The continued use of a constant RBE reflects the significant uncertainties that are associated with variable RBE models. These models not only depend on LET, but also on dose, fractionation and biological factors such as cell-type radiosensitivities, which are especially difficult to predict reliably [2]. In addition, many RBE models differ in their underlying assumptions and use different datasets for fitting, leading to different predictions based on the choice of RBE model [3] [4].

The use of a constant RBE of 1.1 has recently led to much discussion. While the conservative value prevents target underdosing, it presents a risk for OARs, as this value may underestimate the biological effect in critical structures bordering the target distally, where RBE is elevated due to increased LET. Although current evidence for an increase in radiation-induced side-effects due to variable RBE remains mixed, significant correlations have been found [5]. For example, changes observed in healthy brain tissue on MR images have been correlated with elevated RBE and an increased incidence of rib-fractures has been associated with increased end-of-range RBE [6] [7].

While focus has been placed on this risk of increased RBE in OARs, increased RBE could also be strategically harnessed. By confining high RBE inside the tumor, therapeutic efficacy may be improved, while potentially shifting high RBE away from OARs. Alternatively, by enhancing RBE within the target, a sufficient therapeutic effect may be reached even in cases where dose escalation is not possible due to the proximity of critical structures, offering a potential strategy for overcoming anatomical constraints.

A promising strategy for modulating RBE is through optimization of the LET distribution instead, since it has been shown to correlate strongly with RBE and, being a physical quantity, it can be more reliably calculated, and does not contain the uncertainties associated with RBE models [5].

The potential for manipulation of the LET distribution stems from the inherent degeneracy of the proton therapy optimization problem, especially present in Intensity Modulated Proton Therapy (IMPT) which uses Pencil Beam Scanning (PBS) for delivery [8]. This degeneracy describes the fact that similar dose distributions can be achieved with very different sets of pencil beams and weights, resulting in very different LET distributions. This means that, in principle, one should be able to obtain a similar quality dose distribution with a more favorable LET distribution [9].

While LET optimization is a promising strategy for elevating target RBE, it also comes at a cost. This is because LET starts to rise sharply in the distal-fall off region of a spread-out Bragg peak, meaning to increase target LET, the dose part with steep gradients would have to move closer to, or inside the target. Consequently, target dose robustness with regards to setup and range uncertainties would inherently suffer, as a small displacement of the tumor will quickly result in a much lower target dose.

While there seems to exist an inherent trade-off between target LET and target dose robustness, it has been reported that allowing target overdosing (i.e. decreasing the uniformity of target dose) can improve this trade-off, by either increasing target dose robustness or target LET [10][11][12]. This relationship between target dose uniformity, target LET and target dose robustness has been described as the so-called LET 'trilemma', where only two of these three qualities can be sufficiently achieved in a treatment plan [13].

Quite a few target-LET optimization studies have been performed that indicate varying amounts of target LET elevation, from a few percent in mean LET [14] to more than one hundred percent when employing many beam angles [15]. However, few IMPT studies include robustness analyses [11][16], with only a single study performing such analyses for a representative group size [17]. This in spite of the known trade-off between target LET and robustness. This could mean that the large increases in LET that are reported could come at the cost of clinically unacceptable reductions in target dose robustness. Conversely, in some cases there might exist more potential for target LET elevation, as a higher reduction in target dose robustness could be allowed. Therefore it remains to be assessed what the current maximum potential is for target LET elevation in proton therapy while adhering to clinically acceptable target dose robustness. It is also important to assess whether the possible increases in target LET constitute a clinically relevant effect, reflected by significant increases in Tumor Control Probability (TCP).

Furthermore, many studies aimed to elevate target LET by directly incorporating LET-based objectives for the target into the optimization problem [5], a functionality that not every institution currently has access to. Therefore it is important to examine whether standard dose-optimization functionalities offered by any treatment planning software can be sufficient in shaping the LET-distribution and create clinically acceptable target RBE-enhancement plans.

Finally, the treatment sites studied are often limited to those of brain, prostate, and head and neck (H&N), with only a single case examined for a lung treatment [18]. Despite lung cancer being the most common form of cancer and the greatest contributor to cancer-deaths worldwide [19].

In this study, we therefore (i) investigate the elevation of target LET without relying on LET-based optimization objectives, (ii) determine the maximum feasible increase while maintaining adequate target dose robustness and adherence to clinical goals, and (iii) assess whether such increases translate into meaningful improvements in tumor control probability in a cohort of lung cancer patients treated with IMPT.

2. Methods and materials

2.1. Patient Data

A total of ten stage III-IV non-small cell lung cancer (NSCLC) patients with various histologies and primary tumor locations, recently treated at Holland PTC, were included in this study.

These patients were referred for proton treatment according to a model-based approach currently adopted by Dutch health care authorities [20]. Fractionation schemes varied from 55, 60 to 66 Gy(RBE) in 20, 30 and 24 fractions respectively, using a constant RBE of 1.1. Patient and treatment characteristics are included in Table S1. Plans in which robust target coverage had to be compromised due to OAR-based clinical goals were excluded, as these already represent complex cases for which additional LET optimization would impose further planning challenges.

2.2. Treatment planning system

All treatment plans in this study were created using RayStation (version 2023B RaySearch, Sweden). Dose calculations were performed using the latest RayStation Monte Carlo (MC) dose engine [21] [22] with 4000 ions per spot and 1% uncertainty. Dose grid resolution was set to 0.3 cm.

This engine also allowed scoring dose-averaged LET (LET_d) values, used to quantify LET in this study. The LET_d value in a voxel (i) is computed by summing the LET values of the particle tracks (j) crossing that voxel, weighted by their microscopic dose (d) delivered to that voxel, see Formula 1 [23].

$$LET_{d,i} = \frac{\sum_j d_j \cdot LET_j}{\sum_i d_j} \tag{1}$$

2.3. Clinical treatment plans (Clinical)

Clinical treatment plans (*Clinical*) consisted of Multi Field Optimized (MFO) IMPT plans with three to five coplanar beams. Beam angles were chosen based on protocol and prior experience with additional constraints such as a minimum separation angle between beams of 30 degrees, and avoidance of traversal through the arms, breast sides, table sides, thorax support and contralateral lung.

To mitigate the interplay effect caused by breathing motion [24], range shifters with variable thicknesses (3 or 5 cm) were used, depending on the extent of tumor motion.

To account for breathing-induced tumor motion, dose was prescribed to an Internal Target Volume (ITV). The ITV is defined as the geometric union of Clinical Target Volume (CTV) contours delineated across multiple respiratory phases of a 4D CT scan, rigidly transferred to the average CT, which was used for further dose planning.

Robust minimax optimization was performed to account for setup and range errors, considering a total of 21 error scenarios [25]. These scenarios included seven combinations of setup errors, consisting of the nominal scenario and six combinations of setup errors corresponding to displacements along the three principal directions of ± 7 mm. These setup errors were combined with relative range errors of -3%, 0%, and 3%. Robust min- and max-dose objective functions were used for the ITV. An overview of the clinical goals and initial planning objectives template, including objectives for the OARs, are included in the Supplementary Materials in Tables S2 and S3 respectively.

According to consensus within the Dutch Proton Therapy (DUPROTON) group [26], clinical goals were evaluated based on 4D robust evaluation of the voxel-wise minimum (VWmin) and voxel-wise maximum (VWmax) dose distributions of 28 error scenarios. The VWmin and VWmax distributions were generated from the combined average dose of four breathing phases, transferred by deformable registration to the mid-exhale phase, on which evaluation was performed. The error scenarios, computed separately for each phase, included 14 setup error combinations with directions determined by the six faces and eight vertices of a cube, all normalized to 7 mm. These setup errors were combined with range errors of -3% and 3%. The near-minimum VWmin dose in the CTV had to reach 95% of the prescribed dose ($VWmin_{D_{98},CTV} \geq 0.95 D_{pr}$). The near maximum VWmax dose in the CTV had to remain below 107% of the prescribed dose ($VWmax_{D_2,CTV} < 1.07 D_{pr}$), as per ICRU recommendation [27].

2.4. LET-optimization strategy

2.4.1. LET-painting method

In this study, LET was shaped and confined to the target by employing two sets of opposing beam angles, combined with beam-specific minimum dose objectives for proximal target segments [28]. These segments were created by subdividing the target volume based on proximity to each beam direction, using the center of the primary tumor as a reference point. The primary tumor was set as the target for LET-optimization throughout this study, to ensure high LET was confined to the tumor only. Further details on the automated segmentation process based on beam-direction proximity are provided in Supplementary Materials section S3.

Clinical beam angles served as starting point for the selection of opposing beams. Two sets of opposing beams were selected for each patient, where beam angles were initially copied from the clinical plan, ensuring that adding an opposing beam was feasible given the clinical goals, especially pertaining to sparing of OARs. If the resulting beam angles would lead to violation of OAR-related clinical goals, they were slightly adjusted until all clinical goals could be reached.

The use of opposing beams was motivated by the idea that the required minimum dose in the center of the target can be supplied by the combined distal dose fall-off regions of both beams where LET rises sharply, ”painting” the center of the target with high LET. However, as beam spots are typically placed distal to the target as a consequence of robust optimization, beam-specific minimum dose objectives for proximal target segments were introduced, to encourage the optimizer to place spots closer to or within the target.

2.4.2. Leveraging treatment parameters according to the LET-trilemma

In addition to the beam-specific min dose objectives, a high near-maximum (D_2) dose objective of up to 140% of the prescribed dose was allowed to allow for high target dose inhomogeneity (i.e. low uniformity). According to the LET-trilemma, this should allow for higher increases in target LET_d. While this degree of inhomogeneity is high, such values are not uncommon in stereotactic body radiotherapy (SBRT) for lung cancer [29][30].

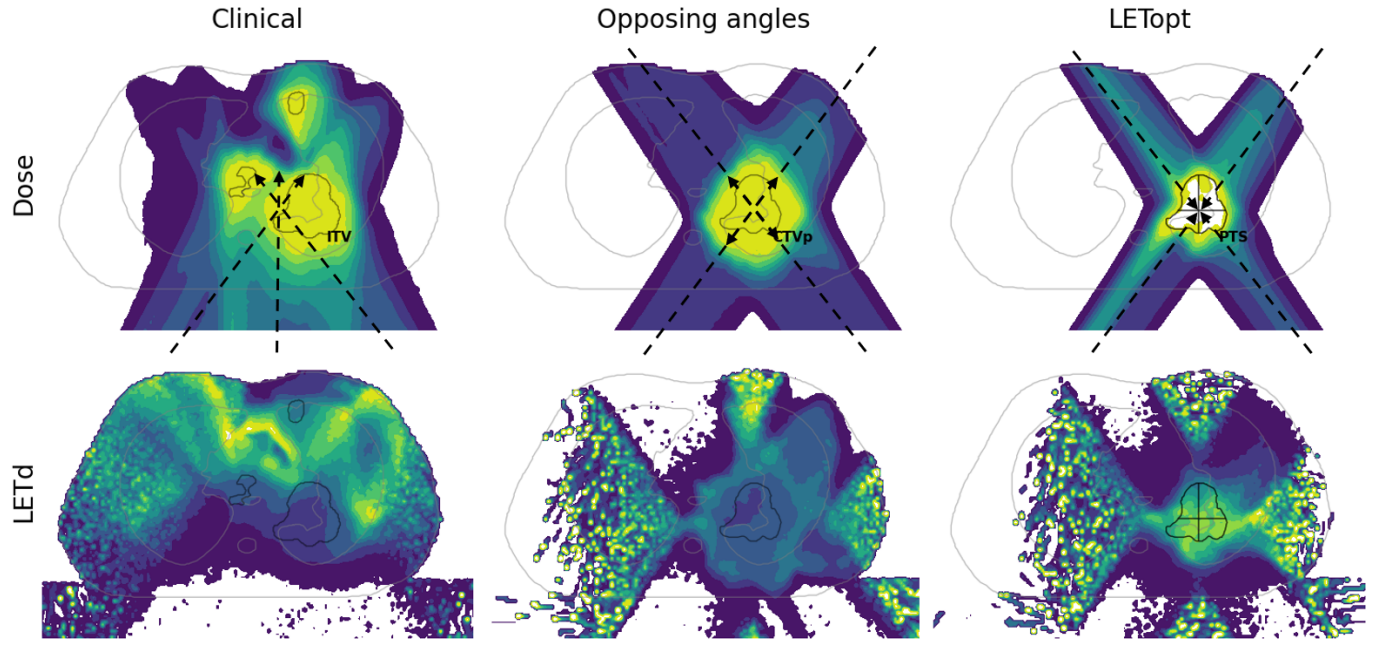

Figure 1: A visualization of the employed target-LET painting method, which deviates from the clinical plan (first column) by using two sets of opposing beams (second column) and subdividing the primary CTV (CTV_p) into beam-specific proximal target segments (PTS) with associated beam-specific minimum dose objectives and allowing a high maximum dose (third column). The arrows indicate the employed beam angles. The tip of the arrows indicate the approximate placements of spots. As indicated by the arrow heads; the introduction of beam-specific minimum dose goals causes spots to be placed further inside the target, concentrating high LET_d inside the target (LETopt case), rather than distally (clinical case) or surrounding the target (middle case).

Furthermore, CTV-based planning of the primary tumor (CTV_p) rather than ITV-based planning was performed, by assuming breathing motion can be neglected through the use techniques that eliminate its effect on the resulting dose distribution, such as breath-hold, respiratory gating or the use of ultra-high dose rates, as used in FLASH radiotherapy. This assumption not only simplifies the planning process, but may also allow for higher target LET_d values to be achieved, as this margin reduction is likely to result in less robust plans. Figure 1 illustrates the described LET-painting method in combination with CTV-based planning and high target dose inhomogeneity to maximize LET in the target.

Finally, to fully leverage optimization freedom in terms of target dose robustness, probabilistic evaluation was used to obtain a more statistically robust assessment of target dose robustness. As this allows very specific adherence to a predefined probabilistic robustness goal, this results in the maximum reduction in target dose robustness in favor of target LET_d. Instead of prescribing dose to the VWmin-D_{98}, as is the current clinical standard, probabilistic planning with Polynomial Chaos Expansion (PCE) was applied with the probabilistic goal of achieving a coverage of 0.95 D_{pr} in 99.8 percent of the volume in 90 percent of the error scenarios ($p_{10}(D_{99.8,CTV}) = 0.95D_{pr}$). A study employing PCE robust evaluation found this to be most consistent on average with photon planning and the associated formulation of Van Herk's margin recipe [31]. The probabilistic maximum dose goal was $p_{95}(D_{2,CTV_p}) \leq 1.4D_{pr}$, reflecting the increased allowance for the maximum dose.

2.4.3. *Polynomial Chaos expansion*

Polynomial Chaos Expansion was used to perform fast probabilistic evaluation, which can be used to fit a plan-specific analytical model of the dose engine for different uncertainty scenarios $\vec{\xi}$, consisting of the combination of different realizations of the three setup errors $\sigma_x, \sigma_y, \sigma_z$ and relative range error ρ ($\vec{\xi} = \left(\sigma_x, \sigma_y, \sigma_z, \rho\right)$). The dose D_i in a voxel is approximated by a finite series expansion of probabilists' Hermite polynomial basis vectors Ψ_k, orthonormal for normally distributed variables, which the uncertainty variables are assumed to be, with expansion coefficients $c_{i,k}$, see Formula 2.

$$D_i \approx \sum_{k}^{P} c_{i,k}\Psi_k(\vec{\xi}) \tag{2}$$

The PCE model is then constructed by determining the expansion coefficients through linear regression based on Monte Carlo dose computations for a sparse set of uncertainty scenarios, where the amount of polynomials and scenarios used for fitting the model is optimized to balance between speed and accuracy [32]. For further technical details, the reader is referred to the Supplementary Materials of the original paper that introduced PCE in the field of proton therapy [32].

2.4.4. *Probabilistic evaluation*

PCE parameter optimization and initial validation was performed for five lung cancer patients, see Supplementary Materials section S4. To perform probabilistic evaluation, a PCE model was constructed using the open-source OpenGPC [32][33], with a maximum Polynomial order of five, a grid order of four and an

extra level of one. The expansion coefficients of a PCE model were determined based on a sparse set of 208 error scenarios, computed by the RayStation MC-dose engine.

To generate these error scenarios, systemic and random setup errors were estimated using the clinical margin $M = 7$ mm and Van Herk's margin recipe ($M = 2.5\Sigma + 0.7\sigma$). Assuming both error types were approximately equal ($\Sigma \approx \sigma$) resulted in an estimated error of 2.1875 mm for each, assumed to be equal for each principal direction. A mean of the relative range error of 1.2% was used, based on a previous lung phantom study [34], with a standard deviation of 1%.

Using the constructed PCE model, dose distributions of 10.000 fully fractionated treatments for each patient were simulated by sampling a fixed isotropic systematic setup error (Σ) and a fixed relative range error (ρ) for each simulated treatment, and a different random setup error (σ) for each individual fraction. In combination with the amount of fractions (N), a total of 10.000 $\times N$ error scenarios were sampled and computed per treatment plan.

2.5. Probabilistically robust LET-optimized treatment plans (RobustLETopt)

Using the outlined LET-optimization strategy, probabilistically robust LET-optimized treatment plans were created. In addition to the beam-specific minimum dose goals for proximal target segments, standard robust min- and max-dose objectives for the full unsegmented CTV_p were used. In general, a more simplified set of planning objectives was used that uses simple dose-fall off functions for OARs, enabling more generalizable dose planning and minimizing inter-plan variability due to planning strategy. The full set of planning objectives used for the *RobustLETopt* plans can be found in Table S4. As the effect of breathing motion was neglected no range shifters were employed.

Robust minimax optimization was performed with settings initially equal to clinical robust optimization ($\sigma = 7$ mm, $\rho = 3\%$). A PCE model was constructed and the $p_{10}(D_{99.8})$ was evaluated based on the simulated 10.000 fully fractionated treatments. To adhere to the coverage goal of $p_{10}(D_{99.8,CTV}) = 0.95D_{pr}$, the dose was scaled to exactly meet this goal. If scaling lead to violation of the probabilistic maximum dose goal by more than one percentage point ($p_{95}(D_{2,CTV}) > 1.41D_{pr}$), the setup error used during robust optimization was increased. Consequently, this increased target dose robustness and therefore the $p_{10}(D_{99.8,CTV})$, resulting in a lower $p_{95}(D_{2,CTV})$ after scaling. Conversely, whenever the $p_{95}(D_{2,CTV})$ was more than one percentage point lower than 140% D_{pr} ($p_{95}(D_{2,CTV}) < 1.39D_{pr}$), the setup error used during robust optimization was lowered. This process was iteratively performed until treatment plans could be successfully scaled, after which the needed scaling factor to achieve the probabilistic coverage goal was used to scale the Monitor Units (MUs) of the four opposing beams equally in the treatment plans in RayStation, to be exported again for PCE for final evaluation.

By aiming to remain within one percentage point of the maximum dose goal of 140 % D_{pr}, treatment plans could be created that fully exploit the optimization freedom in terms of target dose robustness and maximum dose. The target LET_d in these plans therefore represent the maximum values that are feasible.

The creation of the *RobustLETopt*, consisting of this iterative process, is visualized in Figure 2.

2.5.1. Non-robust LET-optimized treatment plans (LETopt)

Additional non-robust LET-optimized plans were created that use the same LET-optimization strategy and planning objectives as used in the *RobustLETopt* plans, except for robust minimax optimization and robustness evaluation. These plans therefore allow to isolate the effectiveness of the LET-painting method in shaping target LET_d. Additionally, they represent maximum target LET_d enhancement achievable with this LET-painting method in the case of no setup or range error, serving as a theoretical benchmark in terms of LET_d elevation. The planning objectives for the *LETopt* plans can be found in table S5.

2.5.2. Dose-optimized treatment plans (RobustDoseOpt)

The LET-optimized treatment plans (*RobustLETopt*, *LETopt*) use different beam angles and target the CTV of the primary tumor instead of the full ITV, excluding possible irradiation of secondary tumors or lymph nodes. As a result, changes in OAR doses due to the employed LET-optimization strategy cannot be compared in a fair manner. Therefore, dose-optimized treatment plans were created that use the same beam angles and CTV_p-based planning as the LET-optimized treatment plans, but other than that follow clinical treatment planning guidelines, including VWmin and VWmax evaluation and the use of range shifters. The planning objectives for the *RobustDoseOpt* plans can be found in Table S6.

2.6. Evaluation and statistical analyses

To judge the success of the outlined LET-optimization strategy and investigate what maximum LET-values are possible, the patient-wise median of the mean, near-minimum ($LETd_{98}$) and near-maximum ($LETd_2$) LET_d values in the CTV_p were compared between the four plan types (*Clinical*, *RobustDoseOpt*, *LETopt*, *RobustLETopt*).

In addition, the patient-wise median of the mean voxel-wise product of LET_d and Dose ($LETd \times D$), as well as the near-minimum ($[LETd \times D]_{98}$) and near-maximum values ($[LETd \times D]_2$), were further evaluated. This quantity can be seen as the additional biological dose due to LET_d, when assuming that LET increases RBE linearly with dose. This assumption underlies the pragmatic Unkelbach RBE model, see Formula 3 [35]. While not a formal RBE model, it can be used to evaluate the potential biological impact of increased LET, without requiring biological input variables such as the difficult-to-predict cell-type radiosensitivity parameters (α,β).

$$RBE_{Unkelbach} = (1 + cLET_d) \tag{3}$$

Multiplying Formula 3 by the physical dose D_p gives the RBE-weighted dose D_{bio}, see Formula 4. This formula includes the LET_d-dose product term, reflecting the added biological dose

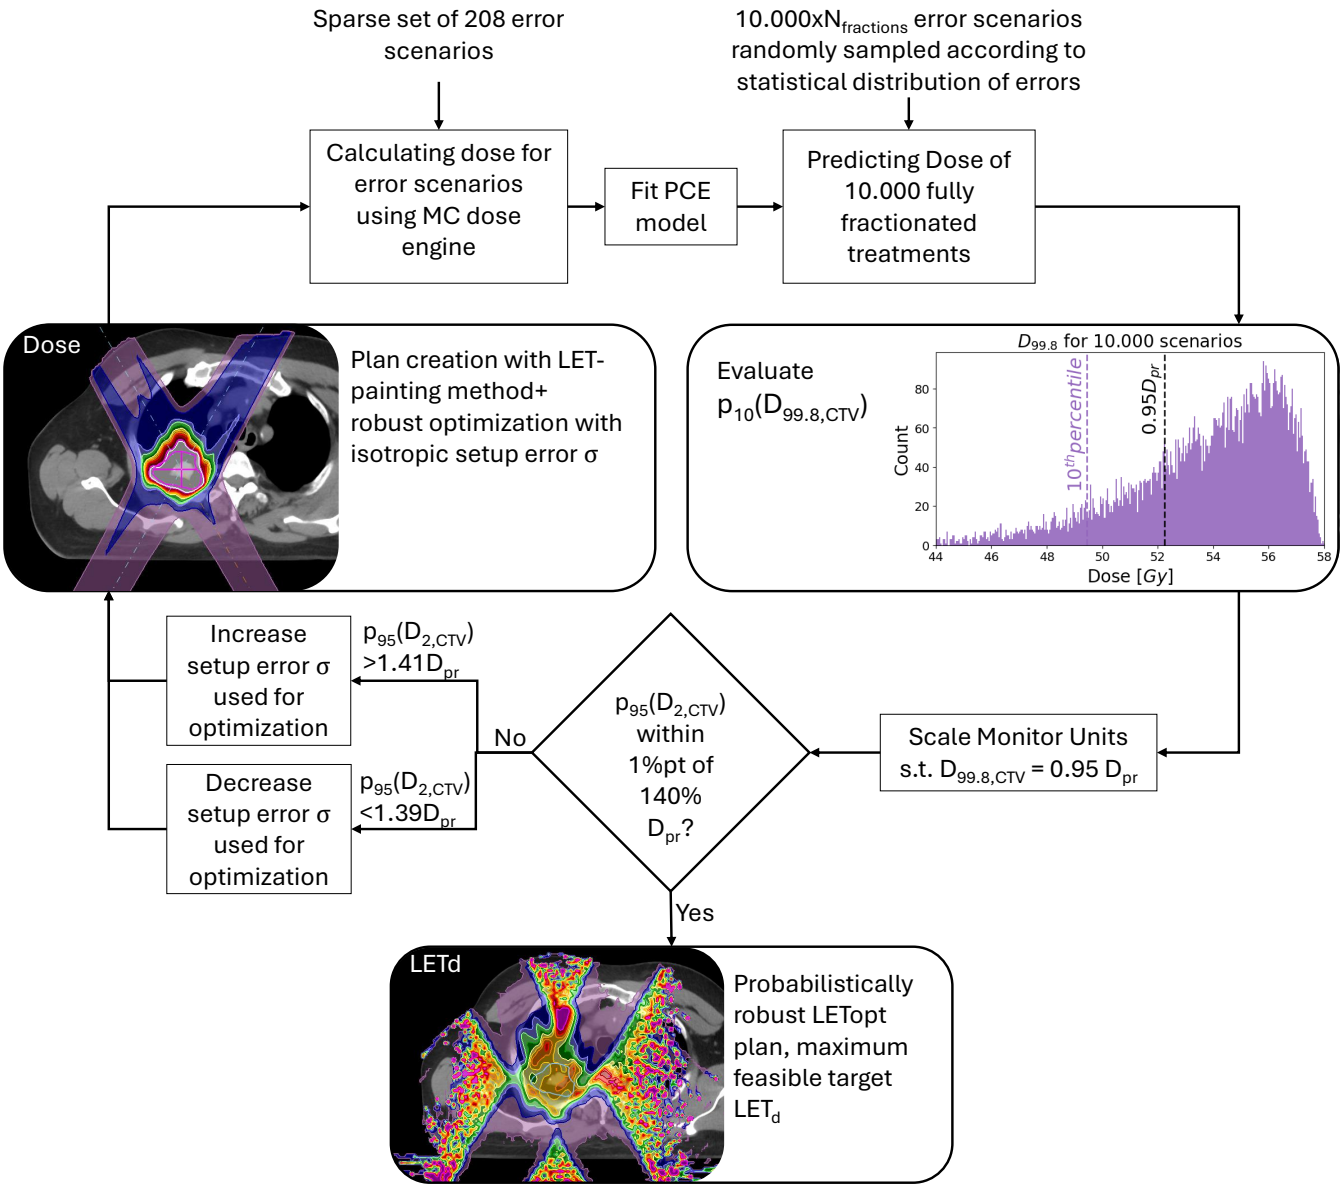

Figure 2: A flow diagram visualizing the process of creating the probabilistically robust LET-optimized treatment plans using PCE.

due to LET_d. The proportionality constant c allows to convert this product into units of Gray. Its value was taken to be 0.04 $\mu m/keV$ throughout this study, resulting in an RBE of 1.1 in the center of a spread-out Bragg peak of 5 cm modulation and 10 cm range [36].

$$D_{bio} = D_p + cD_pLET_d \tag{4}$$

As the final goal of target LET escalation is to achieve an increased target RBE, the patient-wise median of the mean, near-minimum and near-maximum LET_d values for all patients were converted into corresponding RBE values using Formula 3.

Several additional metrics were evaluated to assess the attainment of clinical goals. To this end, nominal target coverage and maximum target dose goals were evaluated based on the D_{98} and D_2 respectively. To provide further information on the dose to the target, the D_{mean} was evaluated as well. Regarding target dose robustness, the $p_{10}(D_{99.8,CTV})$ and $p_{95}(D_{2,CTV})$ were evaluated and compared between treatment plans. The clinical goal values for the OARs were evaluated based on nominal values and probabilistic values. These probabilistic values were derived from the 95th percentile of associated clinical goal dose metric across the 10,000 simulated treatments.

2.6.1. Tumor control probability calculations

The clinical significance of target-LET_d increases in the LET-optimized plans was examined using tumor control probability (TCP) calculations. These calculations were performed based on a TCP-model that incorporates the equivalent uniform dose (EUD) concept, see Formula 5 [37].

$$TCP = \frac{1}{1 + \left(\frac{TCD_{50}}{EUD}\right)^{4\gamma_{50}}} \tag{5}$$

The values for the radio-biological response parameters, including the dose yielding a 50% chance of tumor control (TCD_{50}) and the corresponding slope (γ_{50}), were taken from Okunief et al. (TCD_{50} = 51.24 Gy , γ_{50} = 0.83), which are based on multi-institutional data of NSCLC patients [38]. A value of $a = -10$ was used for EUD calculations [38], see Formula 6.

$$EUD = \left(\sum_i v_i D_i^a\right)^{1/a} \tag{6}$$

With v the fractional volume of the target in a voxel i and D the corresponding dose in that volume. For EUD calculations, the total dose of the combined fractions was used as input.

The effect of LET_d on the resulting TCP was determined by calculating the EUD, used as input for TCP calculation, based on the unweighted physical dose alone ($RBE = 1$) and the Unkelbach RBE-weighted dose (Formula 4). As the latter includes the increased biological dose due to LET_d, the difference in the resulting TCP values reflects the added value of LET_d in terms of therapeutic effect.

In addition, TCP values based on an EUD calculated using the constant-RBE-weighted dose ($RBE = 1.1$) were evaluated to assess whether the current consensus of RBE=1.1 remains accurate for predicting TCP in LET-optimized plans.

2.6.2. Statistical Analyses

Statistical tests from the Python Scipy.stats library were used to test for statistical significance. The paired samples t-test was performed to test for significant difference in mean metric values (e.g. mean LET_d) between plan types. The requirements of this test were performed beforehand, by performing the Shapiro-Wilk and Levene's test to test for normality of the distributions and equal variances respectively. Violation of the assumption of normality resulted in the use of the non-parametric Wilcoxon signed rank test.

3. Results

3.1. Plan creation and LET-painting

The clinical and opposing beam angles, used for the *RobustLETopt*, *LETopt* and *RobustDoseOpt* plans as part of the LET-painting method, are listed in the Supplementary Materials in Table S7. The employed robustness settings and scaling factors used to create the *RobustLETopt* plans can be found in Table S8. The computation times associated with PCE model creation and subsequent dose prediction of 10.000 fully fractionated treatments can be found in Figure S5.

Figure S6 in the Supplementary Materials illustrates how the used LET-painting method leads to spots being placed more proximal to the target, thereby increasing target LET_d, by comparing the separate beam-dose and beam-LET_d distributions.

3.2. LET and RBE elevation

Figure 3 shows for the 10 patients the near-minimum, mean and near-maximum values per plan type for the LET_d, dose, LET_d×D, and RBE, in the CTV_p respectively.

3.2.1. LET_d elevation

Significant differences were found in the population mean of the mean target LET_d between the *RobustLETopt* and both the *Clinical* (p<0.001) and *RobustDoseOpt* (p<0.001) plans. The population median of the mean target LET_d increased from 2.54 (*Clinical*) and 2.62 (*RobustDoseOpt*) to 4.01 $keV/\mu m$. Similarly, the population-median near-minimum target LET_d increased from 2.15 and 2.33 to 3.12 $keV/\mu m$ and the population-median near-maximum target LET_d increased from 3.19 and 3.07 to 4.78 $keV/\mu m$.

The non-robust optimized *LETopt* plans show higher elevation in target LET_d compared to the *RobustLETopt* plans. Significant differences are found between the population-mean of the mean target LET_d in the *LETopt* plans compared to the *Clinical* (p<0.001) and *RobustDoseOpt* (p<0.001) plans. The population median of the mean target LET_d in the *LETopt* plans increased from 2.54 (*Clinical*) and 2.62 (*RobustDoseOpt*) to 5.30 $keV/\mu m$.

A smaller, but statistically significant difference in the population mean of mean target LET_d was found between the *Clinical* and *RobustDoseOpt* plans (p=0.039<0.05), with higher mean and near-minimum LET_d values for *RobustDoseOpt*. In contrast, lower near-maximum LET_d values are observed for the *RobustDoseOpt* plans.

3.2.2. Dose comparison

The comparison of dose metrics shows higher near-minimum, mean and near-maximum dose values for both the *LETopt* and *RobustLETopt* plans, with the highest values observed for the latter plans. Higher target dose inhomogeneity is observed for these plans as well, reflected by an increased distance between the near-minimum and near-maximum dose.

All plans meet the nominal target coverage goal of $D_{98} \geq 0.95D_{pr}$. The *Clinical* and *RobustDoseOpt* plans adhere to the maximum dose goal of $D_2 < 1.07D_{pr}$ in the nominal scenario. The *LETopt* and *RobustLETopt* plans adhere to the relaxed maximum dose goal of $D_2 < 1.4D_{pr}$ in the nominal scenario, except for a single *RobustLETopt* plan. Almost all treatment plans met clinical dose goals for OARs in both the nominal and the probabilistic scenarios, see Figures S8 and S9 respectively in the Supplementary Materials. Some deviations are found for the *Clinical* plans.

3.2.3. Increases in biological dose due to LET_d

Significant differences in the population mean of the mean LET_d×D are found between the *RobustLETopt* and the *Clinical* (p<0.001) and *RobustDoseOpt* (p<0.01) plans. Increases in the population median of the mean LET_d×D of 5.95 Gy and 5.85 Gy are observed in the *RobustLETopt plans*, compared to the *Clinical* and *RobustDoseOpt* plans respectively. Increases in population-median near-minimum LET_d×D of 3.72 Gy and 3.18 Gy, and increases in near-maximum LET_d×D of 7.30 Gy and 7.84 Gy are found.

Significant differences in population-mean of mean target LET_d×D are also found between the non-robust *LETopt* and *Clinical* and (p<0.001) *RobustDoseOpt* (p <0.001) treatment plans. The increases in population-median near-minimum LET_d×D, mean LET_d×D and near-maximum LET_d×D from the *Clinical* and *RobustDoseOpt* plans to the *LETopt* plans are respectively: 5.21 Gy and 4.67 Gy, 8.49 Gy and 8.39 Gy, and 10.73 Gy and 11.28 Gy.

The highest overall near-maximum LET_d×D value is observed for the *LETopt* plans and is slightly higher than 20 Gy (21.69 Gy). The highest overall LET_d×D value of the *RobustLETopt* plans is substantially higher than 15 Gy (17.71 Gy).

3.2.4. RBE enhancement

The population median of mean target LET_d values correspond to population-median mean RBE values of 1.102, 1.105, 1.212 and 1.160. These constitute increases in mean RBE of approximately 5.9 percentage points for the *RobustLETopt* plans when comparing to the *Clinical* plans and 5.6 percentage points when comparing to the reference *RobustDoseOpt* plans. These same increases are approximately 11 and 10.7 percentage points respectively for the *LETopt* plans. Population-median near-minimum RBE values of 1.086, 1.093, 1.153 and 1.125, and near-maximum RBE values of 1.128, 1.123, 1.253 and 1.19 are observed for the *Clinical*, *RobustDoseOpt*, *LETopt* and *RobustLETopt* plans respectively.

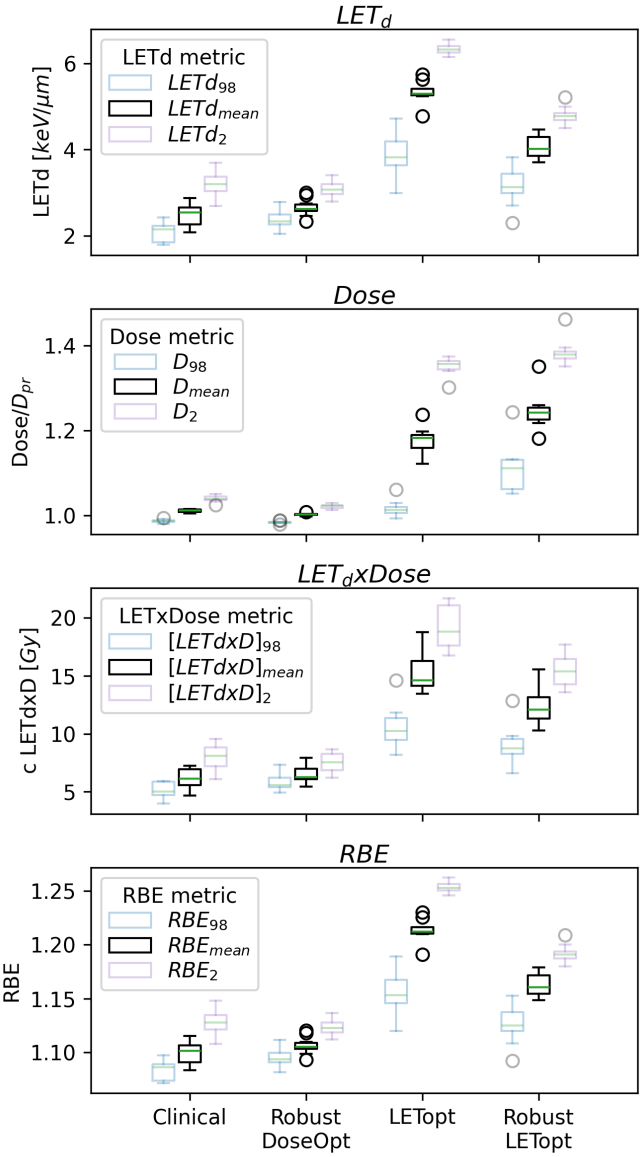

Figure 3: Group-wise comparison between the four plan types of near-minimum, mean and near-maximum values of dose, LET_d, LET_d×D and RBE inside the CTV_p for all patients.

3.3. Example case: 2D dose, LET_d and LET_d×D visualization

Figure 4 shows 2D dose, LET_d and LET_d×D contour plots for all plan types for a representative patient.

As can be seen in the figure, the employed LET-painting method shapes and confines high LET_d to the target, especially in the *LETopt* plans. The clinical constraints regarding beam angle selection in combination with lung geometry results in X-shaped beam arrangements.

Comparing the dose-distribution of the four plans, it can be observed how the additional opposing beam angles lead to an inevitable increased low-dose bath of healthy tissue, in particular the heart in this case. At the same time, a lower dose to healthy tissues and OARs can also be observed as a consequence of the reduced robustness of the LET-optimized plans, which increases the target dose conformity. An increased beam entry dose can

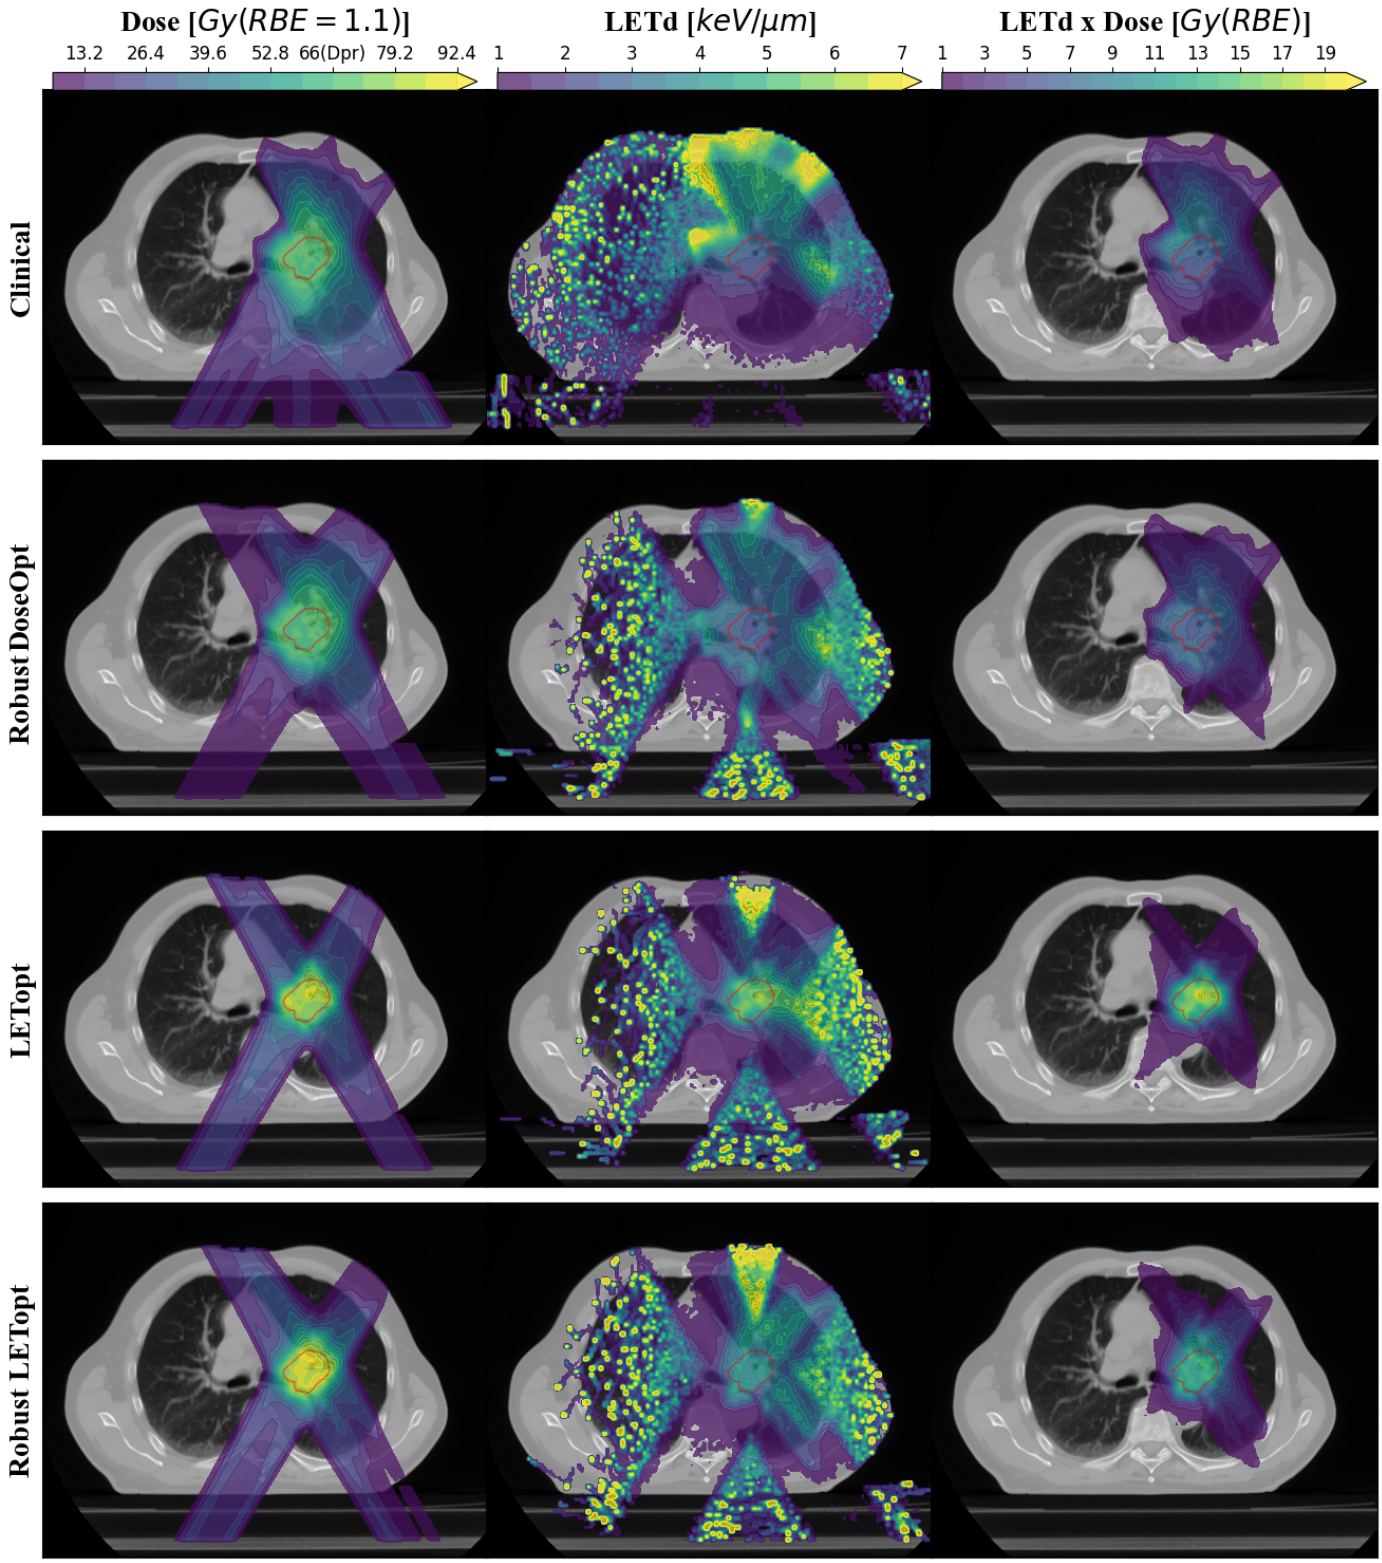

Figure 4: Contour plots of the dose, LET_d and c $LET_d \times D$ distributions of the four plan types for a single representative patient. c is taken to be 0.04 $\mu m/keV$ to convert the product of LET_d and Dose into units of Gray. The contours of the primary tumor CTV are indicated in red

also be observed along the opposing beams in the LET-optimized plans, especially visible for the *LETopt* plans. Finally, one can clearly observe the effect of a higher maximum dose objective for the LET-optimized plans, with a higher and less uniform dose in the target.

Comparing the LET_d distributions of the four plans, it can be observed for the *Clinical* plan how LET_d is highest distal to the target. In combination with non-negligible dose surrounding the target, this leads to an extended distal region of biological dose (LET_d×D) penetrating into healthy tissue of the heart in this case. In the *RobustDoseOpt* plans, hotspots of higher LET_d can be observed surrounding the target in nearly all directions, leading to a band of slightly elevated LET_d-based biological dose surrounding the target. In case of the *LETopt* plans, the high LET_d is mostly confined within the target region, but not completely, as it slightly extends around the target. As a result, a highly confined and elevated LET_d×D is observed in the target. In the *RobustLETopt* plans the target LET_d is less high and dispersed over a larger area compared to the *LETopt* plans, extending further outside the borders of the target. Nonetheless, higher LET_d compared to the conventional dose-optimized plans is observed inside the target. In combination with the elevated dose in the target, a considerable product of LET_d and dose is still observed in the target. One can also notice the small localized peaks in LET_d at the lateral borders of the employed beams, that manifest in voxels where few protons stop, automatically scoring the high LET of those few protons that stop there.

3.4. Comparison of probabilistic robustness

Figure 5 shows for the four plan types the probabilistic 10th percentile $D_{99.8}$ and 95th percentile D_2 of 10.000 simulated treatments for all patients.

Population-median $p_{10}(D_{99.8})$ values of 0.98, 0.98, 0.64 and 0.95 D_{pr}, are observed for the *Clinical*, *RobustDoseOpt*, *LETopt* and *RobustLETopt* treatment plans. Except for a single *Clinical* plan, all *Clinical* and *RobustDoseOpt* plans adhere to the probabilistic coverage goal of $p_{10}(D_{99.8,CTV}) = 0.95D_{pr}$. Most *RobustLETopt* plans adhere to this goal or are extremely close to this goal. No *LETopt* plans adhere to the probabilistic coverage goal. A comparison of the probabilistic $D_{99.8}$ distributions for the different plan types for a single patient is included in Supplementary Materials Figure S10.

Population-median $p_{95}(D_2)$ values of 1.035, 1.020, 1.507 and 1.406 D_{pr}, are observed for the *Clinical*, *RobustDoseOpt*, *LETopt* and *RobustLETopt* treatment plans. The dose-optimized treatment plans adhere to the near-maximum dose goal of 1.07 D_{pr}. Most *RobustLETopt* plans adhere to the relaxed near-maximum dose goal of 140% D_{pr} with a margin of 1%pt, with a single case exceeding this margin. No *LETopt* plan adhered to this goal.

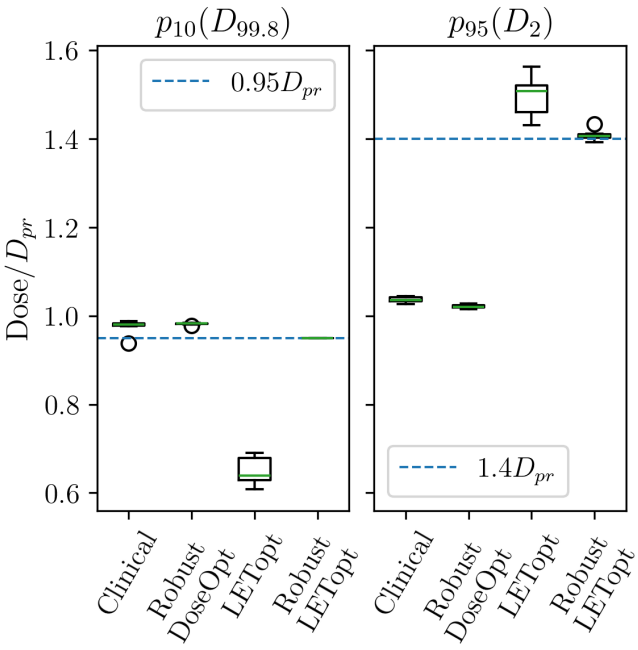

Figure 5: Comparison of the found probabilistic $p_{10}(D_{99.8})$ and $p_{95}(D_2)$ values of the four plan types for all patients

3.5. Comparison of TCP values

Figure 6 shows the comparison between the four plan types of the EUD-based TCP values, with EUDs calculated based on the non-RBE-weighted physical dose, Unkelbach RBE-weighted dose and constant-RBE-weighted (RBE=1.1) dose. For all plan types there was a significant difference between the physical dose-based TCP values and the Unkelbach RBE-weighted dose-based TCP-values, with increases in median values from 0.56 to 0.64 ($p<0.01$), 0.56 to 0.64 ($p<0.01$), 0.65 to 0.78 ($p<0.001$) and 0.70 to 0.79 ($p<0.001$), for the *Clincial*, *RobustDoseOpt*, *LETopt* and *RobustLETopt* plans respectively. These constitute increases in TCP, due to the marginal therapeutic effect of LET_d, of 8, 8, 13 and 9. These increases in relative terms are 14, 14, 20 and 13 percent. The highest TCP value of 0.88 is observed for a *RobustLETopt* plan.

The predicted TCP values based on the Unkelbach RBE-weighted dose are significantly higher compared to the use of a constant RBE of 1.1 in the case of the *LETopt* plans ($p<0.001$) and *RobustLETopt* plans ($p<0.001$). In case of the *Clinical* plans, this difference is not significant ($p>0.05$), while this difference was marginally significant for the *RobustDoseOpt* plans ($p = 0.049$, $p<0.05$).

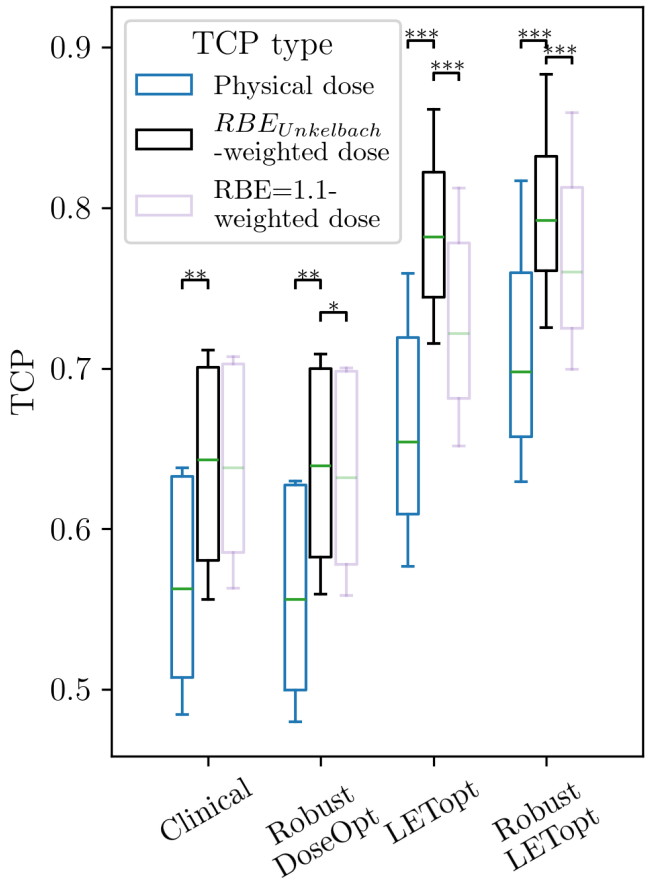

Figure 6: Comparison of TCP-values between the different plan types based on the non-RBE weighted physical dose, Unkelbach RBE-weighted dose and RBE=1.1-weighted dose for all patients. Asterisks denote statistical significance: * $p < 0.05$, ** $p < 0.01$,*** $p < 0.001$

4. Discussion

In this study we investigated the elevation of target LET without relying on LET-based optimization objectives, determined the maximum feasible increase while maintaining adequate target dose robustness and adherence to clinical goals, and assessed whether such increases translate into meaningful improvements in tumor control probability in a cohort of lung cancer patients treated with IMPT.

4.1. Summary of key findings

4.1.1. Target LET enhancement achieved with dose objectives only

The results demonstrated that it was possible to significantly increase target LET_d without using LET-based objectives, while still meeting clinical goals and adhering to adequate target dose robustness. This was achieved using opposing beam arrangements with beam-specific minimum dose objectives for proximal target segments, combined with high dose inhomogeneity (up to 140% of D_{pr}), CTV-based planning (instead of ITV-based), and reduced target dose robustness through probabilistic evaluation using polynomial chaos expansion.

4.1.2. LET_d increases of over 50% are feasible while adhering to clinical goals and robustness

With this strategy, the population-median of the mean target LET_d increased by more than 50% compared to clinical treatment plans and reference dose-optimized plans. When completely sacrificing target dose robustness, mean target LET_d increases could exceed 100%, representing the near upper limit achievable with this method. These increases corresponded to additional biological doses of ~6 Gy and ~8 Gy (based on the Unkelbach RBE model), raising population-median mean RBE from 1.1 to 1.16 in the probabilistically robust LET-optimized plans, and up to 1.21 in the non-robust LET-optimized plans.

4.1.3. Significant increases in TCP due to elevated LET, underestimated by a constant RBE in LET-optimized plans

The increases in biological dose to the marginal effect of LET, led to significantly higher TCP values across all plan types. While absolute TCP gains were larger for the LET-optimized plans, relative increases for the robust LET-optimized plans were comparable to those of the dose-optimized plans. Finally, the assumption of a constant RBE of 1.1 significantly underestimated TCP values reached by the LET-optimized treatment plans.

4.2. Discussion on employed LET-painting method

4.2.1. Lack of control in LET elevation and OAR sparing

While target LET elevation was possible without the use of LET-based objectives, the employed method presents challenges in terms of fine control over increasing target LET at the cost of target dose robustness and a lack of control regarding high LET surrounding the target. The elevated LET, in combination with the non-negligible dose surrounding the target, may increase toxicity to nearby OARs. The use of LET-based objectives would allow for more fine control in LET elevation and would allow maximum LET-objectives for sparing surrounding OARs.

4.2.2. Beam angle selection and ideal tumor location

Another challenge was selecting the two sets of opposing beam angles, since many angles are either clinically unacceptable due to constraints forbidding traversal through certain structures (e.g. arms, sides of breast, thorax support, table edges), or unfavorable in terms of increased OAR doses. In some cases these constraints forced the selection of angles that traverse through important OARs such as the heart. This seems to be at odds with the OAR-sparing ability of proton therapy, where one is able to spare OARs distal to the target, because of the finite range of protons. The location of the tumor in combination with these constraints also lead to variation between patients in the angle between the two sets of opposing beam angles, see Table S7. For some patients more narrow angles between the two sets of opposing beams were chosen, which may limit the confinement capability of high LET to the tumor only. It should also be noted that the strict use of opposing beam angles is not necessary, as some studies use beams in one general direction [39] or use almost opposing beam angles [28]. However, the potential for high-LET confinement to the tumor might decrease in doing so. Using many sets of opposing beam angles, as done in

proton arc therapy, could decrease the radiation dose to critical OARs when a beam has to traverse through it, while increasing the amount of LET elevation that is possible in the tumor [15].

4.2.3. *Effect of LET painting method on OAR dose metrics and clinical plan comparability*

Finally, is important to note that this LET-painting technique only targets the primary tumor, meaning to form a full clinical treatment plan that would normally also target secondary tumors and/or lymph nodes, additional fields should be employed. It is therefore also worth noting that the OAR doses observed in the *Clinical* treatment plans in this study are thus higher not only because of ITV-based planning, but also because of nodal irradiation that was included in some of the plans. The *RobustDoseOpt* plans therefore allow for the most fair comparison in terms of OAR-dose increases due to the LET-optimization strategy. As can be observed in Figures S8 and S9, there is no clear pattern between the LET-optimized and reference dose-optimized plans in terms of OAR-dose increases. In some cases the median OAR clinical goal values in the LET-optimized plans are higher, and in some cases even lower. This is likely caused by an interplay between increased entry doses for the LET-optimized treatment plans and lowered dose surrounding the target due to decreased target dose robustness. Depending on the clinical goal value and location of the OAR, one or the other effect could dominate in the resulting clinical goal value.

4.3. *Attainable target LET elevation*

4.3.1. *Comparison of LET elevation values to previous studies*

Increases in mean target LET_d of approximately 50% were found for the probabilistically robust LET-optimized treatment plans compared to clinical plans. Various amounts of mean target LET elevation, from a few percent [14] to over 100 percent [15], are noted in previous studies [11; 12; 14–18; 39–43], with various treatment sites and methodologies, including a.o. differences in the amount of inhomogeneity and the amount of beam angles. A direct comparison with values observed in other studies is therefore difficult. Nonetheless, to the best of our knowledge, no study explored the very limit of LET enhancement possible in terms of the available freedom in target dose robustness and inhomogeneity, which are linked to target LET in what has been previously described as a trilemma [13].

4.3.2. *The effect of target size and geometric constraints on target LET elevation*

Differences in target LET between patients could have partly been caused by differences in target volume. Multiple studies have found a significant negative correlation between target volume and target LET enhancement in both proton therapy [12][15] and carbon ion radiotherapy (CIRT) [44][10]. This relationship was further explored for the probabilistically robust LET-optimized plans in this study, see Figure S12. A linear regression was performed between mean target LET_d and target volume, resulting in a non-significant relationship ($p = 0.087 > 0.05$, $R^2 = 0.32$). It should be noted that the low sample size of this study might not allow for enough statistical power to detect such an effect.

Tumor location might have also influenced the amount of LET elevation possible per patient, as it influenced the employed beam angles and may influence the spots placed in and surrounding the tumor. For example, OARs close to, or in a beam path, in combination with OAR-sparing planning objectives, could prevent placement of spots that would have increased target LET otherwise and might be placed elsewhere in favor of OAR sparing. As discussed earlier, the patient geometry also dictates the chosen opposing beam angles, with different angles between the opposing sets potentially leading to less high-LET confinement to the tumor.

4.3.3. *Future potential for LET elevation through technological advances in imaging, optimization and delivery*

Although a maximum feasible increase of around 50% in mean target LET_d was found here, it is likely that this number will increase as technology advances. For example, the use and further development of new CT imaging techniques as Dual Energy CT and photon counting CT could allow for better stopping power estimation, thereby reducing the relative range error and allowing higher target LET increases [45][46]. The more these errors are reduced, the closer elevation levels as observed for the non-robust LET-optimized treatment plans in this study may be approached, with increases of over 100% in mean target LET_d. As hardware capabilities advance and more efficient Monte Carlo algorithms are developed, beam angle optimization at a more clinical timescale could become possible, allowing one to find more optimal solutions in terms of both dose and LET [41] [42]. Alternatively, deep learning, currently used for the fast prediction of optimal beam angles based on dose only, could also be used to predict optimal angles based on both dose and LET_d, by training on LET-optimized treatment plans. Finally, new treatment modalities such as proton arc, could enable the use of many angles, allowing for better confinement of high LET and higher amounts of LET elevation inside the tumor [15].

4.4. *Increases in TCP*

4.4.1. *Diminished LET-based TCP returns due to elevated dose*

An interesting finding is that despite the significantly higher target LET_d of the probabilistically robust LET-optimized treatment plans compared to the clinical and reference dose-optimized plans, the relative increase in TCP was comparable. This can be explained by the increased physical dose that was associated with the probabilistically robust LET-optimized plans and the decreased steepness of the EUD-based TCP curve for higher EUD values, shown in Figure S13. The increased dose in these plans was necessary to adhere to probabilistic coverage and causes the additional biological dose due to increased LET_d to have less of an effect on the TCP. This might indicate that with lower inhomogeneity at the same level of target dose robustness, resulting in lower target LET_d, might cause relatively larger increases in TCP due to LET_d, even though the absolute increase in LET_d might be lower.

4.4.2. *Robustness of TCP*

Although the highest TCP values were observed for the probabilistically robust LET-optimized treatment plans, it is impor-

tant to realize that these represent values in the nominal scenario, i.e. without the consideration of any uncertainties. However, due to the decreased target dose robustness of these plans, the LET_d, $LET_d \times D$ and therefore TCP values, are expected to vary more than those of the dose-optimized plans. To further investigate this, the robustness of these quantities for the different plan types for a single patient have been compared in Supplementary Materials Figure S14. As shown in this figure, as the target dose robustness decreases, so does the robustness of target LET_d values. Similarly, the decreased target dose robustness also gives rise to a larger spread in observed TCP values.

These findings highlight the complexity of determining whether LET-optimized treatment plans represent better clinical alternatives to simple dose escalation strategies, which would represent with higher robustness of dose and TCP. Furthermore, in this study probabilistic evaluation was used to maximize target LET, but it could also be used to reduce target dose robustness in order to improve target dose conformity. As this would reduce dose to neighboring OARs, this would allow to further escalate target dose, making dose escalation as a strategy more viable. These insights further highlight the need to compare dose-only versus dose-and-LET escalation strategies, with robustness analyses not limited to dose, but extending to robustness analysis of LET_d, $LET_d \times D$ and TCP.

4.4.3. Challenges in predicting the biological effect of LET elevation

Finally, it is good to note that predicting the exact biological effect due to an increase in LET is difficult. In the first place because of the dose-averaged LET quantity itself, which aims to quantify a microscopic phenomenon in a macroscopic voxel. As a result, different LET spectra inside a voxel can result in the same dose-averaged LET value, but give rise to different biological effects [47]. Furthermore, the linear Unkelbach RBE model was used throughout this study, but the increase in biological effect due to increased LET is not necessarily linear. The TCP values could have also been calculated on the basis of a non-linear RBE model such as the McNamara model, however choosing an adequate value for the radiosensitivity parameter α/β is not trivial and evaluation based on such models was therefore not included in this study.

4.5. Clinical Relevance and implications

4.5.1. Clinical implementation and discussions

Several components of the LET-enhancement strategy as proposed in this study are not yet ready for routine clinical application, including probabilistic evaluation and CTV-based planning for lung cancer, due to the non-perfect motion mitigation of gating and breath-hold, and the fact that the use of ultra high dose rates is still under investigation.

Although not directly implementable in the clinic yet, this study demonstrates that clinically significant improvements in terms of TCP are or could be possible due to target LET enhancement, while maintaining acceptable target dose robustness and adherence to clinical goals. Nonetheless, these increases in LET are associated with increased target dose inhomogeneity, reduced target dose robustness, increased beam entry dose and higher doses to certain OARs, depending on tumor location. Whether the potential gains in TCP justify these trade-offs should be a subject of clinical discussion, particularly in light of future developments that may increase target LET elevation even further or reduce the compromises needed to reach the same amount of LET enhancement. Key questions in such a discussion could include: How much dose inhomogeneity can be tolerated clinically? What level of reduced target dose robustness is acceptable for a given increase in tumor control probability? What amount of additional dose to OARs can be justified for a certain amount of increase in tumor control probability? Under what circumstances could beams be allowed to traverse OARs to enhance target LET? Finally; can we responsibly base such decisions given the uncertainty in biological modeling of increased LET effects?

The results also highlighted that a constant RBE of 1.1 could underestimate the biological effectiveness of LET-optimized treatment plans. As clinical interest in target LET escalation may grow further, there will be a need for consensus on reporting and evaluating RBE-weighted dose in tumors for better study comparison, similar to harmonization that is taking place in the case of RBE-weighted dose in OARs, in reaction to increased end-of-range RBE effects [48].

4.5.2. LET versus dose escalation

An important further point for discussion and investigation is whether dose escalation alone might be preferable to LET escalation, as it can be performed while retaining high amounts of target dose robustness and uniformity. Furthermore, LET escalation in this study was also associated with some dose escalation already, diminishing the gain in TCP due to LET alone. Nonetheless, raising the total dose is different from increasing the inhomogeneity, meaning OAR sparing for the same therapeutic effect could be superior in LET-escalated treatment plans, depending on the amount of dose and LET escalation involved when enhancing target LET. Given these potential advantages and trade-offs, further investigation and comparison studies are warranted.

4.5.3. Use of LET-based objectives

While the method used in this study to elevate target LET does not rely on explicit LET-based objectives, incorporating such objectives during optimization could improve control over the LET distribution, enabling more subtle LET increases, better sparing of OARs from high LET, and the integration of LET-based beam angle optimization. Furthermore, the use of LET-based objectives is suspected to allow for even higher LET increases. Still, such increases will always come at the cost of dose uniformity, robustness or increased OAR doses.

4.5.4. Ideal tumor location

For lung cancer in particular, a promising application of target RBE elevation via LET optimization could be treating large, centrally located tumors. These often lie near critical structures like the heart and spinal cord, limiting the ability to

deliver sufficient dose. Elevating LET, rather than dose, can improve biological coverage without exceeding OAR constraints. However, results showed LET elevation often coincided with increases in mean and beam entry dose, raising OAR dose exposure. Additionally, the opposing beam strategy used in this study will inevitably lead to beams traversing OARs in the case of centrally located tumors. These findings imply that centrally located tumors are likely to be difficult targets for successful target LET elevation. Accordingly, tumors located more peripherally might be better targets, as in such cases beams less often have to traverse through important OARs. Additionally, increases in mean or beam entry doses are less problematic, since critical OARs such as the heart or esophagus are not as close in proximity compared to centrally located tumors.

4.5.5. Sub-targeting to reduce LET enhancement drawbacks

To reduce drawbacks of the LET painting technique used in this study, LET elevation could be restricted to subregions of the tumor or GTV as a biological boost fraction, while robust dose is delivered to the full CTV. Sub-targeting also reduces the risk of high-LET and high-dose hotspots shifting outside the CTV due to reduced robustness, and lowers entry dose. Furthermore, since smaller targets might allow for higher LET to be reached, such an approach can also result in greater target LET enhancement.

4.5.6. Overcoming hypoxia-induced treatment resistance

Finally, an interesting application of target LET escalation is elevating LET in hypoxic subregions of a large primary tumor, to overcome treatment resistance. However, LET levels achievable with protons (around 1-10 $keV/\mu m$) are likely insufficient to significantly reduce the oxygen enhancement ratio (OER) and overcome hypoxia [49], for which heavier ions such as carbon would be needed.

4.6. Assumptions and limitations

4.6.1. Included patients and group size

Complex cases where target coverage was compromised to adhere to OAR-based clinical goals were excluded, which were mostly large centrally located tumors. This means that the included cases were mostly peripheral tumors with volumes typically on the lower side. As such, the observed LET increases may not generalize to (more) centrally located tumors or very large tumors. Nonetheless, there was a good amount of variation in tumor location and volume, which enhances the generalizability of the findings with respect to these variables.

The sample size of 10 patients, while limited, is comparable to or larger than most existing LET-optimization studies.

4.6.2. CTV-based planning

This study assumed ideal motion mitigation through techniques like breath-hold, gating, or ultra-high dose rates. While breath-hold and gating are established techniques, they are not error-free [50]. This means the full reduction from ITV to CTV contours for targeting the tumor is not completely justified. Still, tumor motion in the included cases was limited, resulting in minimal ITV-CTV differences. Furthermore, future work might help reduce errors related to these techniques. Alternatively, the employment of ultra high dose rates could enable CTV-only planning in the future, although its current clinical implementation remains under investigation.

4.6.3. Probabilistic evaluation

The probabilistic coverage goal used in this study was calibrated based on head and neck patients. Consequently, a slightly different probabilistic goal might have been required, based on similar calibration for a group of lung cancer patients. This means that the degree of robustness of the probabilistically robust treatment plans in this study might not be sufficient, or might be more than strictly necessary. This in turn influences the achievable LET values.

As shown in Supplementary Materials section S4, PCE for lung cancer patients showed larger errors for smaller volume metrics, overestimating the near-minimum dose (D_{98}) and underestimating near-maximum dose (D_2), though these errors were typically below 0.5 Gy. The more extreme near-minimum volume of 99.8 could therefore present with a higher error, overestimating the actual dose and therefore overestimating the robustness. Consequently, the robust LET-optimized treatment plans may have required slightly higher robustness. However, even with an error of 1 Gy, such errors are around one percent of the prescribed dose, meaning the effect on the overall LET elevation would have remained small.

Another assumption that could have induced some amount of error was assuming that the systematic and random setup errors were equal. In reality they may differ. Extracting these errors from patient registration data would have been more accurate.

4.6.4. TCP calculations

Some error in the TCP predictions is expected to exist, as no fractionation correction was performed in the determination of the radiobiological response parameters that form a part of the EUD-based TCP model [38]. Accordingly, no correction for fractionation was applied for the patient cohort in this study, although different fractionation schemes are present in this cohort, see Figure S1.

The choice for the exponent in the EUD also strongly influences TCP. As a value of -10 was taken for the exponent, low values will have a higher weight in the EUD computation. Consequently, the near-minimum LET_d values in the target could be important for reaching higher TCP values based on Unkelbach RBE-weighted dose. This implies that the use of range shifters, not used for the LET-optimized plans, could have increased their associated TCP values, as range shifters decrease near-maximum LET, but increase near-minimum LET.

4.6.5. Beam angle errors

In this study only setup and range errors were considered, even though more errors exist in proton therapy treatment. For example, gantry angle errors can exist, where the planned gantry angle is slightly different from the actual gantry angle during delivery. Furthermore, as different beam angles traverse different parts of the patient, the relative range error can differ for the different beam angles. These errors can cause differences in both

the dose and LET distribution that were not considered in this study.

4.7. Future work and recommendations

4.7.1. Assessment of biological dose to OARs

Further investigation should take place regarding potential increases in biological dose to OARs surrounding the target when using the LET-optimization method as employed in this study, seeing as the results suggest that elevation of LET was not fully confined to the target only.

4.7.2. Comparing LET-painting and LET-based optimization

A potential solution for increased biological dose to OARs surrounding the target due to elevated LET, is the use of LET-based objectives that can impose upper limits on LET in these structures. Furthermore, the use of LET-based objectives could also lead to more optimal solutions in terms of both dose and LET. Therefore it would be valuable to directly compare achievable target LET and biological doses in OARs between treatment plans created with LET-painting and those created with LET-based optimization objectives.

4.7.3. Comparison of LET-and-dose vs. dose-only escalation

As LET escalation was also associated with dose escalation, which reduced the relative contribution of LET to improving TCP, an important next step is to identify the optimal combination of dose and LET escalation. This combination should maximize the benefits of LET elevation, while minimizing its drawbacks. Such a strategy should be directly compared to simple dose escalation, by evaluating OAR doses and achievable TCP, especially under uncertainty. This is because LET-elevated plans are likely to be less robust, meaning the LET distribution and TCP will vary more. To this end, polynomial chaos expansion could be used to properly assess the robustness of not only dose, but also LET_d, $LET_d \times D$ and TCP.

4.7.4. Phantom measurements of LET-optimized plans

Finally, extensive QA of LET-optimized treatment plans on phantoms could already be performed, to examine whether the highly inhomogeneous dose distributions with reduced robustness can be delivered safely, and to measure whether the target LET elevations predicted in silico are also achieved in real-life measurements [51].

4.8. Conclusion

Clinically significant increases in tumor control probability due to enhanced target LET are possible without the use of LET-based objectives, by leveraging high inhomogeneity and strictly meeting robustness levels as found in photon therapy through probabilistic evaluation. Whether the potential gains in TCP justify these trade-offs should be a subject of clinical discussion, particularly in light of future developments that may increase target LET elevation even further or reduce the trade-offs needed to reach the same amount of LET enhancement. Future work should further identify the ideal combination of both dose and LET escalation, which maximizes the benefits of LET elevation and minimizes its drawbacks. Such a strategy should be directly compared to simple dose escalation, in terms of OAR doses and attainable TCP, especially under uncertainty. While unlikely to be implemented in the clinic soon, target LET enhancement could have the potential to fully exploit the particle properties of protons, widening the therapeutic window.

Acknowledgments

I am deeply grateful to my supervisors for their involvement in this project and their guidance throughout.

Jenneke's daily support helped me navigate the many challenges of writing a thesis, both in terms of content and the process itself. Her participation in the project made it easy to reach out when stuck, and made me feel like I always had someone by my side during the project.

Steven's contagious curiosity and enthusiasm, coupled with his great knowledge and experience in the field of (medical) physics, helped shape the project and offered invaluable directions to create an interesting and coherent piece of work.

I would also like to express my appreciation to Mischa for proposing and helping set up this project, offering behind-the-scenes support, and providing a fresh, neutral perspective that helped provide directions and ensure coherence of the project.

A special word of thanks also goes out to my fellow students, who made daily thesis life enjoyable, as well as occasionally helping out and being there for me. I would also like to thank all the lovely staff at HollandPTC who made it possible for me to perform my thesis, in particular Koen Krama, who introduced me to proton therapy planning.

Lastly, I want to thank HollandPTC for the opportunity that I was given to finish my studies at a truly unique place, where decades of research and interdisciplinary collaboration come together to provide patients with the best treatment possible. I still remember the fascination I felt during my bachelor's in Applied Physics when I first learned about proton therapy. Having been able to work on this topic as my final project has therefore felt truly special.

References

[1] Eirini Terpsi Vitti and Jason L Parsons. The radiobiological effects of proton beam therapy: Impact on dna damage and repair. *Cancers*, 11(7), 2019.

[2] Harald Paganetti, Eleanor Blakely, Alejandro Carabe-Fernandez, David J. Carlson, Indra J. Das, Lei Dong, David Grosshans, Kathryn D. Held, Radhe Mohan, Vitali Moiseenko, Andrzej Niemierko, Robert D. Stewart, and Henning Willers. Report of the aapm tg-256 on the relative biological effectiveness of proton beams in radiation therapy. *Medical Physics*, 46(3):e53–e78, 2019.

[3] Eivind Rørvik, Lars Fredrik Fjæra, Tordis J Dahle, Jon Espen Dale, Grete May Engeseth, Camilla H Stokkevåg, Sara Thörnqvist, and Kristian S Ytre-Hauge. Exploration and application of phenomenological rbe models for proton therapy. *Physics in Medicine amp; Biology*, 63(18):185013, Sep 2018.

[4] Stephen J McMahon. Proton rbe models: Commonalities and differences. *Physics in Medicine amp; Biology*, 66(4), Feb 2021.

[5] Melissa McIntyre, Puthenparampil Wilson, Peter Gorayski, and Eva Bezak. A systematic review of let-guided treatment plan optimisation in proton therapy: Identifying the current state and future needs. *Cancers*, 15(17):4268, Aug 2023.

[6] Christopher R. Peeler, Dragan Mirkovic, Uwe Titt, Pierre Blanchard, Jillian R. Gunther, Anita Mahajan, Radhe Mohan, and David R. Grosshans. Clinical evidence of variable proton biological effectiveness in pediatric patients treated for ependymoma. *Radiotherapy and Oncology*, 121(3):395–401, 2016.

[7] Chia-Chun Wang, Aimee L. McNamara, Jungwook Shin, Jan Schuemann, Clemens Grassberger, Alphonse G. Taghian, Rachel B. Jimenez, Shannon M. MacDonald, and Harald Paganetti. End-of-range radiobiological effect on rib fractures in patients receiving proton therapy for breast cancer. *International Journal of Radiation Oncology, Biology, Physics*, 107(3):449–454, 2020.

[8] Clemens Grassberger, Alexei Trofimov, Anthony Lomax, and Harald Paganetti. Variations in linear energy transfer within clinical proton therapy fields and the potential for biological treatment planning. *International Journal of Radiation Oncology, Biology, Physics*, 80(5):1559–1566, Aug 2011.

[9] Drosoula Giantsoudi, Clemens Grassberger, David Craft, Andrzej Niemierko, Alexei Trofimov, and Harald Paganetti. Linear energy transfer-guided optimization in intensity modulated proton therapy: Feasibility study and clinical potential. *International Journal of Radiation Oncology Biology Physics*, 87(1):216–222, 2013.

[10] Mansure Schafasand, Andreas Franz Resch, Ankita Nachankar, Joanna Gora, Giovanna Martino, Erik Traneus, Lars Glimelius, Dietmar Georg, Piero Fossati, Antonio Carlino, and Markus Stock. Dose averaged linear energy transfer optimization for large sacral chordomas in carbon ion therapy. *Medical Physics*, 51(6):3950–3960, 2024.

[11] M. Vieceli, J. Park, W. C. Hsi, M. Saki, N. P. Mendenhall, P. Johnson, and M. Artz. Potential therapeutic improvements in prostate cancer treatment using pencil beam scanning proton therapy with letd optimization and disease-specific rbe models. *Cancers*, 16(4), 2024.

[12] H. S. Wan Chan Tseung, J. Ma, C. R. Kreofsky, D. J. Ma, and C. Beltran. Clinically applicable monte carlo–based biological dose optimization for the treatment of head and neck cancers with spot-scanning proton therapy. *International Journal of Radiation Oncology Biology Physics*, 95(5):1535–1543, 2016.

[13] Albin Fredriksson, Lars Glimelius, and Rasmus Bokrantz. The let trilemma: Conflicts between robust target coverage, uniform dose, and dose-averaged let in carbon therapy. *Medical Physics*, 50(12):7338–7348, 2023.

[14] Qingkun Fan, Xiaoyuan Zhang, Riao Dao, Yujia Qian, Lewei Zhao, Xiaoqiang Li, Xuanfeng Ding, Gang Liu, and Shuyang Dai. Optimizing linear energy transfer distribution in intensity-modulated proton therapy using the alternating direction method of multipliers. *Frontiers in Oncology*, 14, 2024.

[15] X. Li, X. Ding, W. Zheng, G. Liu, G. Janssens, K. Souris, A. M. Barragán-Montero, D. Yan, C. Stevens, and P. Kabolizadeh. Linear energy transfer incorporated spot-scanning proton arc therapy optimization: A feasibility study. *Frontiers in Oncology*, 11, 2021.

[16] Mei Chen, Wenhua Cao, Pablo Yepes, Fada Guan, Falk Poenisch, Cheng Xu, Jiayi Chen, Yupeng Li, Ivan Vazquez, Ming Yang, X. Ronald Zhu, and Xiaodong Zhang. Impact of dose calculation accuracy on inverse linear energy transfer optimization for intensity-modulated proton therapy. *Precision Radiation Oncology*, 7(1):36–44, 2023.

[17] Chenbin Liu, Samir H. Patel, Jie Shan, Steven E. Schild, Carlos E. Vargas, William W. Wong, Xiaoning Ding, Martin Bues, and Wei Liu. Robust optimization for intensity modulated proton therapy to redistribute high linear energy transfer from nearby critical organs to tumors in head and neck cancer. *International Journal of Radiation Oncology Biology Physics*, 107(1):181–193, 2020.

[18] Ya-Nan Zhu, Weijie Zhang, Jufri Setianegara, Yuting Lin, Erik Traneus, Yong Long, Xiaoqun Zhang, Rajeev Badkul, David Akhavan, Fen Wang, Ronald C. Chen, and Hao Gao. Proton arc based lattice radiation therapy: feasibility study, energy layer optimization and let optimization. *Physics in Medicine and Biology*, 69(21), 2024.

[19] Iwona Ziółkowska-Suchanek. Mimicking tumor hypoxia in non-small cell lung cancer employing three-dimensional in vitro models. *Cells*, 10(1), 2021.

[20] Johannes A. Langendijk, Philippe Lambin, Dirk De Ruysscher, Joachim Widder, Mike Bos, and Marcel Verheij. Selection of patients for radiotherapy with protons aiming at reduction of side effects: The model-based approach. *Radiotherapy and Oncology*, 107(3):267–273, 2013.

[21] Francesco Fracchiolla, Erik Engwall, Martin Janson, Fredrik Tamm, Stefano Lorentini, Francesco Fellin, Mattia Bertolini, Carlo Algranati, Roberto Righetto, Paolo Farace, Maurizio Amichetti, and Marco Schwarz. Clinical validation of a gpu-based monte carlo dose engine of a commercial treatment planning system for pencil beam scanning proton therapy. *Physica Medica*, 88:226–234, 2021.

[22] Andries N. Schreuder, Daniel S. Bridges, Lauren Rigsby, Marc Blakey, Martin Janson, Samantha G. Hedrick, and John B. Wilkinson. Validation of the raystation monte carlo dose calculation algorithm using a realistic lung phantom. *Journal of Applied Clinical Medical Physics*, 20(12):127–137, 2019.

[23] Fredrik Kalholm, Leszek Grzanka, Erik Traneus, and Niels Bassler. A systematic review on the usage of averaged let in radiation biology for particle therapy. *Radiotherapy and Oncology*, 161:211–221, 2021.

[24] Laleh Kardar, Yupeng Li, Xiaoqiang Li, Heng Li, Wenhua Cao, Joe Y. Chang, Li Liao, Ronald X. Zhu, Narayan Sahoo, Michael Gillin, Zhongxing Liao, Ritsuko Komaki, James D. Cox, Gino Lim, and Xiaodong Zhang. Evaluation and mitigation of the interplay effects of intensity modulated proton therapy for lung cancer in a clinical setting. *Practical Radiation Oncology*, 4(6):e259–e268, 2014.

[25] Albin Fredriksson, Anders Forsgren, and Björn Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Medical Physics*, 38(3):1672–1684, 2011.

[26] Erik W. Korevaar, Steven J.M. Habraken, Daniel Scandurra, Roel G.J. Kierkels, Mirko Unipan, Martijn G.C. Eenink, Roel J.H.M. Steenbakkers, Stephanie G. Peeters, Jaap D. Zindler, Mischa Hoogeman, and Johannes A. Langendijk. Practical robustness evaluation in radiotherapy – a photon and proton-proof alternative to ptv-based plan evaluation. *Radiotherapy and Oncology*, 141:267–274, 2019.

[27] Douglas Jones. Icru report 50—prescribing, recording and reporting photon beam therapy. *Medical Physics*, 21(6):833–834, 1994.

[28] Marcus Fager, Iuliana Toma-Dasu, Maura Kirk, Derek Dolney, Eric S. Diffenderfer, Neha Vapiwala, and Alejandro Carabe. Linear energy transfer painting with proton therapy: A means of reducing radiation doses with equivalent clinical effectiveness. *International Journal of Radiation Oncology, Biology, Physics*, 91(5):1057–1064, 2015.

[29] Matthias Guckenberger, Nicolaus Andratschke, Karin Dieckmann, Mischa S. Hoogeman, Morten Hoyer, Coen Hurkmans, Stephanie Tanadini-Lang, Eric Lartigau, Alejandra Méndez Romero, Suresh Senan, and Dirk Verellen. Estro acrop consensus guideline on implementation and practice of stereotactic body radiotherapy for peripherally located early stage non-small cell lung cancer. *Radiotherapy and Oncology*, 124(1):11–17, 2017.

[30] Karin Lindberg, Vitali Grozman, Kristin Karlsson, Sara Lindberg, Ingmar Lax, Peter Wersäll, Gitte Fredberg Persson, Mirjana Josipovic, Azza Ahmed Khalil, Ditte Sloth Moeller, Jan Nyman, Ninni Drugge, Per Bergström, Jörgen Olofsson, Lotte Victoria Rogg, Christina Ramberg, Charlotte Kristiansen, Stefan Starup Jeppesen, Tine Bjørn Nielsen, Britta Lödén, Hans-Olov Rosenbrand, Silke Engelholm, André Haraldsson, Charlotte Billiet, and Rolf Lewensohn. The hilus-trial—a prospective nordic multicenter phase 2 study of ultracentral lung tumors treated with stereotactic body radiotherapy. *Journal of Thoracic Oncology*, 16(7):1200–1210, 2021.

[31] Jesús Rojo-Santiago, Erik Korevaar, Zoltán Perkó, Stefan Both, Steven J.M. Habraken, and Mischa S. Hoogeman. Ptv-based vmat vs. robust impt for head-and-neck cancer: A probabilistic uncertainty analysis of clinical plan evaluation with the dutch model-based selection. *Radiotherapy and Oncology*, 186:109729, 2023.

[32] Zoltán Perkó, Sebastian R van der Voort, Steven van de Water, Charlotte M H Hartman, Mischa Hoogeman, and Danny Lathouwers. Fast and accurate sensitivity analysis of impt treatment plans using polynomial chaos expansion. *Physics in Medicine Biology*, 61(12):4646, may 2016.

[33] Sebastian van der Voort, Steven van de Water, Zoltán Perkó, Ben Heijmen, Danny Lathouwers, and Mischa Hoogeman. Robustness recipes for minimax robust optimization in intensity modulated proton therapy for oropharyngeal cancer patients. *International Journal of Radiation Oncology, Biology, Physics*, 95(1):163–170, 2016. Particle Therapy Special Edition.

[34] A Meijers, Oria C Seller, J Free, D Bondesson, Carmen Seller Oria, M Rabe, K Parodi, G Landry, J A Langendijk, S Both, C Kurz, and AC Knopf. Assessment of range uncertainty in lung-like tissue using

a porcine lung phantom and proton radiography. *Physics in Medicine Biology*, 65(15):155014, jul 2020.

[35] Jan Unkelbach, Pablo Botas, Drosoula Giantsoudi, Bram L. Gorissen, and Harald Paganetti. Reoptimization of intensity modulated proton therapy plans based on linear energy transfer. *International Journal of Radiation Oncology, Biology, Physics*, 96(5):1097–1106, 2016.

[36] Jan Unkelbach, Pablo Botas, Drosoula Giantsoudi, Bram L. Gorissen, and Harald Paganetti. Reoptimization of intensity modulated proton therapy plans based on linear energy transfer. *International Journal of Radiation Oncology, Biology, Physics*, 96(5):1097–1106, 2016.

[37] Andrzej Niemierko. Reporting and analyzing dose distributions: A concept of equivalent uniform dose. *Medical Physics*, 24(1):103–110, 1997.

[38] Paul Okunieff, David Morgan, Andrzej Niemierko, and Herman D. Suit. Radiation dose-response of human tumors. *International Journal of Radiation Oncology, Biology, Physics*, 32(4):1227–1237, 1995.

[39] Jiasen Ma, Sonja Dragojevic, Nicholas B. Remmes, Nicole L. Mendelson, Jake A. Kloeber, Daniel K. Ebner, Zheming Wu, Heather J. Gunn, Kenneth W. Merrell, Christopher L. Hallemeier, and et al. Linear energy transfer optimized proton therapy for rectal cancer. *Radiotherapy and Oncology*, 207:110850, Jun 2025.

[40] Hadis Moazami Goudarzi, Gino Lim, David Grosshans, Radhe Mohan, and Wenhua Cao. Incorporating variable rbe in impt optimization for ependymoma. *Journal of Applied Clinical Medical Physics*, 25(1), 2024.

[41] Wenbo Gu, Dan Ruan, Wei Zou, Lei Dong, and Ke Sheng. Linear energy transfer weighted beam orientation optimization for intensity-modulated proton therapy. *Medical Physics*, 48(1):57–70, 2021.

[42] Xuemin Bai, Gino Lim, Wenhua Cao, Taewoo Lee, Cunjiang Yu, and Ying Lin. *Linear Energy Transfer (LET) - Guided Optimization Incorporating Biological Effectiveness for Intensity-Modulated Proton Therapy*. Thesis, University of Houston, 2019.

[43] Wenhua Cao, Azin Khabazian, Pablo P. Yepes, Gino Lim, Falk Poenisch, David R. Grosshans, and Radhe Mohan. Linear energy transfer incorporated intensity modulated proton therapy optimization. *Physics in Medicine and Biology*, 63(1), 2018.

[44] R. Kohno, M. Koto, H. Ikawa, S. H. Lee, K. Sato, M. Hashimoto, T. Inaniwa, and T. Shirai. High–linear energy transfer irradiation in clinical carbon-ion beam with the linear energy transfer painting technique for patients with head and neck cancer. *Advances in Radiation Oncology*, 9(1), 2024.

[45] Ming Yang, Patrick Wohlfahrt, Chenyang Shen, and Hugo Bouchard. Dual- and multi-energy ct for particle stopping-power estimation: Current state, challenges and potential. *Physics in Medicine amp; Biology*, 68(4), Feb 2023.

[46] Vicki T. Taasti, David C. Hansen, Gregory J. Michalak, Amanda J. Deisher, Jon J. Kruse, Ludvig P. Muren, Jørgen B. B. Petersen, and Cynthia H. McCollough. Theoretical and experimental analysis of photon counting detector ct for proton stopping power prediction. *Medical Physics*, 45(11):5186–5196, 2018.

[47] Lena Heuchel, Christian Hahn, Jakob Ödén, Erik Traneus, Jörg Wulff, Beate Timmermann, Christian Bäumer, and Armin Lühr. The dirty and clean dose concept: Towards creating proton therapy treatment plans with a photon-like dose response. *Medical Physics*, 51(1):622–636, 2024.

[48] Dirk Wagenaar, Steven J.M. Habraken, Ilaria Rinaldi, Daniëlle B.P. Eekers, Miranda Kramer, Jaap P.M. Jaspers, Dik van Gent, Lara Barazzuol, Yvonne L.B. Klaver, Jaap Zindler, Ida Coremans, Inge Compter, Daniel Scandurra, Hiska L. van der Weide, Stefan Both, Mischa Hoogeman, Mirko Unipan, and Alejandra Méndez Romero. Evaluating and reporting let and rbe-weighted dose in proton therapy for glioma – the dutch approach. *Radiotherapy and Oncology*, 202:110653, 2025.

[49] Walter Tinganelli, Marco Durante, Ryoichi Hirayama, Michael Krämer, Andreas Maier, Wilma Kraft-Weyrather, Yoshiya Furusawa, Thomas Friedrich, and Emanuele Scifoni. Kill-painting of hypoxic tumours in charged particle therapy. *Scientific Reports*, 5(1), Nov 2015.

[50] Yuki Tominaga, Masaki Suga, Mikuni Takeda, Yuki Yamamoto, Takashi Akagi, Takahiro Kato, Sunao Tokumaru, Michinori Yamamoto, and Masataka Oita. Comparing interplay effects in scanned proton therapy of lung cancer: Free breathing with various layer and volume rescanning versus respiratory gating with different gate widths. *Physica Medica*, 120:103323, 2024.

[51] V.A. Pan, L.T. Tran, T. Madden, B. Clasie, H. Kooy, A.B. Rosenfeld, and N. Depauw. Letd optimization verification with an soi microdosimeter. *International Journal of Radiation Oncology, Biology, Physics*, 119(3):946–956, 2024.

[52] Mark Geurts. 3D gamma computation in MATLAB. Code accessible at Github: https://github.com/mwgeurts/gamma, July 2024. DOI: 10.5281/zenodo.12735617.

[53] Daniel A. Low, William B. Harms, Sasa Mutic, and James A. Purdy. A technique for the quantitative evaluation of dose distributions. *Medical Physics*, 25(5):656–661, 1998.

Supplementary Materials

S1. Patient and treatment characteristics

Table S1: Patient and treatment characteristics

Patient Alias	Histology	Location primary tumor	Target Volume [cm^3]	Stage	Prescribed dose [Gy(RBE=1.1)]	N fractions
Long001	Adenocarcinoma	Right upper lobe	28,25	IIIA	55	20
Long005	Squamous cell carcinoma	Right lower lobe	80,72	IIIB	55	20
Long007	Squamous cell carcinoma	Right main bronchus	112,36	IIIA	60	30
Long009	Squamous cell carcinoma	Left upper lobe	38,34	IIIC	66	24
long010	Adenocarcinoma	Left lower lobe	142,83	IVA	55	20
long012	Squamous cell carcinoma	Right upper lobe	8,09	IIIC	66	24
long015	Neoplasm, malignant	Mediastinum	19,28	-	60	30
long016	Unknown	Right lower lobe	25,97	IIIB	60	30
long017	Adenocarcinoma	Left upper lobe	148,48	IIIC	66	24
long018	Adenocarcinoma	Left lower lobe	38,8	IIIA	66	24

S2. Clinical goals and planning objectives

Table S2: Clinical goals

ROI	**Clinical Goal**
CTVp	$D_{98} \geq 0.95 D_{pr}$
CTVp	D2 $\leq 1.07 D_{pr}$
Esophagus	$D_{avg} \leq$ 34Gy(RBE) $D_{1cc} \leq$ 60Gy(RBE)
Lungs −GTV	$D_{avg} \leq$ 20Gy(RBE) D35 $\leq$ 20Gy(RBE)
Heart	$D_{avg} \leq$ 0Gy(RBE) D30 $\leq$ 30Gy(RBE)
Spinalcord	$D_{0.03cc,core} \leq$ 50Gy(RBE) $D_{0.03cc,surf} \leq$ 60Gy(RBE)
Skin	$D_{0.03cc} \leq$ 50Gy(RBE)

Table S3: Initial set of planning objectives used for creation of the clinical treatment plans.

ROI	Planning objective	Initial Weight	Robust	Constraint
ITV	Min dose D_{pr} Gy (RBE)	300	X	
ITV	Max dose D_{pr} Gy (RBE)	100	X	
Brachial Plexus (left)	Max dose 60 Gy (RBE)	0	X	
Brachial Plexus (right)	Max dose 60 Gy (RBE)	0	X	
Spinal Cord Surface	Max dose 60 Gy (RBE)	0	X	
Spinal Cord Core	Max dose 54 Gy (RBE)	0	X	
Thyroid gland-ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Liver-ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Trachea-ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Heart-ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Lung (lef) - ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Lung (right) - ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Esophagus-ITV	Max EUD D_{pr} Gy (RBE), A=1	1		
Trachea	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
Heart	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
Esophagus	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
Lung (left)	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
Lung (right)	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
Liver	Dose fall-off D_{pr} to 20 Gy (RBE), Low dose distance 3.00 cm	1		
External	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External	Max dose X* (RBE)	1		X
External	Max dose 1.05 D_{pr}	1	X	

**Clinical goal is applied for every beam separately, with its value dependent upon the amount of beam angles that are employed*

Table S4: Set of planning objectives used for creation of the probabilistically robust LET-optimized (*RobustLETopt*) treatment plans.

ROI	Planning objective	Initial Weight	Robust	Constraint
CTVp	Min dose D_{pr} Gy (RBE)	2500	X	
CTVp	Max dose D_{pr} Gy (RBE)	1500	X	
Esophagus	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Heart	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Lungs-GTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Spinal Cord	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Max Dose 1.05 D_{pr}	5000		
Proximal target segment	Min dose 5 Gy (RBE)*	100		

**Clinical goal is applied for every beam and corresponding target segment separately.*

Table S5: Set of planning objectives used for creation of the non-robust LET-optimized (*LETopt*) treatment plans.

ROI	Planning objective	Initial Weight	Robust	Constraint
CTVp	Min dose D_{pr} Gy (RBE)	2500		
CTVp	Max dose D_{pr} Gy (RBE)	1500		
Esophagus	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Heart	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Lungs-GTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Spinal Cord	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Max Dose 1.05 D_{pr}	5000		
Proximal target segment	Min dose 5 Gy (RBE)*	100		

**Clinical goal is applied for every beam and corresponding target segment separately*

Table S6: Set of planning objectives used for creation of the robust dose-optimized reference treatment plans (*RobustDoseOpt*).

ROI	Planning objective	Initial Weight	Robust	Constraint
CTVp	Min dose D_{pr} Gy (RBE)	2500	X	
CTVp	Max dose D_{pr} Gy (RBE)	1500	X	
Esophagus	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Heart	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Lungs-GTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
Spinal Cord	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Dose fall-off D_{pr} to 0 Gy (RBE), Low dose distance 2.00 cm	1		
External-CTV	Max Dose 1.05 D_{pr}	5000		

S3. Target contour subdivision into proximal target segments

The primary CTV (CTV_p) was divided into four segments, based on spatial proximity to each beam direction, with the center of the CTV_p as reference point. This segmentation relied on the fact that for four beams, the boundaries between segments are simply their bisectors. To create the contours of a proximal target segment for a certain beam angle, the angles of the bisectors between the current beam angle and the other two beam angles of the other set of opposing beam angles were determined, knowing that the angle of a bisecting line is the average of its associated beam angles.

Next, a cube-shaped ROI was defined, with its origin equal to that of the center of the CTV_p. The cube-shaped ROI was then rotated using matrix multiplication of its spatial position matrix with a rotation matrix R, see Formula A.1.

$$R_{\theta,(p_x,p_y)} = \begin{bmatrix} cos(\theta) & -sin(\theta) & 0 & p_x \\ sin(\theta) & cos(\theta) & 0 & p_y \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{A.1}$$

The smallest angle of the two bisecting lines was used for rotation. Then, the cube-shaped ROI was translated in the positive direction of both bisecting angles by half the width of the cube-shaped ROI, using matrix multiplication of the spatial position matrix with a translation matrix T, see Formula A.2.

$$T_{\theta,d} = \begin{bmatrix} 1 & 0 & 0 & dcos(\theta) \\ 0 & 1 & 0 & dsin(\theta) \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{A.2}$$

This causes two of the edges of the cube-shaped ROI to become exactly aligned with the two bisecting lines. Then, using RayStation ROI algebra, a new ROI is created that follows from the geometric union of the full CTV contours and the cube-shaped ROI. The size of the cube-shaped ROI was set to of 15x15x15 cm, which was large enough compared to the size of the tumors in all cases.

This process was repeated and automated using RayStation scripting for all beam angles, to obtain all four proximal target segments. Figure S1 illustrates this method.

Note that in the case of more beam angles, this geometrical trick becomes less straightforward, and exact division of segments using e.g. voxel-wise distance calculations are preferable. It is also important note is that gantry angles (θ) in RayStation are defined with respect to the patient's coordinate system in the coronal plane. As such, an angle of 0 degrees corresponds to the negative anterior direction, i.e. when the beam is pointing exactly into the patient's front. Furthermore, angles increase in a clockwise fashion when viewed from the foot of the patient (head-first supine position). For example, 90 degrees corresponds to a lateral beam pointing from left to right, 180 degrees corresponds to a posterior beam (from back to front) and 270 degrees corresponds to a lateral beam pointing from right to left.

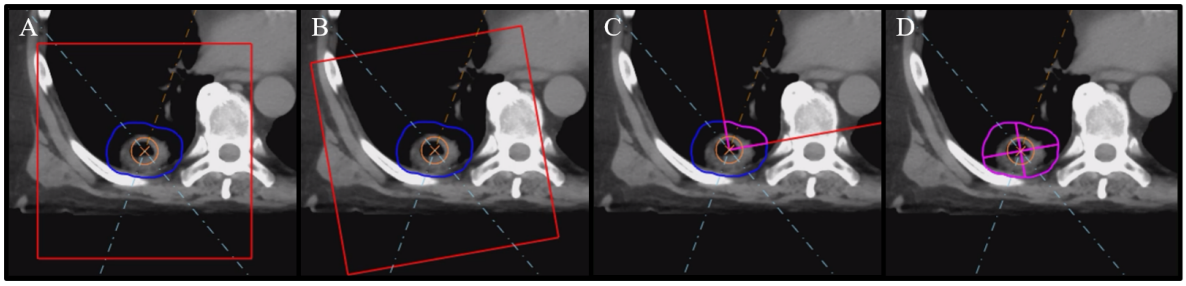

Figure S1: A visualization of the segmentation algorithm that divides the primary tumor CTV (CTV_p) into four proximal target segments (PTS), which involves per beam direction and associated PTS: (A) The creation of a cube-shaped ROI with its center equal to that of the CTV_p, (B) rotation of the cube-shaped ROI according to the bisecting angle between the beam-direction angle and the closest angle of the other set of two opposing beam angles, (C) translation of half the cube's width into the positive direction of the bisecting angles between the beam-direction angle and both beam angles of the other set of opposing beam angles, followed by the geometric union of the cube-shaped ROI and full CTV to form the PTS, (D) and repeating this process for the other beam-direction angles to create all four PTSes.

S4. PCE validation

S4.1. Methods

Validation of PCE for the lung treatment site largely followed the validation performed by the original study that introduced PCE to the field of proton therapy [32]. To this end, dose predictions made by PCE models of five clinical lung cancer IMPT treatment plans were compared to dose distributions as calculated by the RayStation MC dose engine for 425 uncertainty scenarios $\vec{\xi}$. These scenarios were sampled using the combination of 7 equally distanced points along each error axis ($\xi_{AP}, \xi_{LR}, \xi_{SI}, \rho$), spanning their respective 90% confidence intervals. Only points falling inside the contour that encompasses 90% of the probability space of the combined multivariate normal distribution of the individual errors were considered, which can be found analytically according to expression A.3.

$$\left(\vec{\xi} - \vec{\mu}\right) \mathbf{\Sigma^{-1}} \left(\vec{\xi} - \vec{\mu}\right) \leq \chi_p^2(\alpha) \tag{A.3}$$

With $\vec{\xi}$ the 4D-uncertainty vector, $\vec{\mu}$ the vector containing the statistical mean of the errors making up the uncertainty vector, $\mathbf{\Sigma}$ the covariance matrix which, due to statistical independence, is equal to a diagonal matrix containing the variance of the error variables on its diagonal and $\chi_p(\alpha)$ the 1-αth percentile of the chi-square distribution with p degrees of freedom, in this case equal to the amount of error variables (p=4) [**?**]. α was chosen to be 0.1 to include 90% of possible error scenarios in a grid-like structure, therefore equally representing both small error scenarios and relatively larger error scenarios. This resulted in a total of 425 valid scenarios used for validation. Figure S2 illustrates the selection of these valid scenarios using 7 equally distanced points along each uncertainty variable axis for two of the principal setup-error axes and the relative range error dimension.

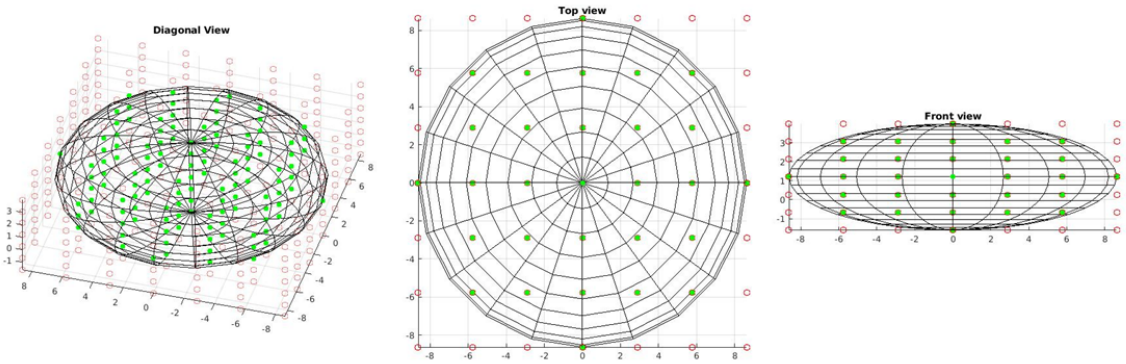

Figure S2: Visualization of the included sample points used for validation in the uncertainty space, consisting of three dimensions of setup error and a single dimension for the relative range error. Pleas note that the actual uncertainty space is four-dimensional, however here it is presented as being 3D, by including only two of the three setup error dimensions and the single relative range error dimension.

The difference between PCE dose predictions and the ground truth MC-based dose calculations of the different error scenarios were compared based on 3D average absolute difference, the worst 2% 3D difference and based on 3D gamma-analysis with strict dose-difference criterion and distance-to-agreement (DTA) criterion of 1% of the prescribed dose and 1 mm, as implemented in MatLab by [52], based on the formalism as used by D.A. Low et al. [53].

Furthermore, several DVH-parameter values were compared for the target, including the D_{98}, D_2, D_{mean}, D_{min} and D_{max}. Similarly, DVH-parameters used in the clinical goals for the heart, esophagus and spinal cord were compared, offering a wide variety of different DVH-parameters for validation.

S4.2. Results

Figure S3, shows cumulative volume histograms of the average absolute difference, worst 2% difference and gamma analysis results for the 425 uncertainty scenarios. As can be observed, for most of the scenarios average differences remain well below 1 Gy, worst 2% differences remain well below half a Gray and gamma analysis results indicate that for all scenarios, the passing rate is well above 97 percent (compared to the goal that is often stated for gamma analyses of 95%). One can observe average differences of up to 3 Gy for a single patient in a single uncertainty scenario between PCE and MC dose engine predictions. Differences in accuracy of the PCE-model predictions between the different patients can be observed.

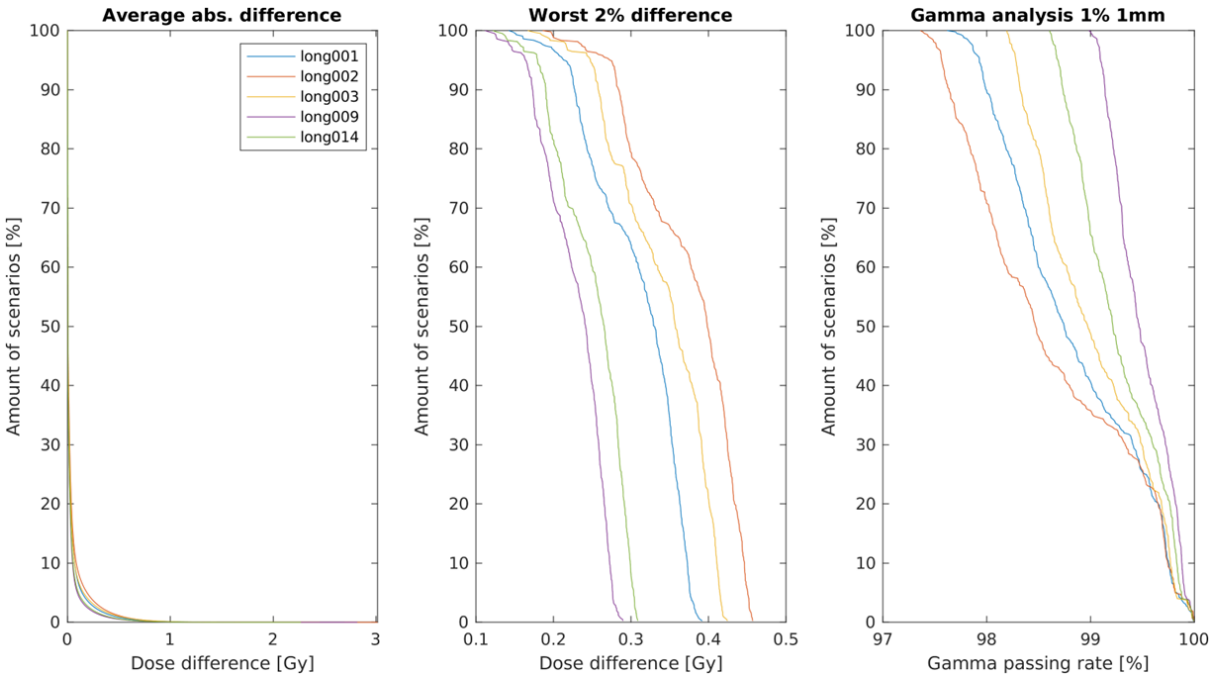

Figure S3: Cumulative volume histograms displaying the average absolute differences, worst 2% differences and gamma analysis results (1%,1 mm) between dose predictions made by PCE models for five clinical lung cancer IMPT treatment plans.

Figure S4 shows the difference in DVH-parameters for the target and OARs between 425 uncertainty scenario PCE dose predictions and corresponding MC dose calculations for the five lung cancer IMPT treatment plans. Looking at the near-minimum (D_{98}) and near-maximum (D_{98}) values for the target, we can see that the PCE predictions seem to consistently predict higher near-minimum values and lower near-maximum values, however these differences are mostly below an absolute value of 0.5 Gray, which corresponds to less than one percent the prescribed doses of these plans. Especially the predicted mean target dose values seem to align very well with MC-computations. The smaller the volume that is considered in the dose metric, the higher the error becomes, as can be seen by looking at the min and max values, where absolute differences of up to 2 Gray are observed, although most remain within 0.5 Gray. Again, we see for the maximum values a general trend underestimation and for the minimum values a trend of overestimation.

This high accuracy for larger volume metrics and mean values is also observed for the clinical goal metrics of the heart (D_{30},D_{mean}) and the mean dose metric for the esophagus. Again, larger errors are observed for smaller volume dose metrics with errors ranging from -1 to 1 Gray for the D_{1cc} metric of the esophagus and $D_{0.03cc}$ metric of the core of the spinal cord.

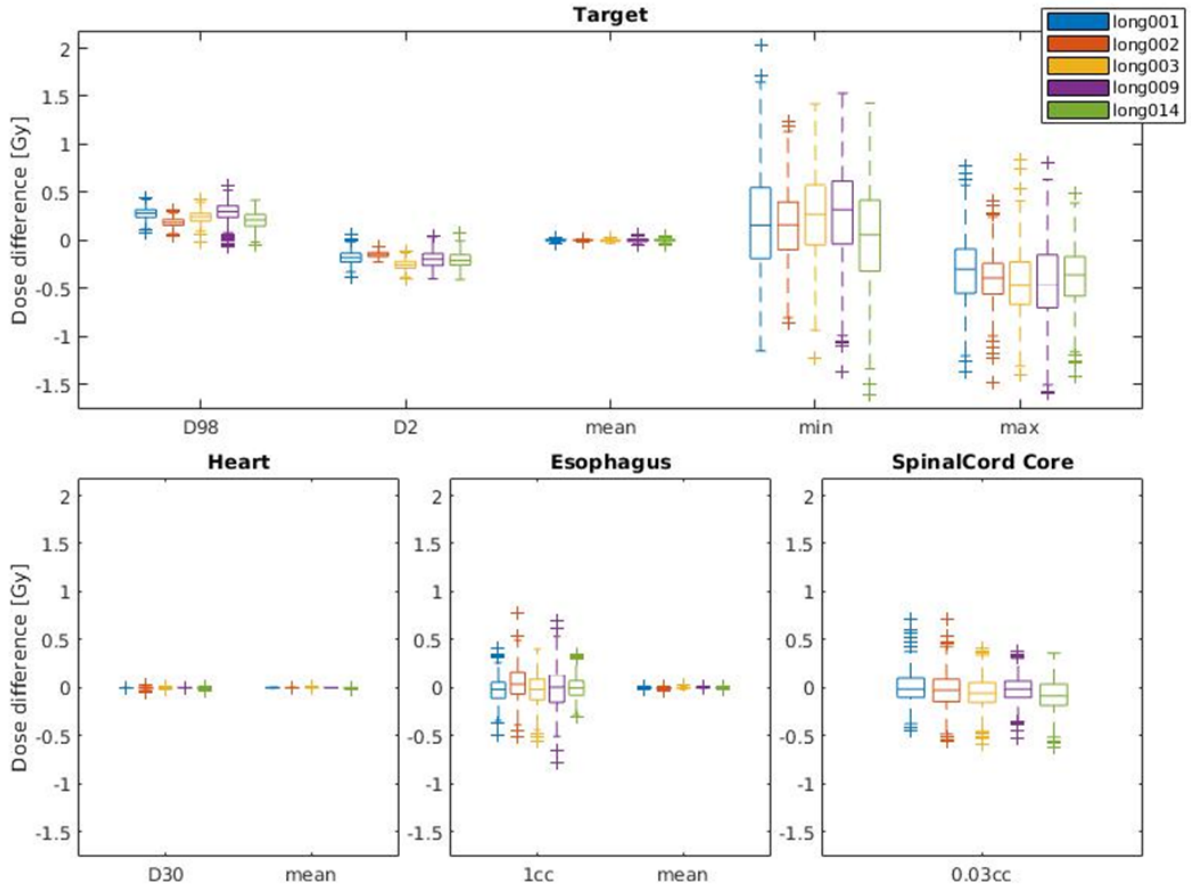

Figure S4: Absolute dose differences for five lung cancer patients between PCE predictions and MC-dose engine calculated dose distributions for 425 uncertainty scenarios. The differences for several target dose DVH-parameters and clinical goal DVH-parameters for important lung OARs are included.

S5. Employed sets of opposing beam angles

Table S7: Original (*Clinical*) and new coplanar gantry angles of the two sets of opposing beam angles used for the *RobustDoseOpt*, *LETopt* and *RobustLETopt* plans.

Patient Alias	Original Gantry angles† [°]	New Gantry angles† [°]	
		Opposing Set 1	Opposing Set 2
Long001	(180,150,210)	(30,210)	(150,330)
Long005	(10,340,205)	(20,200)	(140,320)
Long007	(180,150,210)	(30,210)	(150,330)
Long009	(180,150,210)	(30,210)	(150,330)
long010	(180,150,15,345,210)	(35,215)	(145,325)
long012	(180,350,210)	(30,210)	(150,330)
long015	(30,0,330)	(30,210)	(150,330)
long016	(180,150,210)	(30,210)	(150,330)
long017	(180,150,25)	(25,205)	(150,330)
long018	(170,220,190)	(15,195)	(165,345)

† *Gantry angles as listed here are defined with respect to the patient's coordinate system in the coronal plane. An angle of 0 degrees corresponds to the negative anterior direction, i.e. when the beam is pointing exactly into the patient's front. Angles increase in a clockwise fashion when viewed from the foot of the patient (head-first supine position). For example, 90 degrees corresponds to a lateral beam pointing from left to right, 180 degrees corresponds to a posterior beam (from back to front) and 270 degrees corresponds to a lateral beam pointing from right to left*

S6. Robustness Settings and scaling factors

Table S8: The used robustness settings during optimization and associated scaling factors per patient to reach the probabilistic coverage goal of $p_{10}(D_{99.8,CTV}) = 0.95D_{pr}$, while maintaining within one percentage point of the probabilistic maximum dose goal of $p_{95}(D_{2,CTV}) = 1.4D_{pr}$

Patient Alias	**Robustness Settings**		**D99.8,p10 /Dpr**	**D2,p95 /Dpr**	**Scaling factor**	**D99.8,p10 /Dpr Scaled**	**D2,p95 /Dpr Scaled**
	Range [%]	**Setup [mm]**					
Long001	3	2.5	0,920	1,365	1,033	0,950	1,396
Long005	3	3	0,918	1,360	1,035	0,950	1,392
Long007	3	6.5	0,950	1,405	1,000	0,950	1,405
Long009	3	4	0,941	1,388	1,010	0,950	1,397
long010	3	5.5	0,934	1,387	1,018	0,950	1,403
long012	3	2	0,869	1,311	1,094	0,950	1,392
long015	3	5.75	0,956	1,401	0,993	0,950	1,395
long016	3	4	0,929	1,371	1,023	0,950	1,392
long017	3	4	0,927	1,370	1,025	0,950	1,393
long018	3	3.5	0,927	1,375	1,025	0,950	1,397

S7. Time info

The median time needed for computing the 208 perturbed scenarios using the RayStation MC dose engine was approximately 1020 seconds. The median time needed for predicting the 10.000 fully fractionated treatment plan doses was approximately 420 seconds. The median time for PCE model fitting/creation was 10.3 seconds. This corresponds to an average computation time of a single MC dose calculation of 4.98 seconds, versus a PCE prediction time of just 0.0015 seconds per scenario (when taking the extra fraction scenarios into account that are needed to compute the fully fractionated dose). This shows the comparative speed of more than 3000 (3420.82) times as fast dose computations with PCE once the model is built.

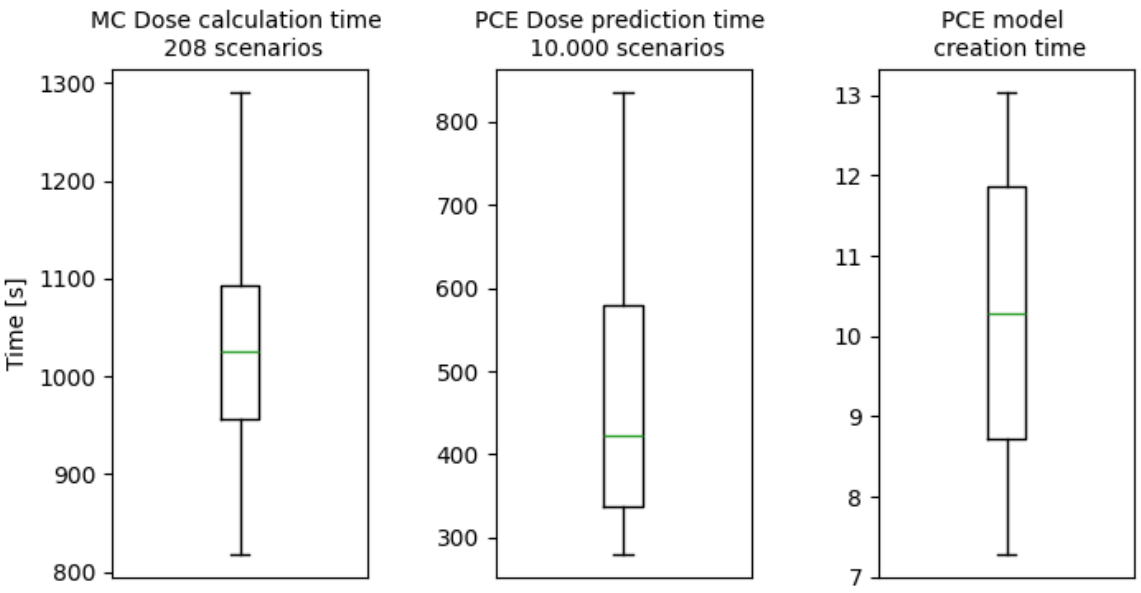

Figure S5: Three panels displaying the time needed for computing the sparse set of 208 scenarios used for fitting the PCE models with the RayStation MC-dose engine, the prediction time of the fully fractionated treatment times using the fit PCE models, and the time needed for fitting the models, for all patients. *Note:* the time for fully fractionated doses using PCE is dependent on the fractionation schemes for the different patients, with more fractions necessitating more scenario calculations and therefore a higher amount of time needed for prediction. See Table S1 for fractionation info.

S8. Beam placement

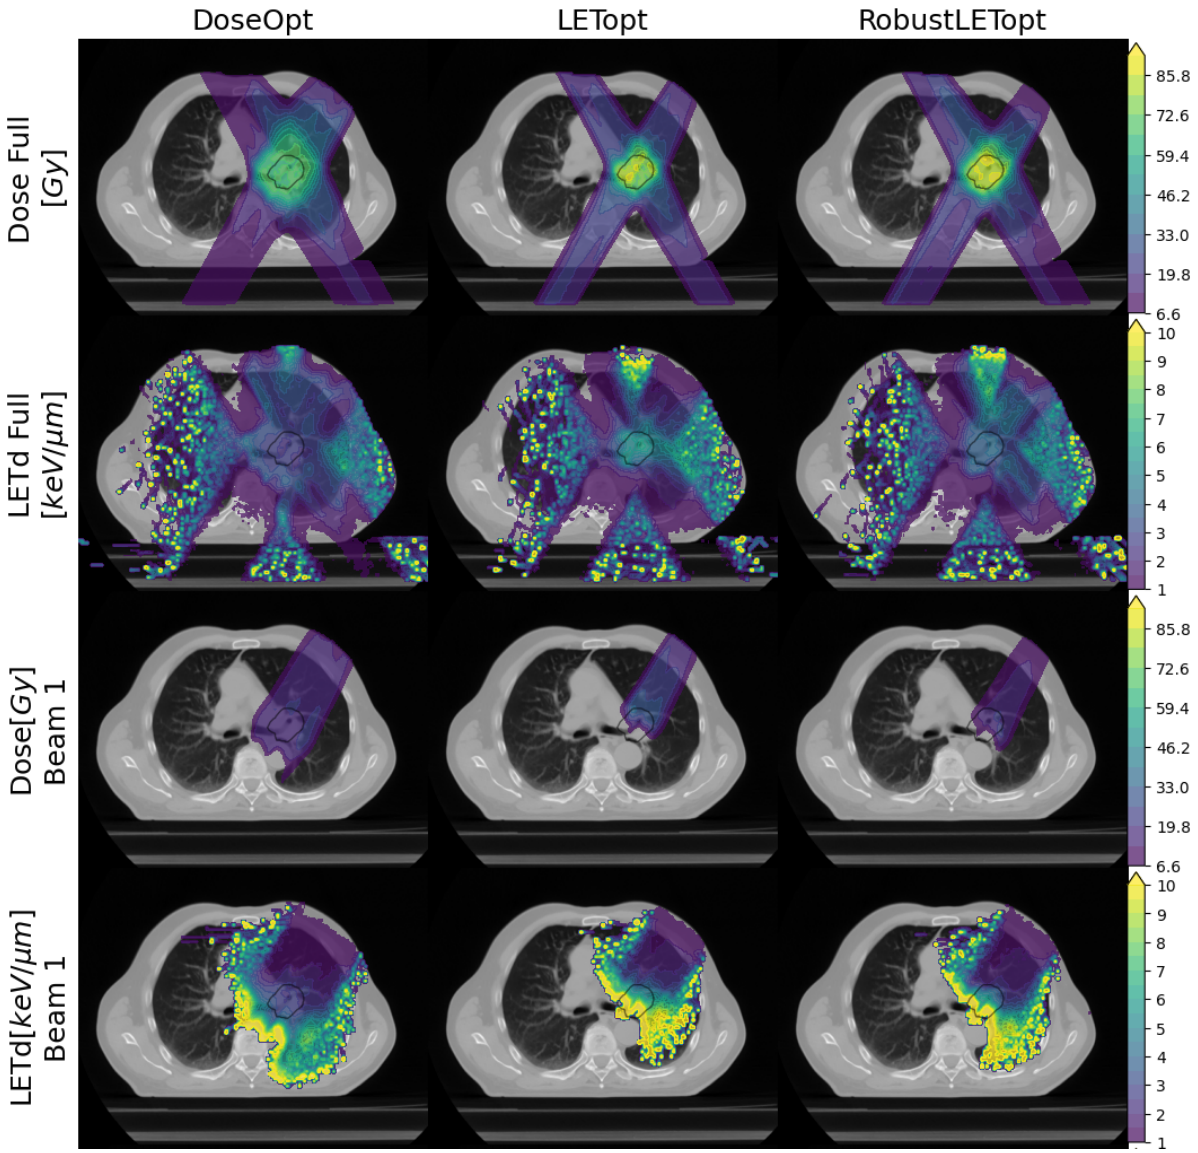

Figure S6: A comparison of the beam-wise dose and LET_d distributions for the reference dose-optimized and LET-optimized treatment plans, showcasing the effect of introducing beam-specific minimum dose objectives for proximal target segments on the placement of spots closer inside or beam-wise proximal to the target (*RobustDoseOpt* vs. *LETopt*). Increases in robustness due to robust optimization again lead to spots being placed more distally, with the *RobustLETopt* plans having spot placement approximately between the placement observed in the *RobustDoseOpt* and *LETopt* plans, reflecting the increased robustness of these plans compared to the *LETopt* plans and decreased robustness compared to the *RobustDoseOpt* plans. As showcased in Figure 5 of the Results section

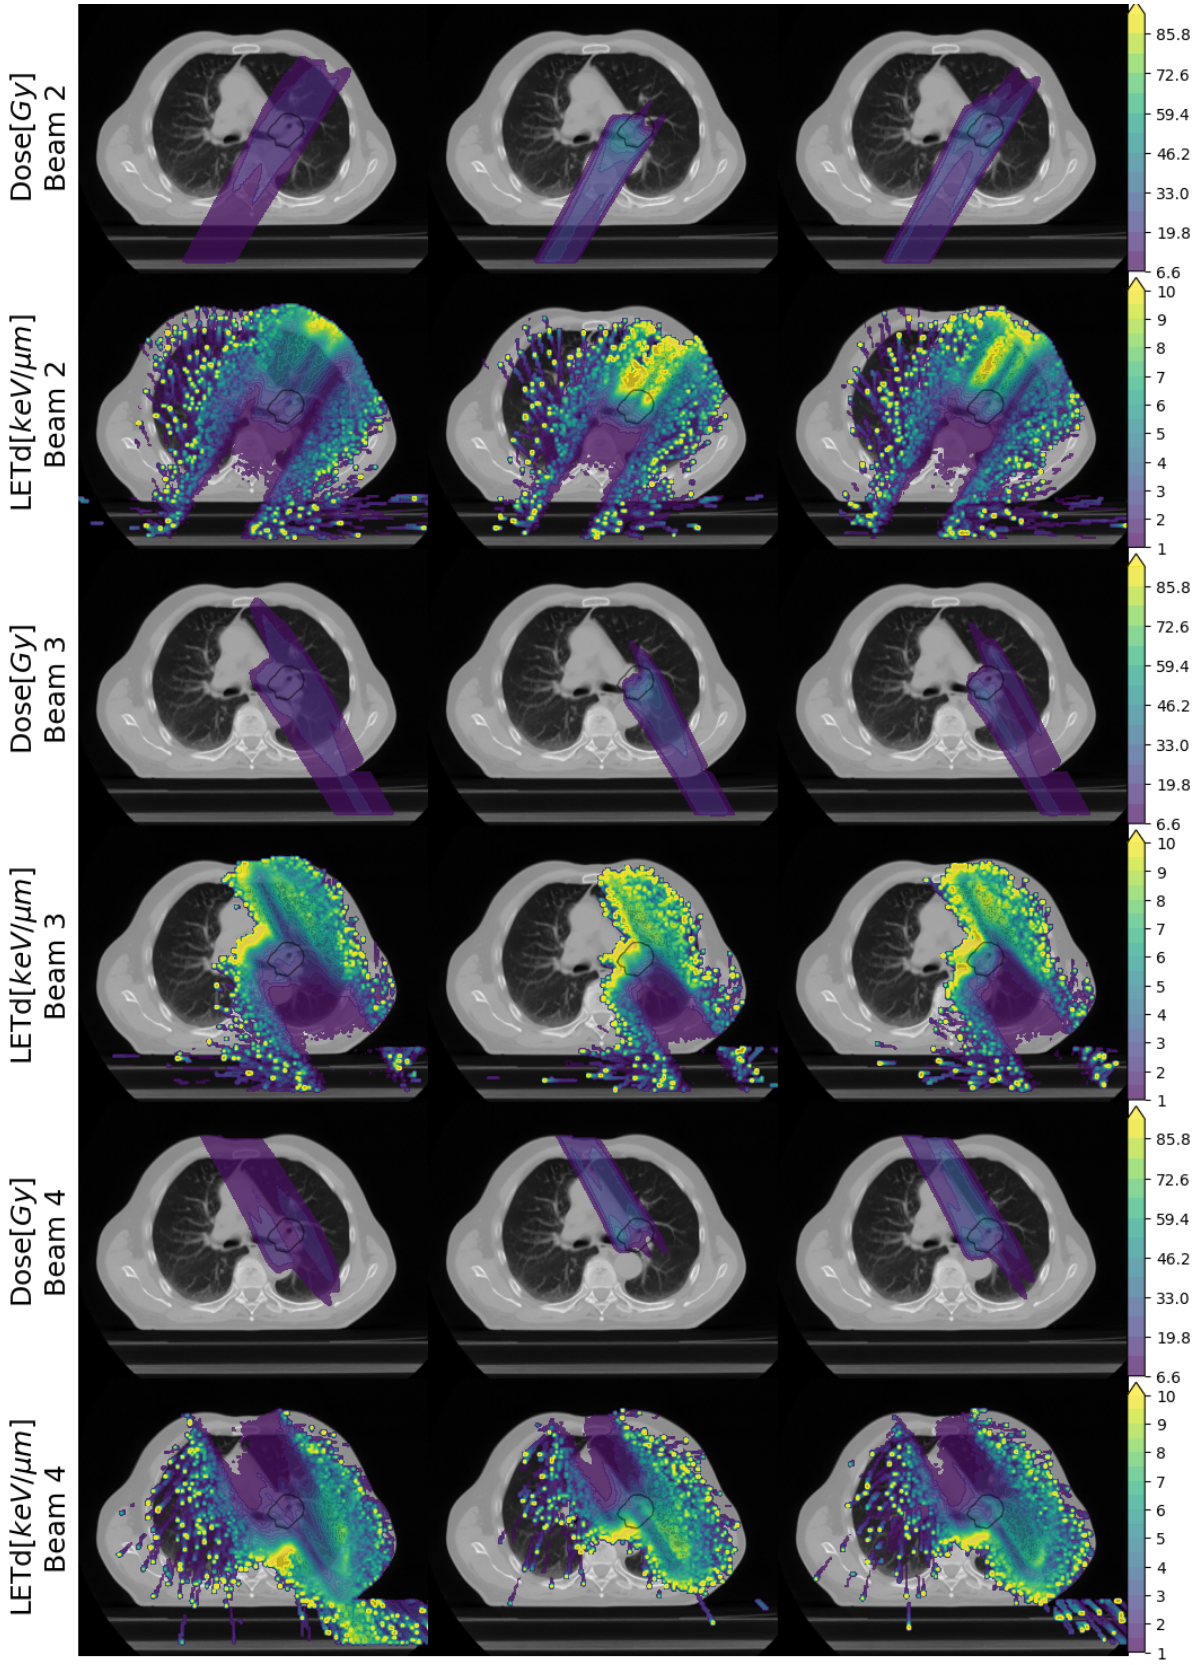

Figure S7: Figure S6 continued.

S9. Adherence to OAR clinical goals: Nominal scenario

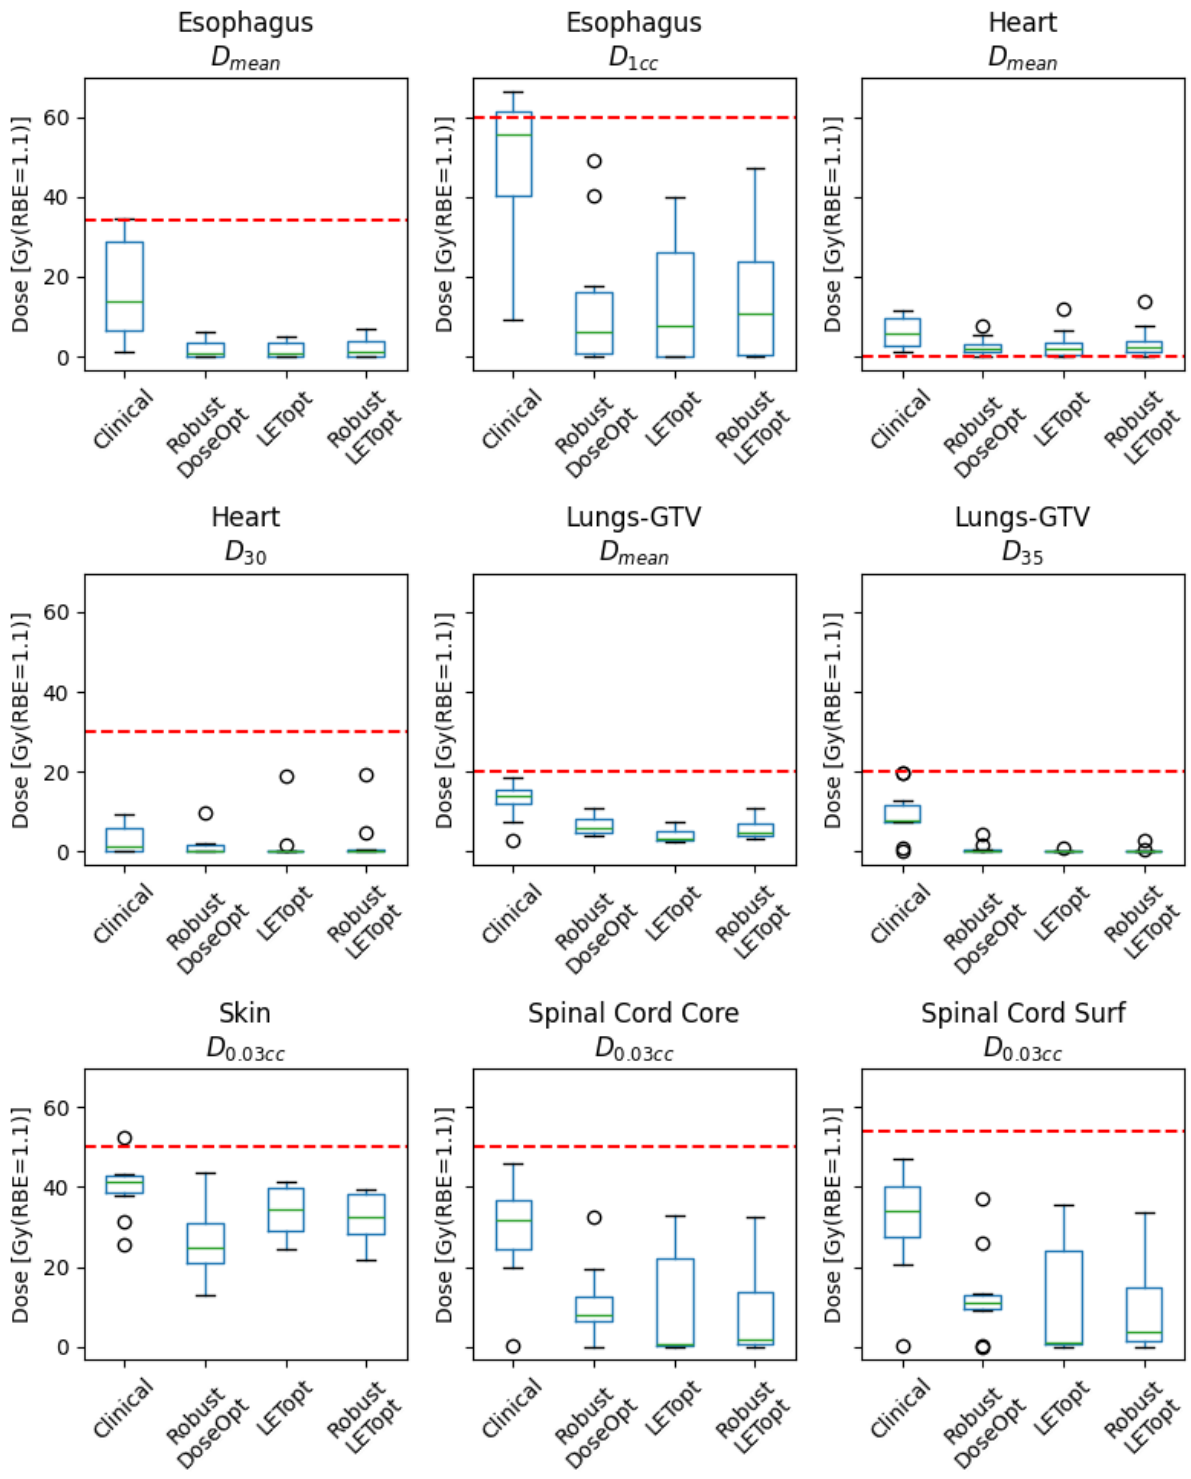

Figure S8: Comparison of the OAR-based clinical goal DVH-parameter values between plan types for all patients. The horizontal dashed lines indicate the clinical goal values, that are not to be exceeded.

S10. Adherence to OAR clinical goals: Probabilistic evaluation of $p_{95}(D_x)$

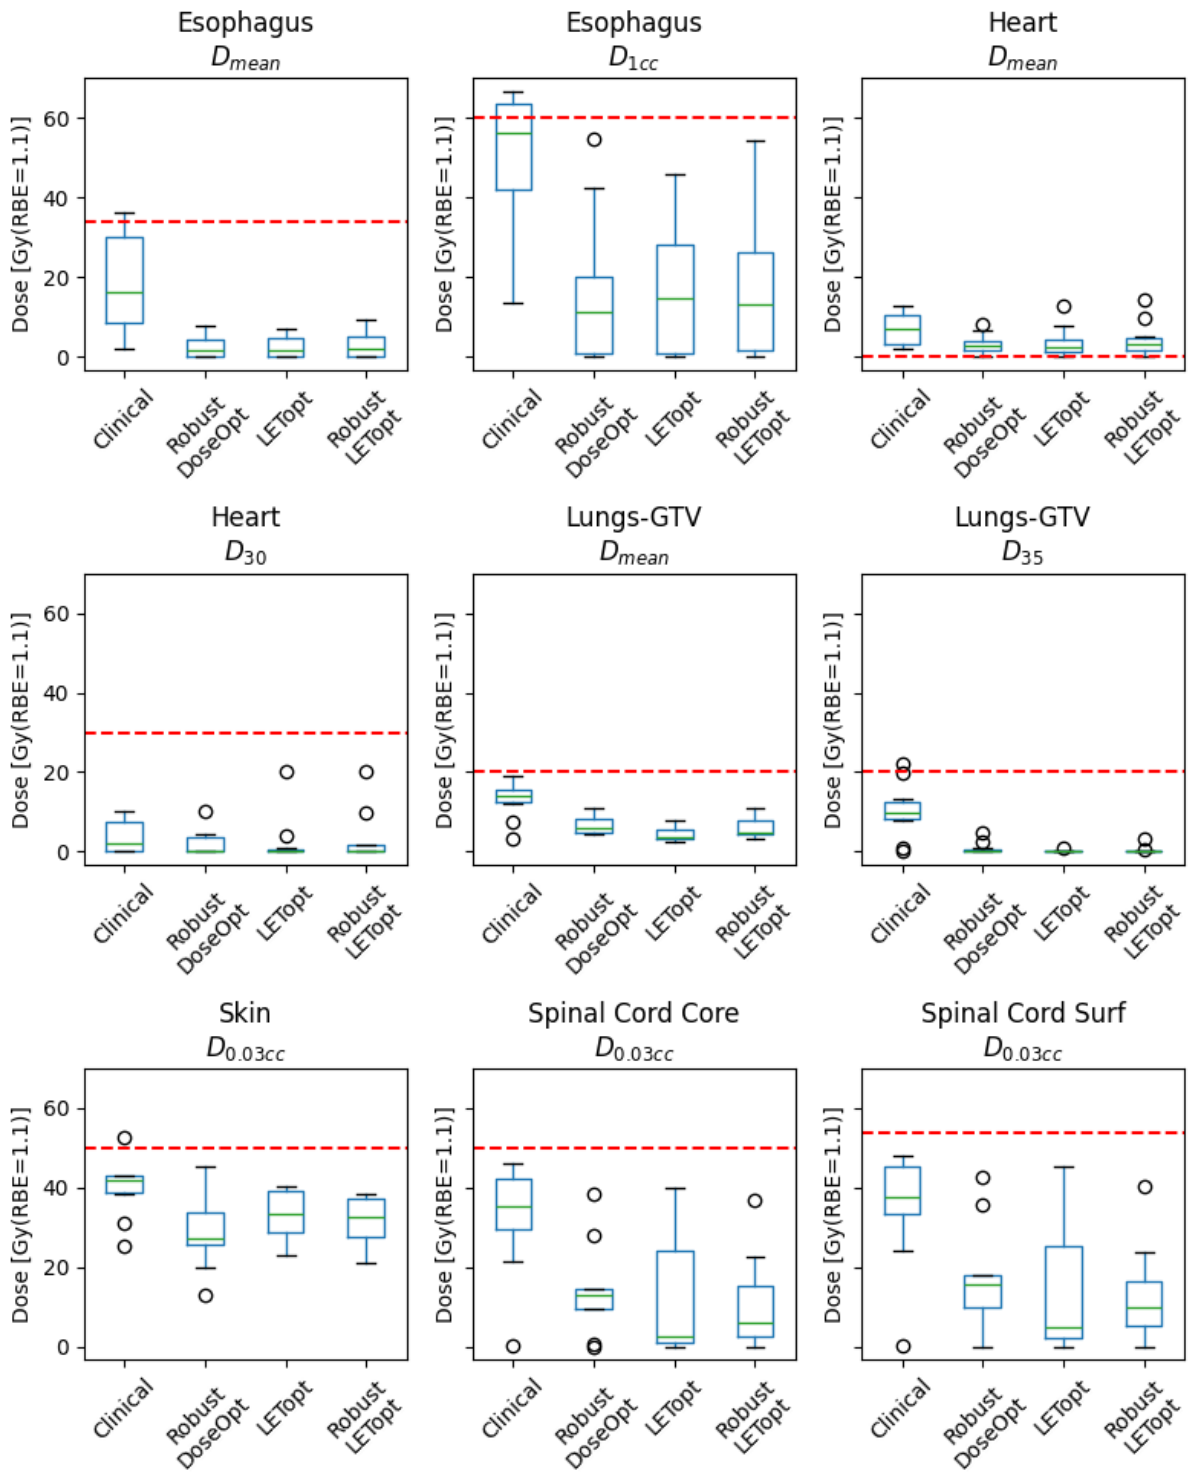

Figure S9: Comparison between plan types of the OAR-based clinical goal values based on the 95th percentile value of associated DVH-metric, based on 10.000 fully fractionated treatments as predicted using PCE. The horizontal dashed lines indicate the clinical goal values, that are not to be exceeded.

S11. $D_{99.8}$ distributions of the four plan types for a single patient

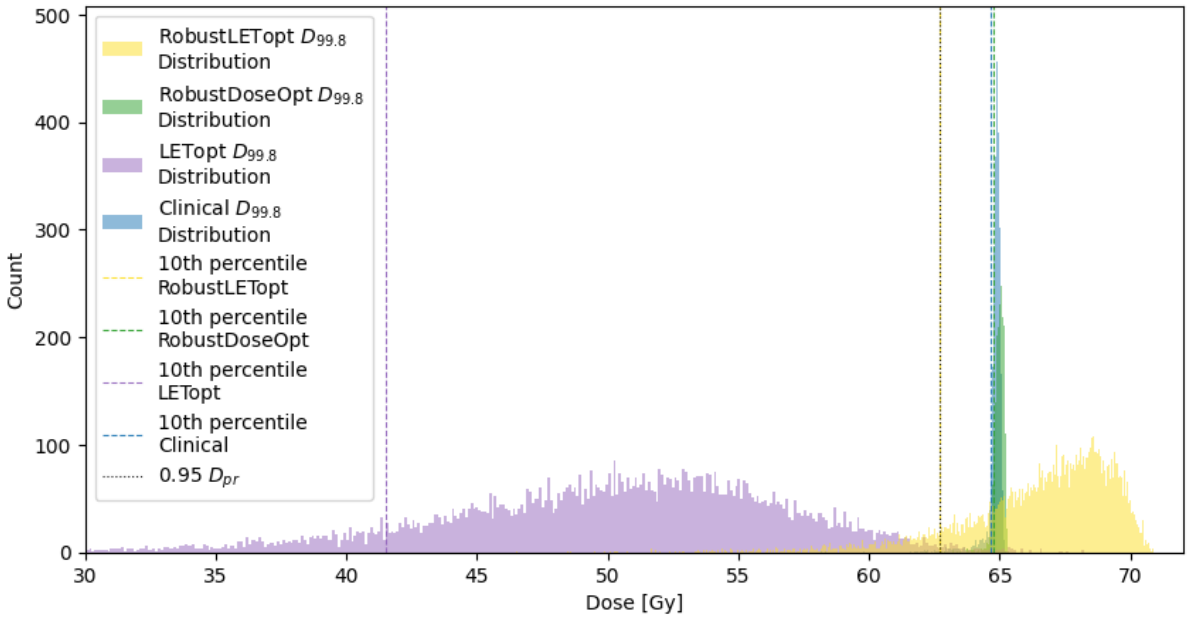

Figure S10: A comparison of the distribution of $D_{99.8}$ values between the dose-optimized (*Clinical, RobustDoseOpt*) and LET-optimized (*LETopt, RobustLETopt*) plans, for 10.000 simulated, fully fractionated treatments for a single patient. Dashed vertical lines indicate the 10th percentile value of the corresponding $D_{99.8}$ same-color distribution for a plan type. A black dashed vertical line indicates the probabilistic coverage goal of $p_{10}(D_{99.8,CTV}) = 0.95D_{pr}$used throughout this study. *Note*: the 10th percentile line of the *RobustLETopt* plan for this patient coincides with the coverage goal line.

S12. TCP curve

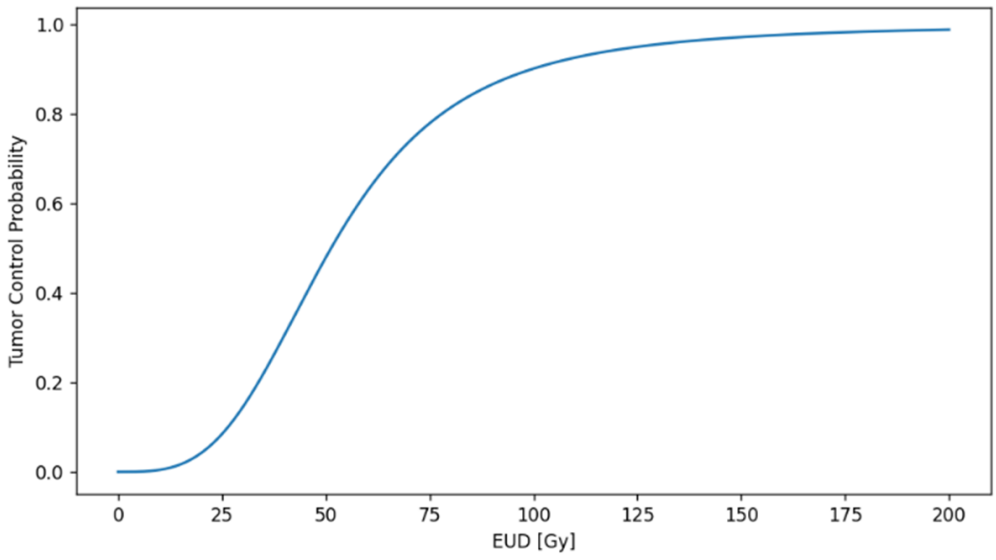

Figure S11: The tumor control probability (TCP) versus equivalent uniform dose (EUD), as calculated by equation 5 in the Methods section.

S13. Correlation between target LET_d and target volume

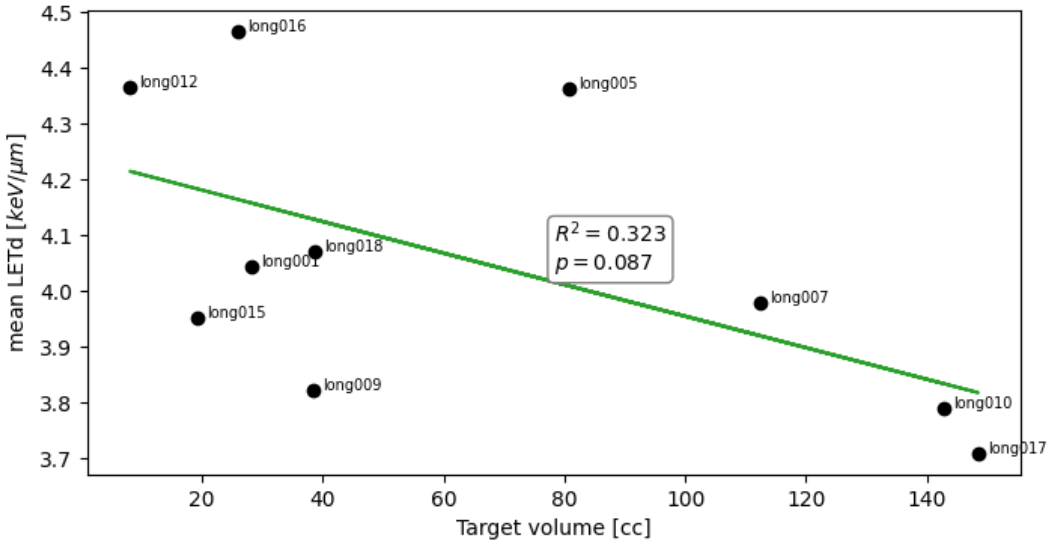

Figure S12: The mean LET_d in the primary CTV versus the target volume in cc, alongside a non-significant linear fit of the data-points.

S14. TCP curve

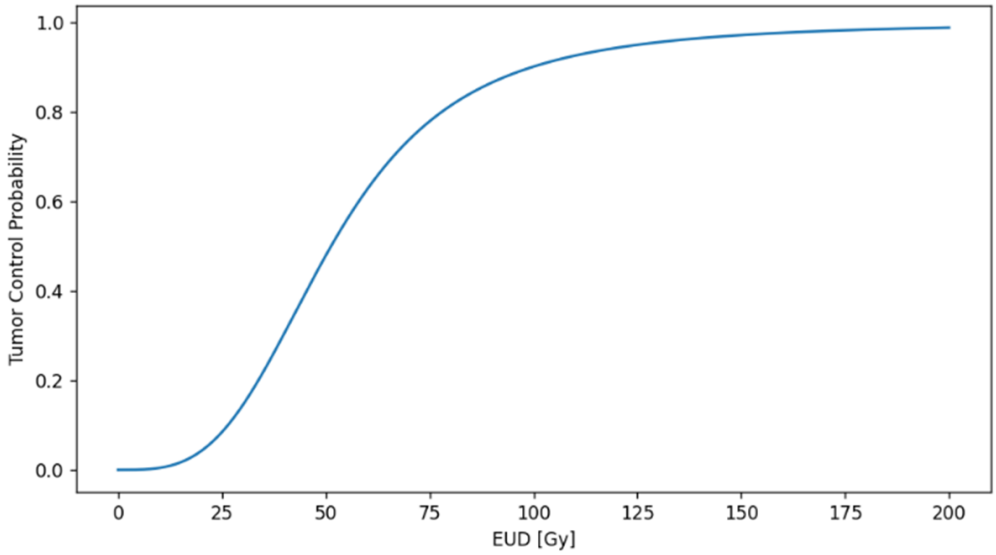

Figure S13: The tumor control probability (TCP) versus equivalent uniform dose (EUD), as calculated by equation 5 in the Methods section.

S15. Robustness of LET_d, $LET_d \times D$ and TCP

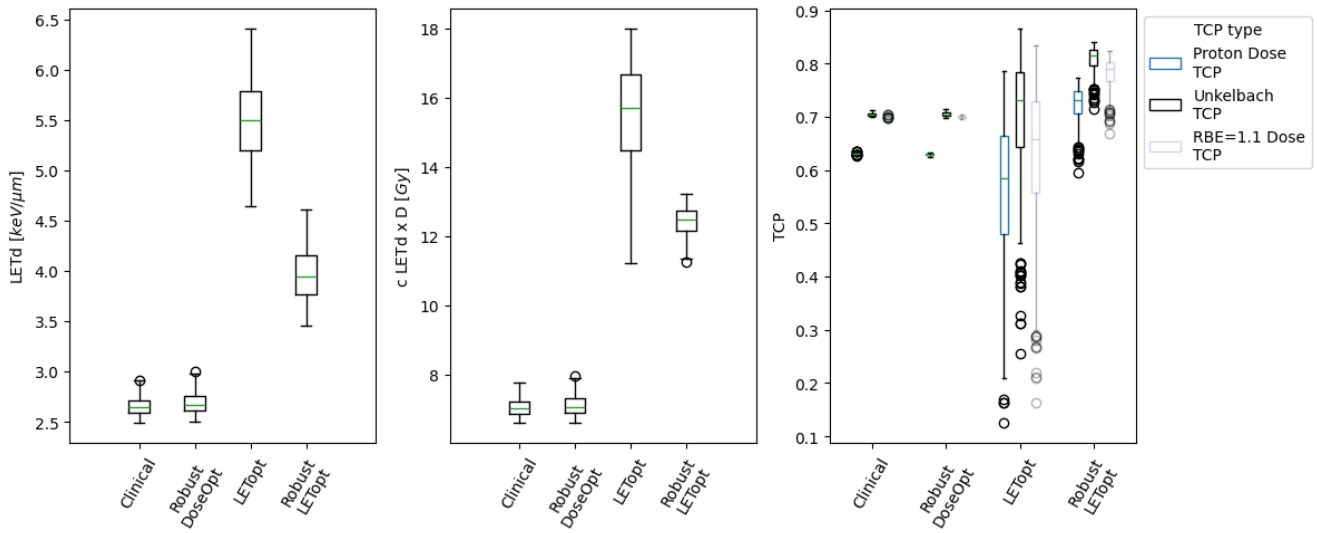

Figure S14: Three panels displaying the spread in MC-engine-based values of target mean LETd, LETdxD and TCP for a single patient and 425 scenarios uncertainty scenarios. These scenarios span the 4D 90% confidence space of the combined setup-errors and relative range errors. These scenarios were the same as used for the initial validation of PCE for lung cancer patients, see Section S4