

RESEARCH ARTICLE

MEDICAL PHYSICS

Direct optimization of the probability of lesion origin in proton treatment planning for low-grade glioma patients

Tim Ortkamp^{1,2,3} | Habiba Sallem^{4,5} | Semi Harrabi^{4,5,6} | Martin Frank^{1,3} |
Oliver Jäkel^{2,3,5,6} | Julia Bauer^{4,5,7} | Niklas Wahl^{2,5}

¹Karlsruhe Institute of Technology (KIT),
Scientific Computing Center,
Eggenstein-Leopoldshafen, Germany

²Department of Medical Physics in Radiation
Oncology, German Cancer Research Center
(DKFZ), Heidelberg, Germany

³Helmholtz Information and Data Science
School for Health (HIDSS4Health),
Karlsruhe/Heidelberg, Germany

⁴Department of Radiation Oncology,
University Hospital Heidelberg, Heidelberg,
Germany

⁵National Center for Radiation Research in
Oncology (NCRO), Heidelberg Institute for
Radiation Oncology (HIRO), Heidelberg,
Germany

⁶Heidelberg Ion Beam Therapy Center (HIT),
Heidelberg, Germany

⁷Department of Radiation Oncology, Charité -
Universitätsmedizin Berlin, Corporate Member
of Freie Universität Berlin and
Humboldt-Universität zu Berlin, Berlin,
Germany

Correspondence

Tim Ortkamp, Scientific Computing Center,
Karlsruhe Institute of Technology (KIT),
Hermann-von-Helmholtz-Platz 1, 76344
Eggenstein-Leopoldshafen, Germany.
Email: tim.ortkamp@kit.edu

Niklas Wahl, Department of Medical Physics
in Radiation Oncology, German Cancer
Research Center (DKFZ), Im Neuenheimer
Feld 280, 69120 Heidelberg, Germany.
Email: n.wahl@dkfz-heidelberg.de

Funding information

Deutsche Forschungsgemeinschaft,
Grant/Award Number: 443188743;
Alois-Hirt-Erben und Wieland Foundation;
Helmholtz Information and Data Science
School for Health

Abstract

Background: In proton therapy of low-grade glioma (LGG) patients, contrast-enhancing brain lesions (CEBLs) on magnetic resonance imaging are considered predictive of late radiation-induced lesions. From the observation that CEBLs tend to concentrate in regions of increased dose-averaged linear energy transfer (LET_d) and proximal to the ventricular system, the probability of lesion origin (POLO) model has been established as a multivariate logistic regression model for the voxel-wise probability prediction of the CEBL origin.

Purpose: To date, leveraging the predictive power of the POLO model for treatment planning relies on hand tuning the dose and LET_d distribution to minimize the resulting probability predictions. In this paper, we therefore propose automated POLO model-based treatment planning by directly integrating POLO calculation and optimization into plan optimization for LGG patients.

Approach: We introduce an extension of the original POLO model including a volumetric correction factor, and a model-based optimization scheme featuring a linear reformulation of the model together with feasible optimization functions based on the predicted POLO values. The developed framework is implemented in the open-source treatment planning toolkit matRad.

Results: Our framework can generate clinically acceptable treatment plans while automatically taking into account outcome predictions from the POLO model. It also supports the definition of customized POLO model-based objective and constraint functions. Optimization results from a sample LGG patient show that the POLO model-based outcome predictions can be minimized under expectable shifts in dose, LET_d , and POLO distributions, while sustaining target coverage ($\Delta_{PTV}D95_{RBE,fx} \approx 0.00$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.03$), even when NTCP is strongly downregulated.

Conclusion: POLO model-based treatment plan optimization for LGG patients can be implemented in a technically feasible way, alleviating the need to hand tune the dose and LET_d distribution. Future work should address multipatient follow-up studies.

KEYWORDS

outcome-based treatment planning, proton therapy, voxel-based outcome analysis

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Medical Physics* published by Wiley Periodicals LLC on behalf of American Association of Physicists in Medicine.

1 | INTRODUCTION

Late radiation-induced contrast-enhancing brain lesions (CEBLs) pose an adverse event after proton therapy of brain tumors and, by tendency, cluster in regions of increased dose-averaged linear energy transfer and near the ventricular system (VS).¹ They are considered predictive of late radiation-induced lesions (LRLs), that is, brain tissue that shows damage from radiation exposure with observable symptoms delayed for months to years after treatment. For the patients, LRLs can—in rare cases—severely impact quality of life through anatomical and functional impairments, such as vascular abnormalities, demyelination, gliosis, white matter necrosis, and cognitive impairment.² Therefore, attention is paid to the prevention of the occurrence of post-treatment LRLs in patients undergoing proton therapy.

Backed by clinical evidence on CEBLs, the probability of lesion origin (POLO) model has been developed,³ a voxel-level model that accurately predicts the localization of CEBLs in low-grade glioma (LGG) patients and allows mapping to a patient-level normal tissue complication probability (NTCP). The POLO model constitutes a multivariate logistic regression function fitted on a cohort of 110 patients with LGG treated with proton therapy at the Heidelberg Ion-Beam Therapy Center (HIT). It incorporates three key risk factors as follows: the voxel-wise total dose d , the dose-averaged linear energy transfer (LET_d), denoted by l_d , and a binary risk factor b measuring the proximity to the cerebral ventricles. Although the POLO model shows reasonably high predictive performance (AUC $\approx$ 0.94, Brier score $\approx$ 2.61×10^{-5}), there is only a limited amount of preceding and follow-up research to date promoting its clinical implementation.^{1,3–8}

Related work includes a series of publications on *voxel-based outcome analysis* (VBA), which takes the full spatial dose distribution as input to detect patterns between local dose and treatment outcome^{9–11} and has already been applied to different sites.^{12–14} VBA establishes as a retrospective tool to perform statistical inference on the dose–response relationship, that is, to generate “significance maps” from voxel-wise hypothesis testing or generalized linear models. However, McWilliam et al.¹¹ even states that “the future direction of VBA is towards optimizing plans directly on the dose response maps.” This is accompanied by developments in the field of *outcome-based treatment planning*,^{15–19} leveraging predictive models from radiobiology and machine learning to directly optimize treatment plans. These models are largely based on extracted dose features rather than the full spatial distribution. Hence, by integrating the POLO model into treatment plan optimization, we position ourselves in the current gap between these two research fields.

Regarding the POLO model, a previous study by Bauer et al.⁵ has already explored the impact of various risk factors, patient characteristics, and treatment plan parameters on the POLO model-based risk prediction, acknowledging the need for “novel biologically guided plan optimization approaches to achieve effective risk minimization in proton therapy for LGG.” As a response, Sallem et al.⁷ developed a risk-minimizing treatment planning optimization strategy, which incorporates a variable LET_d-dependent proton RBE and actively spares the peri-ventricular region. This strategy is currently being applied in the world’s first prospective, randomized phase II trial (ClinicalTrials.gov Identifier: NCT05964569) to evaluate the feasibility of POLO model-based risk minimization in clinical practice. However, this does not yet feature an automated optimization pipeline to directly integrate the POLO model into treatment planning.

In this paper, we leverage the clinical insights gained from Sallem et al.⁷ to address the challenges arising from outcome model integration by introducing a framework for direct optimization of the POLO in intensity-modulated proton therapy planning for LGG patients. To this end, we implemented a forward calculation step to determine the spatial POLO distribution (POLO “map”) together with feasible POLO model-based optimization functions—either based on the original POLO model or on a linear reformulation—and derived the corresponding backward calculation step to yield the fluence gradient. The optimization functions then serve as independent quantities for treatment plan optimization.

We implemented our approach within the open-source treatment planning toolkit matRad,²⁰ using the nonlinear quasi-Newton interior-point methods from IPOPT²¹ to minimize a weighted sum of dose–volume and POLO model-based objectives. While we aim to approach the quality of clinical reference plans, our focus here is on showing the feasibility and overall efficacy of POLO model-based optimization.

The paper is structured as follows: in Section 2, we introduce the theoretical and application-related concepts underlying the integration framework, including a definition of the POLO model, the derived model-based optimization functions, their implementation and finally, some details on the design choices made for treatment planning. Results on the dose/LET_d distributions, POLO maps, as well as treatment outcome predictions are then covered in Section 3. We round off the paper with a discussion and concluding remarks in Sections 4 and 5. For interested readers, Appendix includes mathematical definitions of the dose–volume functions used for conventional plan optimization, the computational costs for all optimization cases, and an example of robust optimization.

2 | MATERIALS AND METHODS

2.1 | Voxel-wise POLO model

The POLO model assigns each voxel from the patient geometry a local risk estimate between 0 and 1, indicating the probability of a voxel being the origin of a CEBL. Its original definition proposed by Bahn et al.³ considers the voxel-wise total dose d and its elementary product with the dose-averaged LET l_d as physical inputs. We further denote the number of voxels by n , that is, $d \in \mathbb{R}_+^n$ and $l_d \in \mathbb{R}_+^n$.

2.1.1 | Definition

Under the assumptions for a multivariate logistic regression, the *probability function* with the linear predictor coefficients taken from Bahn et al.³ reads

$$p(\eta) = \sigma(\eta) = \frac{1}{1 + \exp(-\eta)} \in (0, 1)^n, \\ \eta = -26.3 + \beta_1 \cdot d + \beta_2 \cdot (d \circ l_d) + 1.19 \cdot b, \quad (1)$$

where $\beta_1 = 0.19 \text{ Gy}^{-1}$, $\beta_2 = 0.018 \text{ keV}/(\text{Gy} \cdot \mu\text{m})$, and the binary risk factor

$$b = (\mathbb{1}\{\delta_1 \leq 4\} \quad \dots \quad \mathbb{1}\{\delta_n \leq 4\})^T \in \{0, 1\}^n \quad | \delta_i : \\ \text{distance [mm] of } i\text{th voxel to VS} \quad (2)$$

captures the proximity effect of the local risk in a radius around the VS.

2.1.2 | Extension with volume dependence

One effect that is not taken into account in the above definition is the volume dependence of the local risk estimates. This may be considered in practice because the model parameters from Equation 1 are fitted on a planning CT with a spatial resolution of $0.6 \times 0.6 \times 3.0 \text{ mm}^3$, while treatment planning usually takes place at a lower resolution. We argue that $p(\eta)$ could be adjusted based on the critical element model²² by expanding Equation 1 to

$$p(\eta, k) = 1 - [1 - p(\eta)]^k = 1 - \frac{1}{[1 + \exp(\eta)]^k} \in (0, 1)^n, \quad (3)$$

including the volumetric correction factor

$$k = v_{\text{new}}/v_{\text{old}} \in \mathbb{R}_+ \quad | v : \text{volume [mm}^3\text{]} \text{ of a voxel} \quad (4)$$

which accounts for the rescaling of the local risk in the event of volumetric changes. In Equation 3, we use the point symmetry of the sigmoid function that $1 - p(\eta) = p(-\eta)$, and therefore we describe the volume-corrected probability of lesion origin $p(\eta, k)$ as the counter-probability of the exponentiated probability of no lesion origin $p(-\eta)^k$. For example, if $p(\eta)_i = 0.1$ for the i th voxel, and we halve the spatial resolution, that is, $k = 2$, then the volume-corrected POLO will be $p(\eta, k)_i = 1 - (1 - 0.1)^2 = 0.19$. This effectively reduces biases due to volumetric differences.

2.1.3 | Patient-level NTCP

Under the assumption of a serial tissue response, Bahn et al.³ also derived a patient-level NTCP prediction function, which maps the voxel-wise POLO values $p(\eta, k)_i$ to a risk estimate of developing at least one CEBL:

$$\text{NTCP} = 1 - \prod_{i=1}^n (1 - p(\eta, k)_i) \in (0, 1). \quad (5)$$

2.2 | POLO model-based optimization

Our approach strives to directly optimize intensity-modulated proton treatment plans for LGG patients such that the voxel- and patient-level risk estimates returned by the POLO model are minimized while satisfying conventional dosimetric treatment goals. This necessitates the design of a scalar optimization function $f : D \rightarrow \mathbb{R}$, $D \subseteq \mathbb{R}^n$, and (ideally) the derivation of a corresponding vector-valued gradient function $\nabla_x f : D \rightarrow \mathbb{R}^n$, $x \in D$, to feed into a large-scale nonlinear optimization algorithm, as well as the establishment of a gradient propagation rule from the model to the fluence (decision variable). The function f serves as a generic operator that, when applied to the probability function p , takes the form $f_p : (0, 1)^n \rightarrow \mathbb{R}$ with associated gradient $\nabla_p f_p : (0, 1)^n \rightarrow \mathbb{R}^n$.

Let $\phi \in \mathbb{R}_+^m$ be the fluence vector whose dimensionality equals the number of proton beamlets m . Both f_p and $\nabla_p f_p$ rely on the calculation of the product between d and l_d . In accordance with Unkelbach et al.,²³ we define the i th element of the quantity $(d \circ l_d)$ as follows:

$$d_i \cdot l_{d,i} = d_i \cdot \frac{\sum_j \mathcal{L}_{ij} D_{ij} \phi_j}{\sum_j D_{ij} \phi_j} = d_i \cdot \frac{1}{d_i} \sum_j \underbrace{\mathcal{L}_{ij} D_{ij} \phi_j}_{\mathcal{M}_{ij}} = \sum_j \mathcal{M}_{ij} \phi_j, \quad (6)$$

where i and j index the voxels and proton spots. Equation 6 involves the fluence-independent LET- and dose-influence matrices $\mathcal{L} \in \mathbb{R}_+^{n \times m}$ and $D \in \mathbb{R}_+^{n \times m}$, whose elementwise product $\mathcal{M}$ can be precomputed. While

the forward calculation of f_p is straightforward using Equation 6, there are implications for the backward calculation of $\nabla_\phi f_p$. Before we introduce sample functions for f_p , we therefore like to state a general notation of the fluence gradient.

2.2.1 | Fluence gradient

Using basic calculus, the fluence gradient of a differentiable function f_p reads

$$\begin{aligned}\nabla_\phi f_p &= [\nabla_d \eta \cdot J_d(\phi) + \nabla_{(d \circ l_d)} \eta \cdot J_{(d \circ l_d)}(\phi)]^T \\ (\nabla_\eta p \circ \nabla_p f_p) &= C^T (\nabla_\eta p \circ \nabla_p f_p),\end{aligned}\quad (7)$$

under (i) $\nabla_d \eta = \beta_1$, $\nabla_{(d \circ l_d)} \eta = \beta_2$, (ii) $J_d(\phi) = D$, (iii) as per Equation 6, $J_{(d \circ l_d)}(\phi) = M$, and (iv) $\beta_1 \cdot D + \beta_2 \cdot M =: C \in \mathbb{R}^{n \times m}$. We observe that the right-hand side of Equation 7 only depends on the analytical gradient $\nabla_\eta p$ and the specific structure of $\nabla_p f_p$.

2.2.2 | Linear reformulation

Regardless of the function f_p , integrating the probability function p into proton treatment planning has two potential weak points: the (partial) nonconvexity of the probability function p from Equation 3 in the linear predictor η and the flattening of $\nabla_\eta p$ towards the boundaries of η and k . This implies that gradient-based optimizers may terminate at locally but not globally optimal points and/or be prevented from making sufficient steps towards smaller local risk estimates.

To overcome the disadvantageous properties of the function p , we propose a linear reformulation by exploiting its monotonicity. Since monotonic transformations preserve the order, we can simply invert both the sigmoid function σ and the volumetric correction factor k without changing the minimum point of p , yielding the linear *decision function*

$$\tilde{p}(\eta) = \eta \in [-26.3, \infty), \quad \nabla_\eta \tilde{p} = \mathbf{1}. \quad (8)$$

Accordingly, we define $f_{\tilde{p}} : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\nabla_{\tilde{p}} f_{\tilde{p}} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ on the linearized quantity $\tilde{p}$, and adapt the fluence gradient from Equation 7 to

$$\nabla_\phi f_{\tilde{p}} = C^T \left(\underbrace{\nabla_\eta \tilde{p}}_{=\mathbf{1}} \circ \underbrace{\nabla_{\tilde{p}} f_{\tilde{p}}}_{=\nabla_\eta f_{\tilde{p}}} \right) = C^T \nabla_\eta f_{\tilde{p}} \in \mathbb{R}^m. \quad (9)$$

The latter expression is efficiently evaluable, since C can be precomputed before optimization, that is, $\nabla_\phi f_{\tilde{p}}$ only

depends on $\nabla_\eta f_{\tilde{p}}$. It also follows that the reformulation-based optimization problem is convex if $f_{\tilde{p}}$ is convex in $\tilde{p}$ and the model-free subproblem is convex too.

2.2.3 | Optimization functions

Figure 1 displays the corresponding forward and backward calculation tree. Given the importance of the choice for $f_{p/\tilde{p}}$, we propose feasible candidate functions in the following. For this purpose, we simplify the notation by omitting the arguments in the probability or decision function values, that is, $p_i = p(\eta, k)_i$ and $\tilde{p}_i = \tilde{p}(\eta)_i$.

Seriality-assuming NTCP

We can directly use the patient-level NTCP function from Equation 5 as the optimization function f_p with the probability function p . This function takes on a value close to zero if all p_i are close to zero, and a value close to one if there is at least one p_i close to 1. Differentiation by each p_i yields the gradient function as follows:

$$\nabla_p NTCP_p = \left(\prod_{j \neq 1}^n (1 - p_j) \cdots \prod_{j \neq n}^n (1 - p_j) \right)^T. \quad (10)$$

Optimizing proton treatment plans for LGG patients with $NTCP_p$ has the advantage of acting directly on an interpretable outcome measure, with a degree of sensitivity to single high or many mid-range POLO values.

Log-sum-exp approximation

The probabilistic nature of $NTCP_p$ excludes the substitution of $\tilde{p}$ for p , that is, the convexification of the optimization problem. To resolve this, we apply a smoothed maximum function via the logarithmic sum of exponentials (Log-sum-exp, LSE) approximation. The corresponding optimization function reads as follows:

$$LSE_{\tilde{p}} = \log \left(\sum_{i=1}^n \exp(\tilde{p}_i) \right), \quad (11)$$

where $\max\{\tilde{p}_1, \dots, \tilde{p}_n\}$ could be added to account for numerical under-/overflow. $LSE_{\tilde{p}}$ is convex with gradient function

$$\nabla_{\tilde{p}} LSE_{\tilde{p}} = \left(\frac{\exp(\tilde{p}_1)}{\sum_{j=1}^n \exp(\tilde{p}_j)} \cdots \frac{\exp(\tilde{p}_n)}{\sum_{j=1}^n \exp(\tilde{p}_j)} \right)^T. \quad (12)$$

Hellinger distance

By interpreting the probability function values p_i as realizations of random variables, one can also adopt statistical distance measures to design optimization functions f_p . The Hellinger distance is such an example, where the distance between the observed values




FIGURE 1 Forward and backward calculation tree for POLO model-based optimization functions. The black solid arrows symbolize forward value passes, while the violet dashed arrows indicate backward gradient passes. Our linear reformulation bypasses the probability function p and builds a direct path connection from η to $f_{p/\tilde{p}}$ (purple arrows). We denote the fluence ϕ , dose d , dose-averaged LET l_d , binary risk factor b , linear predictor η , volume correction factor k , POLO p , optimization function $f_{p/\tilde{p}}$, model coefficients β_1 and β_2 , dose-influence matrix $\mathcal{D}$, and the “LET-times-dose-influence matrix” $\mathcal{M}$.

$(p_1 \cdots p_n)^T$ and the reference values $(q_1 \cdots q_n)^T$ is measured by the following:

$$H_p = \frac{1}{\sqrt{2}} \sqrt{\sum_{i=1}^n (\sqrt{p_i} - \sqrt{q_i})^2} \in \left(0, \sqrt{\frac{n}{2}}\right). \quad (13)$$

Intuitively, the reference distribution for POLO model-based optimization should be $q_i = 0$ for all $i \in \{1, \dots, n\}$, that is, should be the single-point zero distribution. By inserting $q_i = 0$, removing the constant factor $\frac{1}{\sqrt{2}}$, and inverting the root function due to monotonicity, we obtain the transformed (and convex) optimization function

$$\tilde{H}_p := \sum_{i=1}^n p_i \in (0, n), \quad (14)$$

which yields the constant gradient

$$\nabla_p \tilde{H}_p = \mathbf{1}. \quad (15)$$

We note that above gradient does not depend on the scaling of p_i , that is, we could interchangeably use the decision function values $\tilde{p}_i$ in Equation 14, even though the function $\tilde{H}_p$ would no longer be predicated on a probabilistic measure. The Hellinger distance can generally be thought of as a mean value reduction scheme, with less sensitivity on single high input function values than $NTCP_p$ and $LSE_{\tilde{p}}$.

2.3 | Implementation

To show the applicability of our approach towards treatment plan optimization, we implemented the concepts introduced in Sections 2.1 and 2.2 in matRad (v2.10.1, <http://www.matRad.org>),^{20,24} a MATLAB-based open-

source software toolkit for radiotherapy treatment planning in education and research. This toolkit supports patient loading, dose/LET_d calculation, and optimization for intensity-modulated particle therapy, as well as plan evaluation and visualization.

For dose calculation, we use matRad’s pencil-beam algorithm with voxel raytracing^{25,26} to compute the dose-influence matrix $\mathcal{D}$. Specifically, for protons, the voxel dose is determined as the product of a tabulated integrated depth-dependent property with a tabulated lateral distribution (forming the dose kernel), where the lateral distribution is convolved with the impinging lateral fluence distribution. The product between dose and dose-averaged LET is calculated according to Equation 6, where matRad simplifies the computation by internally operating on the local product of LET and dose contributions (from the kernel database) and then directly mapping onto the quantity $(d \circ l_d)$ via the fluence ϕ .

Our implementation of the POLO model-based optimization framework includes new classes for model handling and the interface to the optimization algorithm. First, we added a “virtual” POLO segmentation over all voxels inside the whole brain and outside the gross tumor volume (GTV) and the 4-mm ventricular fringe, as well as precomputed b from Equation 2. Second, we declared the POLO model from Equation 1 and its linear reformulation from Equation 8 as subclasses of matRad’s *BackProjection*, which facilitates iteration-wise forward calculation and backward differentiation up to the optimization function, according to Figure 1, and inserted an index filter to calculate the POLO prediction values only in voxels where the physical dose is at least 2.0 Gy. Third, we extended matRad’s objective function catalogue with $NTCP_p$, $LSE_{\tilde{p}}$, $\tilde{H}_p$, and $\tilde{H}_{\tilde{p}}$. It is then sufficient to import the LGG test patient data, assign a POLO model-based optimization function to the POLO segment and start the treatment plan optimization process.

2.4 | Treatment planning

This section introduces the LGG sample patient case, the baseline treatment plan and the optimization scenarios covered later. For the remainder, the physical proton dose will be given in units of Gy, and the biologically weighted proton dose in units of Gy(RBE). The displayed dose, dose-averaged LET, and POLO distributions are consistently interpolated back to the treatment planning CT resolution ($0.6055 \times 0.6055 \times 1.0 \text{ mm}^3$) for visualization. Volume contours are drawn thicker than the isolines referring to: 50, 90, 100, and 110 % of the prescribed fractional dose to the CTV/PTV; 4, 5, 8, and 10 keV/ μm for the dose-averaged LET; and 10, 25, 50, and 80 % of the maximum observed POLO.

2.4.1 | Patient

We selected a sample patient case from the cohort described by Sallem et al.,⁷ who received proton therapy for LGG treatment in 2022. Particle radiation was delivered using three beams, with a total dose of 54 Gy(RBE) at constant RBE of 1.1, applied in 30 fractions at the Heidelberg Ion Therapy Center (HIT), Germany. The clinical dose distribution was calculated using the Monte Carlo dose engine in RayStation 11B IonPG (RaySearch Laboratories, Stockholm, Sweden), configured with a dose uncertainty of 0.5 % and a spatial dose resolution of $2.0 \times 2.0 \times 2.0 \text{ mm}^3$. GTV and cerebral ventricles were manually delineated by a radiation oncologist on a T2-weighted MRI scan, while the remaining structures were contoured on the planning CT. In this case, the tumor volume was localized in the right temporo-insular region, with visible overlap between the PTV and the cerebral ventricles.

2.4.2 | Baseline plan

We configured the baseline proton plan for the LGG sample patient in matRad with three couch angles of 175° , 240° , and 310° at fixed gantry angles. The number of fractions was set to 30. Further, we determined the spatial dose grid resolution to be $3.0 \times 3.0 \times 3.0 \text{ mm}^3$ and applied a constant RBE scheme of 1.1. For the optimization step, we compiled the clinical target functions from a template for the sample patient, generated by clinical experts, into a near-equivalent matRad treatment plan. Table 1 summarizes the dose objectives and Table 2 summarizes the dose constraints selected for the plan.

The underlying inverse planning problem is inherently multicriteria, and we handle the nature of this problem by describing the total objective function f as a weighted sum of the partial dose objectives f_r , that is, the inverse

planning problem reads as follows:

$$\min_{\phi} f(\phi) = \sum_r w_r f_r(D_r \phi), \text{ s.t. } c_s^l \leq c_s(D_s \phi) \leq c_s^u \quad \forall s, \quad (16)$$

where w_r denotes the weight of the objective associated with the r th segmented structure, D_r and D_s are the submatrices of the dose-influence matrix D row-filtered for the respective structure, and c_s describes the constraint function for the s th structure with the lower and upper bounds c_s^l and c_s^u . To solve the inverse planning problem, we use the large-scale nonlinear interior-point quasi-Newton IPOPT solver with the *ma57* algorithm.^{21,27}

Figure 2 shows the optimal slice images of the dose (left), dose-averaged LET (middle) and POLO distribution (right). On the underlying CT scan, the tumor volumes are marked by violet to purple contours, while the VS (ventricles and 4-mm fringe) is outlined in white and gray. We observe the high-dose region covering all of the GTV, with a largely homogeneous transition to the clinical and planning target volumes (CTV, PTV), and a moderate to severe overlap to the VS. The dose-averaged LET distributes in layers, with increased values around the contour edge of CTV/PTV and local peaks at the distal ends of the proton beams. Accordingly, hot spots on the POLO map are found in regions where dose and dose-averaged LET scale up, or close to the VS. Beyond the PTV, the POLO map reveals cold spots, due to the steepness of the dose and dose-averaged LET gradients.

2.4.3 | Optimization scenarios

Our analysis on the feasibility and overall efficacy of POLO model-based treatment planning builds on a number of optimization scenarios. Here, we incorporated the POLO model-based optimization functions from Section 2.2 one by one as additional objectives in the weighted sum term from Equation 16 while retaining the structure of the baseline proton plan, and manually tuned the weights associated with the POLO model-based objectives to generate treatment plans at near-equivalent NTCP levels of 29.7 % – 30.5 %, 19.6 % – 20.3 %, and 9.8 % – 10.2 % (calculated with $NTCP_p$). In total, this yielded 12 scenarios (four objective functions, three risk levels), which allow to compare the optimal slice images and histograms of the RBE-weighted fractional dose $d_{RBE,fx}$, the dose-averaged linear energy transfer I_d and the probability function values p_i , as well as to evaluate the outcome prediction graphs by iterating the POLO maps and feeding them to $NTCP_p$.

TABLE 1 Dose objectives for the baseline plan. The function names correspond to the dose objectives implemented in matRad. Appendix provides a mathematical definition of the objective functions used.

Segment	Objective function f_r	Objective weight w_r	Dose parameters
Clinical target volume (CTV)	Squared deviation	100	54 Gy(RBE) ref.
	Dose uniformity	100	—
Planning target volume (PTV)	Squared deviation	100	54 Gy(RBE) ref.
	Dose uniformity	100	—
	Minimum DVH	250	$\geq 97\%$ vol. at 52 Gy(RBE)
PTV–CTV	Squared underdosing	250	52.5 Gy(RBE) min.
	Maximum DVH	250	≤ 56 Gy(RBE) at 0 % vol.
3 mm fringe in PTV	Maximum DVH	25	≤ 50.25 Gy(RBE) at 0 % vol.
Ventricular fringe in PTV–GTV	Squared underdosing	150	49 Gy(RBE) min.
	Maximum DVH	250	≤ 52.3 Gy(RBE) at 0 % vol.
	Squared deviation	250	51.3 Gy(RBE) ref.
	Dose uniformity	250	—
Cochlea (right)	Maximum DVH	100	≤ 45 Gy(RBE) at 0 % vol.
Inner ear (right)	Maximum DVH	150	≤ 44 Gy(RBE) at 0 % vol.

TABLE 2 Dose constraints for the baseline plan. The function names correspond to the dose constraints implemented in matRad. Appendix provides a mathematical definition of the constraint functions used.

Segment	Constraint function c_s	Dose parameters c_s^l / c_s^u
Optical nerve (left)	Minimum/Maximum dose	0 Gy(RBE) min. / 52 Gy(RBE) max.
Optical nerve (right)	Minimum/Maximum dose	0 Gy(RBE) min. / 52 Gy(RBE) max.
Chiasm	Minimum/Maximum dose	0 Gy(RBE) min. / 52 Gy(RBE) max.
Brainstem	Minimum/Maximum dose	0 Gy(RBE) min. / 53 Gy(RBE) max.
External contour	Minimum/Maximum dose	0 Gy(RBE) min. / 56 Gy(RBE) max.
Planning target volume (PTV)	Minimum/Maximum dose	0 Gy(RBE) min. / 59.4 Gy(RBE) max.
Optical system + 2 mm	Minimum/Maximum dose	0 Gy(RBE) min. / 52 Gy(RBE) max.




FIGURE 2 Optimal 2D slice images of the RBE-weighted fractional dose $d_{RBE,fx}$, the dose-averaged linear energy transfer l_d , and the probability function values p_i for the baseline plan. The dose distribution (left) shows high target coverage with decreasing exposure beyond the CTV/PTV margins, but also visibly overlaps with parts of the VS. Meanwhile, for the dose-averaged LET (middle), layered homogeneity in the target volumes and local maxima at the distal ends of the proton beams can be observed. This results in more pronounced levels of POLO values within the target volumes, particularly in the presence of local dose-averaged LET maxima and in the environment of the VS. Referring back to the POLO model from Equations 1 and 3, we can expect exactly this type of coherence between the model output, the physical input feature profiles and the increased local risk near the VS. The baseline treatment plan exhibits a complication risk of 37.6 % after evaluation of the POLO distribution with $NTCP_p$.

3 | RESULTS

3.1 | Dose, LET_d , and POLO distributions

In this section, we present the optimal slice images and histograms for the RBE-weighted fractional dose $d_{RBE,fx}$, the dose-averaged linear energy transfer l_d and the probability of lesion origin p , as well as their interplay.

3.1.1 | RBE-weighted fractional dose $d_{RBE,fx}$

Figure 3 shows the slice images with the optimal dose distributions, where the POLO model-based objective functions are displayed vertically and the NTCP levels horizontally. Throughout this part, we measure the loss in target coverage with respect to the baseline plan by the difference in RBE-weighted fractional doses received by 95 % of the GTV and the PTV, denoted by $\Delta_{GTV}D95_{RBE,fx}$ and $\Delta_{PTV}D95_{RBE,fx}$ with units of Gy(RBE).

We observe intrafunctional similarities of the objectives from a reduction of the dose values with decreasing NTCP level, especially in the overlapping regions to the VS, but also within the target volumes. $NTCP_p$ manages to reduce the risk estimate by almost 18 p.p compared to the baseline plan without affecting target coverage ($\Delta_{PTV}D95_{RBE,fx} \approx 0.01$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.00$), which can be attributed to the dose contraction away from the VS. This is even more pronounced when lowering the NTCP level to 10 %, accompanied again by a preservation of target coverage ($\Delta_{PTV}D95_{RBE,fx} \approx 0.00$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.03$).

From an interfunctional perspective $\tilde{H}_p$ seems to behave very similarly, with only minor deviations in the dose slice. In contrast, we find that $LSE_{\tilde{p}}$ and $\tilde{H}_{\tilde{p}}$ modulate the dose more strongly, with visible local inhomogeneities from the NTCP level of 30 %. Here, $LSE_{\tilde{p}}$ only marginally reduces the dose in the lower part of the VS but restricts the dose exposure in the target volume ($\Delta_{PTV}D95_{RBE,fx} \approx 0.37$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.09$). Even more, target coverage breaks down towards an NTCP level of 10 % ($\Delta_{PTV}D95_{RBE,fx} \approx 0.6$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.23$). One reason could be the focus of $LSE_{\tilde{p}}$ on maximum points, causing “notches” in the high dose region around the overlap of CTV/PTV and VS from the NTCP level of 20 %.

$\tilde{H}_p$, on the other hand, results in an immediate dose contraction, with higher concentration on the GTV with lower NTCP level. Similar to $LSE_{\tilde{p}}$, target coverage is drastically affected at 10 % NTCP ($\Delta_{PTV}D95_{RBE,fx} \approx 0.72$, $\Delta_{GTV}D95_{RBE,fx} \approx 0.19$), which may be due to the increased scale of the function values and gradients of the linearly reformulated POLO model and therefore the higher impact of $LSE_{\tilde{p}}$ and $\tilde{H}_{\tilde{p}}$ on the plan.

Above results are underlined by Figure 4 which displays the corresponding dose–volume histograms (DVHs) and selected dosimetric values in each NTCP interval for the volumes of interest (VOIs) to allow for a global analysis of the dose distributions and their progression.

3.1.2 | Dose-averaged linear energy transfer l_d

In light of the dose distribution changes and the interdependency between dose and dose-averaged LET, we can identify patterns in the slice images with the optimal dose-averaged LET distributions shown in Figure 5.

First, we note that l_d hotspots do not manifest in the high dose regions—this would be a risk driver because of the product term in the POLO model from Section 2.1. Generally, some trade-off between $d_{RBE,fx}$ and l_d is evident, where the optimizer allows for higher values of l_d when $d_{RBE,fx}$ has been sufficiently reduced, and vice versa.

For $NTCP_p$, the slight reduction of $d_{RBE,fx}$ outside the target volume coincides with a redistribution of l_d : higher values in the target volume are prevented, and instead induced along the margins of the PTV. Once again, $\tilde{H}_p$ behaves almost identically. Following the dose/ LET_d trade-off, the almost unchanged values for $LSE_{\tilde{p}}$ can also be explained: $LSE_{\tilde{p}}$ minimally reduces $d_{RBE,fx}$ in the region of overlap with the VS, meaning that high l_d values there would be detrimental to the outcome prediction. Consequently, l_d could only approach the “dose gap” that opens up within the target volume, or remain constant while the shifts in the dose distribution account for the reduction in NTCP.

Finally, for $\tilde{H}_{\tilde{p}}$, we observe a sharp envelopment of the lower side of the PTV margin, consistent with the fall-off in $d_{RBE,fx}$ there. At an NTCP level of 10 %, the high l_d region has virtually manifested itself with hot spots partially overlapping the VS. However, revisiting the slice image for $d_{RBE,fx}$, this is not surprising, as in addition to the lower values in the target volume, the dissipation of $d_{RBE,fx}$ to the outside is only enhanced towards the upper end, leaving space for the l_d band on the lower end, between the PTV margin and the VS.

Again, we can confirm the slice images by the LET_d -volume histograms (LET -VHs) from Figure 6. For $NTCP_p$ and $\tilde{H}_p$, the values decrease with lower NTCP for all VOIs. $LSE_{\tilde{p}}$ almost does not change the distributions, only a slight increase at the left margin of the target volume is visible both in the slice images and in the respective LET -VH curves. Meanwhile, $\tilde{H}_{\tilde{p}}$ reduces the values in all VOIs, just as $NTCP_p$ and $\tilde{H}_p$, but in contrast, for the VS, the lower range is more attenuated and the upper range is enriched due to the redistribution along the PTV margin.





FIGURE 3 Optimal slice images of the RBE-weighted fractional dose $d_{RBE,fx}$ (in Gy(RBE)) for $NTCP_p$, LSE_p , H_p , and $\bar{H}_p$ at different NTCP levels. The objective functions behave similarly in that dose is reduced to achieve a lower NTCP value. However, they differ in the strength of the dose reduction inside and the impact on the dose outside the target volume, particularly in the VS. $NTCP_p$ and $\bar{H}_p$ show comparable dose results, while LSE_p and H_p more aggressively eliminate dose in the target volume.

3.1.3 | Probability of lesion origin p

Figure 7 shows the POLO maps resulting from voxel-wise evaluation of the sigmoid-transformed linear pre-

dictor η from Equation 1 with the input slice images for $d_{RBE,fx}$ and I_d .

For $NTCP_p$ and $\bar{H}_p$, we observed a (slight) reduction of $d_{RBE,fx}$ in the target volume and in the lower region



Opt. case	Segment	ventricles		GTV		CTV		PTV		VS 4mm	
		D_{mean}	D_{max}	D_5	D_{95}	D_5	D_{95}	D_5	D_{95}	D_{mean}	D_{max}
Baseline ($NTCP = 37.6\%$)		0.74	1.77	1.80	1.78	1.80	1.72	1.80	1.71	0.67	1.81
+ $NTCP_p$ (30.2%)		0.72	1.77	1.80	1.78	1.80	1.72	1.80	1.71	0.64	1.81
+ $NTCP_p$ (20.0%)		0.67	1.76	1.80	1.78	1.80	1.72	1.80	1.70	0.60	1.79
+ $NTCP_p$ (10.2%)		0.61	1.78	1.80	1.75	1.79	1.71	1.79	1.71	0.56	1.78
+ $LSE_{\bar{p}}$ (30.5%)		0.66	1.84	1.84	1.69	1.83	1.61	1.83	1.34	0.58	1.84
+ $LSE_{\bar{p}}$ (20.3%)		0.63	1.81	1.82	1.66	1.77	1.54	1.77	1.25	0.56	1.81
+ $LSE_{\bar{p}}$ (9.8%)		0.57	1.66	1.70	1.55	1.67	1.40	1.67	1.11	0.50	1.66
+ $\tilde{H}_p$ (30.0%)		0.72	1.77	1.80	1.78	1.80	1.72	1.80	1.71	0.64	1.80
+ $\tilde{H}_p$ (20.3%)		0.67	1.76	1.80	1.78	1.80	1.72	1.80	1.70	0.60	1.80
+ $\tilde{H}_p$ (9.9%)		0.59	1.79	1.80	1.74	1.79	1.71	1.79	1.62	0.54	1.79
+ $\tilde{H}_{\bar{p}}$ (29.7%)		0.41	1.80	1.81	1.73	1.82	1.70	1.82	1.51	0.40	1.82
+ $\tilde{H}_{\bar{p}}$ (19.6%)		0.39	1.84	1.80	1.70	1.79	1.64	1.79	1.36	0.38	1.84
+ $\tilde{H}_{\bar{p}}$ (10.2%)		0.38	1.88	1.79	1.59	1.75	1.34	1.75	0.99	0.36	1.88

FIGURE 4 Optimal dose–volume histograms and dosimetrics at different NTCP levels for all POLO model-based objectives. The histogram curves correspond to the target volumes and the VS, at different NTCP levels (38%: —, 30%: ---, 20%: - - -, and 10%: ···), and dosimetric values are reported with units of Gy(RBE) per fraction. We observe a positive correlation between dose values and NTCP for the VS in all cases. Additionally, for $NTCP_p$ and $\tilde{H}_p$, target coverage is not compromised, while for $LSE_{\bar{p}}$ and $\tilde{H}_{\bar{p}}$, the dose deposition in the target volumes decreases with NTCP.

overlapping with the VS, as well as a redistribution of I_d from the target volume to the margin of the PTV. Correspondingly, we see an overall reduction of p in the target volume, together with an increasing attenuation

of the hot spots around the marginal region of the PTV, and an elimination at the lower end of the VS. At strong down-regulation of the NTCP, the values are one order of magnitude smaller, with peaks in the intersections to




FIGURE 5 Optimal slice images of the dose-averaged linear energy transfer I_d (in $\text{keV}/\mu\text{m}$) for $NTCP_p$, LSE_p , $\tilde{H}_p$, and $\tilde{H}_p$ at different NTCP levels. The shifts in I_d coincide with those observed on the slice images for $d_{RBE,fx}$, that is, there seems to be a trade-off between $d_{RBE,fx}$ and I_d that is consistent with the mathematical structure of the POLO model. As a result, the I_d slice images for $NTCP_p$ and $\tilde{H}_p$ are almost equal, while LSE_p and $\tilde{H}_p$ differ. For the latter two objectives, the shift of $d_{RBE,fx}$ outside the target volume appears to be crucial, or, respectively, no shift of I_d is necessary due to the strength of the reduction of $d_{RBE,fx}$.

the 4-mm ventricular fringe and in regions where $d_{RBE,fx}$ and I_d intensify.

LSE_p showed a decay of $d_{RBE,fx}$ inside but not outside the target volume, and almost constant values of I_d , and we recognize this pattern again in the p distributions:

hot spots are softened by the isolation of high-dose and high-LET_d regions, and the p values in the target volume follow the negative trend for $d_{RBE,fx}$ at lower NTCP levels. Even the “dose gap” that occurs at an NTCP of 20% in the lower and upper parts of the target volume can




FIGURE 6 Optimal LET_d -volume histograms at different NTCP levels for all POLO model-based objectives. The histogram curves correspond to the target volumes and the VS, at different NTCP levels (38%: —, 30%: - · -, 20%: - -, and 10%: ···). We restrict the abscissa to values greater than or equal to 0.6, as the histogram curves are coherent below. With the exception of $LSE_{\bar{p}}$, where only a slight increase in I_d within the target volumes can be observed, all objectives reduce the I_d values in the VOIs.

be read from p . Last, for $\tilde{H}_{\bar{p}}$, we can correlate the results for $d_{RBE,fx}$ and I_d again to understand p . The higher I_d along the PTV margin is reflected by larger p values and local hot spots, while the region around the GTV exhibits smaller p values due to the reduction of $d_{RBE,fx}$. At the lowest NTCP level, $d_{RBE,fx}$ contracts around the GTV, leaving only a slightly upward region at the left margin of the PTV, which resembles the high- LET_d region, and hot spots where $d_{RBE,fx}$, I_d , and b interact reinforcingly.

Globally, we observe an analogous shift in the “mass” of the p distribution to lower values from the histograms in Figure 8. This indicates that the overarching goal of reducing p during optimization can be fulfilled.

3.2 | NTCP analysis

We optimized the treatment plan from Section 3.1 to NTCP levels of 29.7% – 30.5%, 19.6% – 20.3% and 9.8% – 10.2%, and while the optimal slice images and histograms show the effect of the POLO model, they are not indicative of the stability and convergence behavior of the outcome predictions derived from each optimization function (by evaluating the iterative POLO map with $NTCP_p$). Figure 9 therefore plots the NTCP estimates over the evaluation steps of the optimization scenarios.

First, we note that the scaling of the weight parameter strongly depends on the optimization function, and is not proportional to the NTCP. It seems that $LSE_{\bar{p}}$ responds

slowly until the NTCP reduction, but requires a smaller weight multiplier to reach lower NTCP levels than the other functions. $NTCP_p$ and $\tilde{H}_{\bar{p}}$ score the weights in the same order of magnitude, but need to be more and more “pushed” to achieve lower NTCP values. $\tilde{H}_{\bar{p}}$ has smaller weights due to the larger function values, but for this, lower NTCP levels are obtained with a rather constant weight multiplier.

Second, we see differences in how smoothly and efficiently the optimization functions approach the respective NTCP level. Initially, $NTCP_p$ and $\tilde{H}_{\bar{p}}$ fluctuate more widely, but then level off after around a third of the iterations. Analyzing the graphs of the optimization functions based on the linear reformulation ($LSE_{\bar{p}}$ and $\tilde{H}_{\bar{p}}$), we find a tendency towards faster convergence rates, but also stagnations on plateaus with the need for local line searches, which could be explained by gradient overshooting.

4 | DISCUSSION

This work introduces a technical framework for direct optimization of the POLO in intensity-modulated proton treatment planning for LGG patients. We implemented the forward and backward calculation steps according to Figure 1 and proposed four different optimization functions ($NTCP_p$, $LSE_{\bar{p}}$, $\tilde{H}_{\bar{p}}$, and $\tilde{H}_{\bar{p}}$), which either scalarize the mapping from an extended version of the original POLO model p or from the linear reformulation $\bar{p}$.




FIGURE 7 Optimal slice images of the probability of lesion origin p for $NTCP_p$, LSE_p , $\bar{H}_p$, and $\bar{H}_p$ at different NTCP levels. Considering the slice images for $d_{RBE,fx}$ and I_d from Figures 3 and 5, the calculated p values can be deduced from the prediction model. For example, hot spots are primarily found where high $d_{RBE,fx}$ values meet high I_d values, and for the voxels inside the 4-mm ventricular fringe where b boosts the probability prediction. We take from these results that the POLO model-based optimization of the LGG sample patient's proton plan achieves its goal, that is, reducing p in the region of interest.

By mimicking the proton treatment plan of a LGG sample patient in matRad, we were able to show the feasibility and overall efficacy of POLO model-based optimization over a number of optimization scenarios with prescribed NTCP levels.

Finally, we observed expectable shifts in dose, dose-averaged LET, and POLO distributions across the optimization functions and found that an NTCP reduction by more than 27 p.p is achievable for the sample patient without compromising target coverage.




FIGURE 8 Optimal POLO histograms at different NTCP levels for all POLO model-based objectives. We observe a consistent shift of the histogram mass to the left hand side, that is, a higher importance of the POLO model correlates with lower optimal p values.




FIGURE 9 NTCP curves for the baseline and optimal plans at different NTCP levels w of the POLO model-based objectives. All objectives achieve the desired NTCP levels, but differ in terms of weight parameters, stability, and convergence. Broadly speaking, the weights depend on the properties of the respective function, and those based on the linear reformulation from Equation 8 tend to converge faster.

Nonetheless, there are several aspects to be discussed which we will address in the following.

4.1 | Dose and LET calculation

For the purpose of this proof-of-concept, we calculate the dose-influence matrix $\mathcal{D}$ and the LET_d -influence matrix $\mathcal{L}$ with pencil-beam kernel algorithms. These are known to be very fast, but rather inaccurate compared to, for example, Markov chain Monte Carlo (MCMC) simulations. We accepted this state of affairs in view of the fact that this paper only represents a feasibility study on POLO model-based treatment plan optimization using a gradient-based framework, and in view of the time saved when running through the optimization scenarios. However, we are planning to fall back on more accurate calculation algorithms in multipatient follow-up studies to avoid artifacts in the optimal dose and LET_d distributions.

4.2 | Linear reformulation of the POLO model

In Section 2.2, we emphasized the nonideal optimization properties of the original POLO model, and therefore we introduced its linear reformulation, not as an approximation of the original model, but rather to provide a means for convex optimization and apply stronger penalties on large POLO values. While the reformulation is advantageous in theory, in practice, there are potential scaling issues to be considered due to the extended value range of the decision function in Equation 8. This leads to variants of the optimization problem from Equation 16 with different optimal points, and we observe, for example, from Figure 3 that target coverage may not be clinically acceptable when using the linearly reformulated POLO model. It should be noted that the upper limit of infinity for the decision function is never reached in real-world planning scenarios, since $d_{\text{RBE},fx}$ and I_d assume finite values, but it is likely that lower weights are necessary to balance the partial functions in the weighted sum term, as well as to avoid gradient overshooting for better convergence.

4.3 | POLO model-based optimization functions

The largest degree of freedom in our technical framework is the rightmost element of the forward and backward calculation tree, the POLO model-based optimization function $f_{p/\bar{p}}$. Naturally, there is an infinite number of functional designs, some of which may target the hot spots, others the centroid, and still others a non-standard, higher-order feature of the p distribution. We

proposed only a small selection of technically feasible, interpretable optimization functions with variable points of attack to reduce p in this paper, but even for these “simple” functions, we found evidence of behavioral differences in terms of stability and convergence in the optimization, along with nonuniform weighting schemes and nonproportional relationships of the weights to the final NTCP estimates.

Therefore, while our framework facilitates direct control of treatment planning based on the POLO model, there are important design decisions left to the users that should be customized depending on the LGG patient case.

4.4 | Mathematical and clinical perspective

The optimization results should be evaluated both mathematically and clinically. First, we were able to demonstrate the technical feasibility of POLO model-based treatment plan optimization for a selected LGG patient.

Second, a good agreement of the optimization results, that is, spatial distributions of $d_{\text{RBE},fx}$ (Figures 3 and 4), I_d (Figures 5 and 6), and p (Figures 7 and 8), with the implications of the POLO model from Section 2.1, can be found. This manifests, for example, in the shifting of $d_{\text{RBE},fx}$ and I_d such that lower values are found proximal to the VS or in the elementary product term in Equation 1, leading to attenuated p values. But, as we stated in Section 2.2, the POLO model-based objective functions touch on different aspects of p (hence also $d_{\text{RBE},fx}$ and I_d), and due to the degeneracy of the solution space for specific NTCP values obtained from the NTCP_p model in Equation 5, it occurs that differences in the concrete results can be observed, for instance, when trading off dose and dose-averaged LET.

Third, the outcome predictions (Figure 9) seem to converge across all scenarios, which indicates well-posed optimization problems. However, the choice of weights associated with the objective functions is crucial for optimization success, and we found that this is not always intuitive and might follow a trial-and-error approach. Considering the varying scales of dose-volume and POLO model-based objective functions, more sophisticated weighting schemes could come into play. Beyond that, although technically feasible, we intentionally omitted further extensions of the optimization method, for example, robust optimization, as none of these were applied in the clinical reference plan for the LGG patient, and their inclusion would limit comparability of our results with those of the ongoing phase II trial and the POLO model integration strategy developed by Sallem et al.⁷ To demonstrate feasibility, however, we added results from a composite worst-case robust optimization in the Appendix.

With regard to clinical application, the integration framework offers the user some flexibility in terms of personalized treatment planning by allowing patient-specific adjustments regarding individual anatomical variations, tumor topography relative to the ventricles, and desired dose levels. At the same time, the impact of the POLO model can be adjusted via the weight of the respective objective function, enabling the clinician to somewhat control the trade-off between dose, dose-averaged LET, and NTCP reduction. This approach is akin to the work of Sallem et al.,⁷ who showed that clinical objectives can be achieved while reducing NTCP levels for a number of LGG patients. Also, the INDIGO (INDividualized, model-Guided Optimization of proton beam treatment plans for LGG treatment) protocol presented there promotes active CTV dose de-escalation, thus we did not strictly enforce target coverage in our approach to observe the resulting trade-offs. For method selection, the capability to preserve target coverage may, however, serve as one of the decision criteria.

Nonetheless, we analyzed our integration framework by picking a single LGG patient, which is sufficient to show the feasibility and potential implications of the approach, but limits the ability to draw conclusions about the generalizability of the framework. In this regard, we are not expecting technical issues when transferring the direct optimization approach to other patients. However, uncertainties about the clinical viability would be best addressed by a full treatment plan comparison study on a larger dataset, preferably on the patient cohort currently being recruited for the ongoing phase II trial. A future study such as this could focus on the more clinical aspects of POLO model-based treatment planning and on the reproducibility of clinical decision making in the ongoing trial.

Adding more complex tumor geometries and more intricately located OARs could render the integration framework even more comprehensive. Beyond this, a comparative analysis with the methodology of Sallem et al.⁷ could provide crucial insights into the clinical feasibility and potential advantages of our approach in real-world settings, which is essential for evaluating its clinical impact and practical implementation. Specifically, comparing the dose distributions between our study and that of Sallem et al.⁷ at equivalent NTCP levels would yield valuable information on target coverage and the sparing of OARs, with a particular focus on the cerebral ventricles.

4.5 | Translation to other models

The methodology described in Section 2 can be transferred to other voxel-based outcome models if (i) the model assumes a dose–response relationship, that is, takes into account dosimetric variables, (ii) the model

provides a gradient with respect to the dosimetric variables, and (iii) the voxel-wise model output can be scalarized with an optimization function f .

If these criteria are met, implementing the outcome model for treatment plan optimization is straightforward and only requires to set up the optimization function (e.g., to transform $p(\eta, k)$ into f_p), as well as to define the projection between the spatial dose and outcome distributions (e.g., to calculate $p(\eta, k)$ from ϕ forwards and derive $\nabla_{\phi} f_p$ backwards). This holds true for a wide range of statistical and machine learning models, including linear models and (deep) neural networks.

5 | CONCLUSION

We presented a technical framework to directly integrate the POLO model from Bahn et al.³ into proton treatment planning for LGG patients by defining the forward calculation and backward differentiation steps required to perform model-based dose optimization.

Implementing feasible optimization functions and running a number of optimization scenarios for fixed NTCP levels, we were able to generate results on the spatial $d_{RBE,fx}$, I_d , and p distributions. From a mathematical standpoint, the results are consistent regarding changes expected from model inspection and changes observed between plans with different weights of the POLO model-based optimization function, and we also observe a consistent reduction in NTCP across all scenarios without affecting target coverage for a sample patient.

As already mentioned, the clinical feasibility of the integration framework, particularly its capacity to balance effective tumor control with the minimization of side effects, still requires comprehensive evaluation to ensure that it is both practical and beneficial in a real-world clinical setting. Conclusively, the integration framework may well contribute to automated POLO model-based risk minimization in proton treatment planning for LGG patients, and should further be evaluated in multipatient follow-up studies.

ACKNOWLEDGMENTS

The present contribution is supported by the Helmholtz Association under the joint research school HIDSS4Health—Helmholtz Information and Data Science School for Health. H.S. and J.B. acknowledge financial support from the “Alois-Hirdt-Erben und Wieland” foundation Heidelberg (Germany). N.W. acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—Project No. 443188743.

Open access funding enabled and organized by Projekt DEAL.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

REFERENCES

- Harrabi SB, Nettelbladt vB, Gudden C, et al. Radiation induced contrast enhancement after proton beam therapy in patients with low grade glioma – how safe are protons? *Radiother Oncol*. 2022;167:211-218. doi:10.1016/j.radonc.2021.12.035
- Greene-Schloesser D, Robbins M, Peiffer AM, Shaw E, Chan MD, Wheeler KT. Radiation-induced brain injury: a review. *Front Oncol*. 2012;2:73. doi:10.3389/fonc.2012.00073
- Bahn E, Bauer J, Harrabi SB, Herfarth K, Debus J, Alber M. Late contrast enhancing brain lesions in proton-treated patients with low-grade glioma: clinical evidence for increased periventricular sensitivity and variable RBE. *Int J Radiat Oncol Biol Phys*. 2020;107(3):571-578. doi:10.1016/j.ijrobp.2020.03.013
- Eulitz J, Troost E, Raschke F, et al. Predicting late magnetic resonance image changes in glioma patients after proton therapy. *Acta Oncol*. 2019;58(10):1536-1539.
- Bauer J, Bahn E, Harrabi SB, Herfarth K, Debus J, Alber M. How can scanned proton beam treatment planning for low-grade glioma cope with increased distal RBE and locally increased radiosensitivity for late MR-detected brain lesions? *Med Phys*. 2021;48(4):1497-1507. doi:10.1002/mp.14739
- Eulitz J, Troost E, Klünder L, et al. Increased relative biological effectiveness and periventricular radiosensitivity in proton therapy of glioma patients. *Radiother Oncol*. 2023;178:109422. doi:10.1016/j.radonc.2022.11.011
- Sallem H, Harrabi SB, Traneus E, Herfarth K, Debus J, Bauer J. A model-based risk-minimizing proton treatment planning concept for brain injury prevention in low-grade glioma patients. *Radiother Oncol*. 2024;201(2):110579. doi:10.1016/j.radonc.2024.110579
- Palkowitsch M, Kaufmann LM, Hennings F, et al. Variable-RBE-induced NTCP predictions for various side-effects following proton therapy for brain tumors – identification of high-risk patients and risk mitigation. *Radiother Oncol*. 2025;202:110590. doi:10.1016/j.radonc.2024.110590
- Palma G, Monti S, Cella L. Voxel-based analysis in radiation oncology: a methodological cookbook. *Physica Medica*. 2020;69:192-204. doi:10.1016/j.ejmp.2019.12.013
- Ebert MA, Gulliford S, Acosta O, et al. Spatial descriptions of radiotherapy dose: normal tissue complication models and statistical associations. *Phys Med Biol*. 2021;66:12TR01.
- McWilliam A, Palma G, Abravan A, et al. Voxel-based analysis: roadmap for clinical translation. *Radiother Oncol*. 2023;188:109868. doi:10.1016/j.radonc.2023.109868
- Acosta O, Dreon G, Ospina JD, et al. Voxel-based population analysis for correlating local dose and rectal toxicity in prostate cancer radiotherapy. *Phys Med Biol*. 2013;58(8):2581. doi:10.1088/0031-9155/58/8/2581
- Palma G, Monti S, D'Avino V, et al. A voxel-based approach to explore local dose differences associated with radiation-induced lung damage. *Int J Radiat Oncol Biol Phys*. 2016;96(1):127-133. doi:10.1016/j.ijrobp.2016.04.033
- Monti S, Palma G, D'Avino V, et al. Voxel-based analysis unveils regional dose differences associated with radiation-induced morbidity in head and neck cancer patients. *Sci Rep*. 2017;7:7220. doi:10.1038/s41598-017-07586-x
- Brahme A. Optimized radiation therapy based on radiobiological objectives. *Semin Radiat Oncol*. 1999;9(1):35-47. doi:10.1016/S1053-4296(99)80053-8
- Semenenko VA, Reitz B, Day E, Qi XS, Miften M, Li XA. Evaluation of a commercial biologically based IMRT treatment planning system. *Med Phys*. 2008;35(12):5851-5860. doi:10.1118/1.3013556
- Kierkels R, Korevaar E, Steenbakkers R, et al. Direct use of multivariable normal tissue complication probability models in treatment plan optimisation for individualised head and neck cancer radiotherapy produces clinically acceptable treatment plans. *Radiother Oncol J Eur Soc Ther Radiol Oncol*. 2014;112:430-436. doi:10.1016/j.radonc.2014.08.020
- Ajdari A, Liao Z, Mohan R, Wei X, Bortfeld T. Personalized mid-course FDG-PET based adaptive treatment planning for non-small cell lung cancer using machine learning and optimization. *Phys Med Biol*. 2022;67(18):185015. doi:10.1088/1361-6560/ac88b3
- Maragno D, Buti G, Birbil SI, et al. Embedding machine learning based toxicity models within radiotherapy treatment plan optimization. *Phys Med Biol*. 2024;69(7):075003. doi:10.1088/1361-6560/ad2d7e
- Wieser HP, Cisternas E, Wahl N, et al. Development of the open-source dose calculation and optimization toolkit matRad. *Med Phys*. 2017;44(6):2556-2568. doi:10.1002/mp.12251
- Wächter A, Biegler LT. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math Program*. 2006;106:25-57. doi:10.1007/s10107-004-0559-y
- Niemierko A, Goitein M. Calculation of normal tissue complication probability and dose-volume histogram reduction schemes for tissues with a critical element architecture. *Radiother Oncol*. 1991;20(3):166-176. doi:10.1016/0167-8140(91)90093-V
- Unkelbach J, Botas P, Giantsoudi D, Gorissen BL, Paganetti H. Reoptimization of intensity modulated proton therapy plans based on linear energy transfer. *Int J Radiat Oncol Biol Phys*. 2016;96(5):1097-1106. doi:10.1016/j.ijrobp.2016.08.038
- Ackermann B, Bangert M, Bennan ABA, et al. e0404/matRad: Blaise v2.10.1. Version 2.10.1; 2020. doi:10.5281/zenodo.7107719
- Siddon RL. Fast calculation of the exact radiological path for a three-dimensional CT array. *Med Phys*. 1985;12(2):252-255. doi:10.1118/1.595715
- Hong L, Goitein M, Bucciolini M, et al. A pencil beam algorithm for proton dose calculations. *Phys Med Biol*. 1996;41(8):1305. doi:10.1088/0031-9155/41/8/005
- Duff IS. MA57—a code for the solution of sparse symmetric definite and indefinite systems. *ACM Trans Math Softw*. 2004;30(2):118-144. doi:10.1145/992200.992202

How to cite this article: Ortkamp T, Sallem H, Harrabi S, et al. Direct optimization of the probability of lesion origin in proton treatment planning for low-grade glioma patients. *Med Phys*. 2026;53:e70395. <https://doi.org/10.1002/mp.70395>

APPENDIX A

Let $\tilde{d} = D_r \phi \in \mathbb{R}_+^{n_r}$ be the vectorized dose distribution in the volume of interest r , using the dose-influence matrix $D \in \mathbb{R}_+^{n \times m}$ and the fluence vector $\phi \in \mathbb{R}_+^m$. Then, the dose optimization functions stated in Section 2.4 are given below.

A.1 | Dose uniformity

Dose uniformity measures the standard deviation of the dose distribution, that is,

$$f_{\text{Dose Uniformity}}(\tilde{d}) = \sqrt{\text{Var}(\tilde{d})}.$$

A.2 | Squared deviation

Squared deviation includes a reference dose parameter d^{ref} , and quantifies the normalized sum of squared differences between the dose values and the reference value, that is,

$$f_{\text{Squared Deviation}}(\tilde{d}) = \frac{1}{n_r} (\tilde{d} - d^{\text{ref}})^T (\tilde{d} - d^{\text{ref}}).$$

A.3 | Squared underdosing

Squared underdosing is structurally similar to squared deviation, except that it uses a minimum dose parameter d^{min} and only penalizes lower values, yielding the function

$$f_{\text{Squared Underdosing}}(\tilde{d}) = \frac{1}{n_r} \varepsilon^T \varepsilon,$$

$$\varepsilon_i = \begin{cases} d^{\text{min}} - \tilde{d}_i, & \tilde{d}_i < d^{\text{min}} \\ 0, & \text{else} \end{cases},$$

where ε “clamps” the dose difference values to only consider underdosed voxels.

A.4 | Maximum DVH

Maximum DVH takes a reference dose d^{ref} and a maximum volume fraction v^{max} as input parameters, calculates the dose quantile d^v at $v = v^{\text{max}}$ and then determines

$$f_{\text{Maximum DVH}}(\tilde{d}) = \frac{1}{n_r} \varepsilon^T \varepsilon,$$

$$\varepsilon_i = \begin{cases} \tilde{d}_i - d^{\text{ref}}, & \tilde{d}_i \in (d^{\text{ref}}, d^v] \text{ and } d^{\text{ref}} < d^v \\ 0, & \text{else} \end{cases},$$

where ε is only nonzero if the respective dose value contributes to moving the DVH quantile beyond the reference value.

A.5 | Minimum DVH

Minimum DVH, analogous to maximum DVH, utilizes the value function

$$f_{\text{Minimum DVH}}(\tilde{d}) = \frac{1}{n_r} \varepsilon^T \varepsilon,$$

$$\varepsilon_i = \begin{cases} \tilde{d}_i - d^{\text{ref}}, & \tilde{d}_i \in (d^v, d^{\text{ref}}] \text{ and } d^v < d^{\text{ref}} \\ 0, & \text{else} \end{cases},$$

TABLE A.1 Optimization runtimes in matRad for baseline and POLO model-based treatment plans. The baseline plan requires 0.14 s (CPU) or 0.35 s (wall clock) per iteration, excluding LET and POLO calculation. By adding the POLO model-based optimization components as described in Section 2, time per evaluation increases by a factor of 2.9–6.4 (CPU) or 2.0–4.4 (wall clock).

Optimization case	# IPOPT evaluations	CPU time (s)	Wall-clock time (s)
Baseline (NTCP = 37.6%)	878	82	128
+NTCP _p (30.2%)	1211	719	771
+NTCP _p (20.0%)	1452	665	711
+NTCP _p (10.2%)	1362	584	627
+LSE _p (30.5%)	279	87	120
+LSE _p (20.3%)	515	141	175
+LSE _p (9.8%)	420	120	152
+H _p (30.0%)	1162	675	726
+H _p (20.3%)	1572	662	706
+H _p (9.9%)	646	289	326
+H _p (29.7%)	381	114	147
+H _p (19.6%)	1977	532	569
+H _p (10.2%)	983	264	298

but defines a minimum volume fraction v^{min} instead of v^{max} .

A.6 | Minimum/Maximum dose

Minimum/Maximum dose implements a box constraint on the dose, that is, the dose values are lower- and upper-bounded by d^{min} and d^{max} . The constraint function is given by the identity and therefore the constraint reads

$$d^{\text{min}} \leq \tilde{d} \leq d^{\text{max}}.$$

Internally, matRad translates such a constraint by applying the log-sum-exp approximation.

A.7 | Computational costs

Table A.1 lists the number of IPOPT evaluations, CPU, and wall clock runtimes for optimizing the baseline and POLO model-based treatment plans from Section 2.4.

A.8 | Robust optimization

To address clinical requirements for plan stability, we also performed robust POLO model-based treatment plan optimization. The uncertainty set comprises nine discrete scenarios (eight perturbed and one nominal), incorporating $\pm 2.5\%$ range and 2-mm setup uncertainty. We employed a composite worst-case approach; however, dose uniformity objectives were kept nominal to prevent excessive “dose smearing.” As proof-of-concept, we integrated the POLO model via NTCP_p. Figure A.1 illustrates the optimization results for the robust baseline plan (nominal NTCP: 22%) and the robust POLO model-based plan (nominal NTCP: 13%), including slice images and DVH/NTCP uncertainty bands. Our results




FIGURE A.1 Composite worst-case robust optimization results. Top panels show slice images for dose, LET_d , and POLO. Bottom panels display the respective DVHs (baseline: —, POLO: - -) and NTCP curves (including scenario bands). We observe a dose reduction and LET redistribution, leading to an overall reduction in the nominal NTCP from 22% (baseline) to 13% (POLO model-based).

indicate reduced maximum doses to the tumor volumes and the VS, alongside an LET shift away from the latter. Although NTCP decreased, earlier optimizer termination suggests objective competition and the attainment of a Pareto-optimal trade-off between tumor coverage and OAR sparing. More detailed investigation, however, goes

beyond the scope of this paper, especially since the clinical reference protocol does not include robustness.