



Treatment planning of scanned proton beams in RayStation

Martin Janson*, Lars Glimelius, Albin Fredriksson, Erik Traneus, Erik Engwall

RaySearch Laboratories AB, Stockholm, Sweden

ARTICLE INFO

Article history:

Received 7 September 2023

Revised 17 October 2023

Accepted 22 October 2023

Keywords:

Proton therapy

Treatment planning

RayStation

Robust optimization

4D optimization

Monte Carlo dose calculation

Proton arc therapy

Deep learning auto-planning

Robustness evaluation

Scripting

ABSTRACT

The use of scanned proton beams in external beam radiation therapy has seen a rapid development over the past decade. This technique places new demands on treatment planning, as compared to conventional photon-based radiation therapy. In this article, several proton specific functions as implemented in the treatment planning system RayStation are presented. We will cover algorithms for energy layer and spot selection, basic optimization including the handling of spot weight limits, optimization of the linear energy transfer (LET) distribution, robust optimization including the special case of 4D optimization, proton arc planning, and automatic planning using deep learning. We will further present the Monte Carlo (MC) proton dose engine in RayStation to some detail, from the material interpretation of the CT data, through the beam model parameterization, to the actual MC transport mechanism. Useful tools for plan evaluation, including robustness evaluation, and the versatile scripting interface are also described. The overall aim of the paper is to give an overview of some of the key proton planning functions in RayStation, with example usages, and at the same time provide the details about the underlying algorithms that previously have not been fully publicly available.

© 2023 The Authors. Published by Elsevier Inc. on behalf of American Association of Medical Dosimetrists.

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>)

Introduction

RayStation is a treatment planning system (TPS) for external beam therapy, brachytherapy, and boron neutron capture therapy (BNCT) that is developed, sold, and maintained by the Swedish software company RaySearch Laboratories AB. RaySearch was founded in 2000 and initially provided various plugins to other vendors' TPSs, perhaps most notable the IMRT optimization engine for the Pinnacle TPS (Philips). In 2008, RaySearch decided to develop and sell a complete TPS directly to the end user, and the first clinical version of RayStation was released in 2010. This first release was limited to planning for IMRT and VMAT photon treatments, but support for proton planning was part of the original scope and the very first RayStation contract was indeed signed with a proton clinic, the Westdeutsches Protonentherapie-zentrum Essen (WPE), in 2009. Mainly due to the lack of available dosimetric data at these early days of proton therapy, the clinical release of the proton planning in RayStation was delayed, but, in 2014, the first proton treatment with a plan created in RayStation was delivered at the Provision Center for Proton Therapy in Knoxville, Tennessee. Since then, more than 100 clinics around the world

have chosen RayStation for their proton treatment planning, a figure representing a majority of proton centers that are in operation, under construction, or in the planning stage.¹ This rapid success can likely be explained by the fact that RayStation emerged just in time to support the near explosive growth in the number of proton centers worldwide, and that RayStation was built with the needs of this relatively new modality in mind from day one. Robust optimization, 4D optimization, fast Monte Carlo (MC) dose calculations, and, in later years, linear energy transfer (LET) evaluation and optimization are just a few examples of features originally introduced in RayStation, that have added to its popularity. Another reason for the wide spread of RayStation in the proton community is the fact that RayStation has been adapted to a plurality of proton delivery systems. Today, RayStation is in clinical use at sites with delivery systems from: IBA, Varian, Mevion, Hitachi, Sumitomo Heavy Industries, P-Cure, Mitsubishi, ProNova, plus a variety of "home-built" synchrotron-based systems for regular treatments and low-energy single scattering systems for ocular treatments. RayStation further supports the following delivery techniques: pencil beam scanning (PBS), quasi-discrete PBS, line scanning, double/single scattering (DS/SS), uniform scanning, and wobbling.

In this paper, we mainly focus on RayStation functionalities that are related to optimization and dose calculation of scanned pro-

* Reprint requests to Martin Janson, Box 45169, SE-104 30 Stockholm, SWEDEN.
E-mail address: martin.janson@raysearchlabs.com (M. Janson).

ton beams. Other vital functionalities in the planning process, such as patient modeling, deformable image registration, and fallback planning have to some degree been covered in a previous general overview of RayStation² and will not be included here. The aim of this publication is to present the planning process in general, and to give a relatively detailed explanation of the fundamental algorithms and their underlying mechanisms, descriptions that previously have only been partly available through product documents (e.g., users and reference manuals) and white papers. All descriptions in this paper relate to RayStation 2023B, the currently latest released clinical version.

Plan Optimization

This section will describe the creation of a proton plan by optimization. We will start from the point where the empty beams of a plan have already been created, but the patient-specific beam devices and optimization settings are yet to be defined.

Range shifters, apertures, and ridge filters

For most proton delivery systems, an energy absorbing range shifter must be used for shallow targets (ranges below 4–7 cm). In RayStation, the use of a range shifter is normally selected manually per field but may optionally be automatically added when needed.

RayStation supports the use of patient-specific collimators with PBS beams. The collimator may be a milled-out block aperture or a multi leaf collimator (MLC). The dynamic MLC of the Mevion Hyperscan delivery system³ is also supported. The aperture/MLC shapes are determined prior to the optimization as the geometrical projections of the target(s) including user-defined, non-isotropic margins. Blocking of risk organs may also be defined for the collimation. When (MLC) collimation is done for each energy layer individually, the aperture openings are determined as the projection of the target at the Bragg peak depth of each energy.

Some delivery systems have very narrow pristine Bragg peaks at lower energies and may benefit from using a ridge filter to decrease the number of energy layers and thereby the delivery time. This is common for synchrotron-based systems, but a recent study shows that ridge filters also can be used to significantly decrease the delivery time for cyclotron systems and thereby increase the window for cases tolerating breath hold techniques.⁴ RayStation does not have explicit support for ridge filters, but the effect can be implicitly included via the beam data used for beam modeling. For such ridge filter models, it is possible to include the identity of the ridge filter in the exported DICOM plan.

Spot and energy layer selection settings

Before an optimization starts, the target is analyzed in terms of radiological depth to the most proximal and distal points over the lateral plane, as well as its projection on the fluence plane for each field. Based on this target information and a multitude of user-defined settings, the energy layers and spot positions are selected. The energy layer spacing can either be determined automatically, based on the 80% widths of the Bragg peaks in the machine model scaled by a user-defined factor, or be set to a constant water-equivalent thickness (WET). The user may further define any number of additional proximal and distal energy layers. While these extra layers normally are not needed, there can be cases with inhomogeneous geometries and small targets where the raytracing used to define target depths is insufficient.

The spot pattern can be chosen as hexagonal (default), or square, and the direction of the pattern can be chosen freely to, e.g., align it with the fast scan direction of the delivery system. Just like for the energy layer selection, the spot spacing can be chosen

as automatic or as a fixed value. For the automatic option, the spot spacing is determined as 1.06 times the average spot size (1σ) in the patient at the Bragg peak depth, multiplied by a user-defined scaling constant. For most patients and delivery systems, the automatic option with a scaling constant of 1 works well, but for systems with very large spots, a decrease of the spot spacing scaling factor to ~ 0.6 may improve the plan, especially when used in combination with apertures. The lateral margin of the spot placement relative to the target volume can also be determined using an automatic with scale option, or as a constant value. The lateral margin is needed to achieve full target coverage without creating hot spots at the target border, but if a higher dose at the target edge can be tolerated, a target margin of 0 may be used to minimize the lateral penumbra. When robust optimization is used, additional target margins based on the magnitude of the uncertainties are automatically added to these user-defined values.

Other ways to control the spot placement in RayStation include the OAR range margin, where Bragg peaks may only be placed up to a given proximal distance from selected ROIs, as well as the option to set minimum and maximum radiological depths. The latter function may be used to avoid placing spots at the skin when shallow tumors are treated, or when employing a field-in-field technique where a range shifter is only used for the shallow field.⁵

To mitigate effects from organ motion, it is further possible to specify a layer repainting strategy. The number of paintings can either be specified by a constant number per beam or be based on a maximum meterset per layer or spot.⁶ In the former case, the number of paintings is the same for all layers in the beam, whereas in the latter cases, the number of paintings varies over the different energy layers. The repainting instruction can either be communicated to the delivery system by a single number given per energy layer, or as an explicit delivery sequence with all repainted layers included in the DICOM file.

Optimization algorithm

The dose-based optimization of a treatment plan, also referred to as inverse planning, is a technique where the clinical goals of a plan are achieved through iterative adjustments of the plan parameters using a numerical optimization method. The optimization principle of scanned proton beams is straightforward since the optimization variables are the spot weights, which contribute to the dose in a linear fashion. The optimization engine in RayStation is gradient-based and uses a sequential quadratic programming algorithm.⁷ The algorithm employs a quasi-Newton approximation of the Hessian of the Lagrangian that is updated using the Broyden-Fletcher-Goldfarb-Shanno approach. In each iteration, a better approximation of the Hessian is obtained, but it also becomes more computationally costly to use.⁸ As a consequence, the optimization time increases superlinearly with the number of iterations, an effect that explains the sometimes rather extended optimization times for plans with a high number of spots and where the number of iterations has passed ~ 100 .

When a new optimization is started in RayStation, the target is analyzed, and energy layers and spot positions are selected as described in the previous section. A starting guess of the spot weights is then determined by analytical means,⁹ where it is assumed that the beam doses will completely overlap the target. This method works very well for cases like single field uniform dose (SFUD) optimization¹⁰ of a prostate but becomes less effective for cases where the beam doses do not overlap fully, such as cranio-spinal plans, which may require a few extra iterations for the dose to stabilize. The spot doses are then computed according to the settings for the optimization dose calculation (see [Sec. Statistics](#)). With the GPU-based MC algorithm this should normally be a fast process (~ 10 seconds), but to avoid computing the

same dose twice, a so-called spot cache is implemented in RayStation which stores spot doses between consecutive optimizations. The spot doses are stored in the cache until a change is made that would either affect the dose calculation (e.g., a change in the patient geometry or the snout position) or the spot pattern (e.g., some change in the spot selection parameters, the definition of the target volume through the target optimization functions, or robustness settings). The importance of the spot cache becomes greater for demanding cases such as large volume/very small voxel plans, or a robust optimization using the option of accurate scenario dose calculation (see further [Sec. Approximate and accurate scenario doses](#)).

The optimization continues until any of the 3 following criteria is fulfilled: (1) the change of the optimization objective function value has fallen below the optimization tolerance for 3 consecutive iterations, (2) the maximum number of iterations has been reached, or (3) no improving direction can be found. The tolerance and the max number of iterations are user-defined parameters. The user can also terminate the optimization manually at any time.

After an optimization has stopped, it may be started again using the *Continue optimization* function which resumes the optimization from the point it was stopped. This can be useful for example after smaller adjustments to the optimization functions, or when spots have been manually removed. However, it should be noted that a repeated use of continued optimizations may lead to a different solution compared to that of a full restart. One reason being that the spot filter function might remove spots several times. When possible, it is advisable to always start an optimization from scratch, at least when a significant change to the plan setup or the optimization problem has been made.

Optimization functions

An optimization problem in RayStation is formulated with an objective function whose value is to be minimized iteratively by alteration of the spot weights. The objective function consists of one or several user-defined optimization functions of the type: min/max/uniform dose, min/max DVH, min/max/target EUD, and dose fall-off, that operate on the dose within one ROI specified for each function.^{11,12} The dose fall-off objective is a max dose objective that penalizes dose above a dose level that changes as a function of the distance to the target and is therefore ideal to use for penalizing the dose outside of the target volume, omitting the need for creating ring structures around the target. The optimization functions included in the objective function are assigned weights reflecting their relative importance. The value of an optimization function in RayStation is further divided by the volume of the ROI the function is specified for. This means that a hot spot in a small ROI will have a larger impact on the optimization than a similarly sized hot spot in a larger ROI. The value is also normalized to the dose level of the function, meaning that the value of the function is determined by the relative dose deviation. For example, an overdose of 5 Gy for a max dose function with a dose level of 50 Gy will give the same penalty as an overdose of only 1 Gy if the dose level is set to 10 Gy. Optimization functions with small dose levels must thus be given small weights in order not to dominate the optimization.

A unique feature in RayStation is that an optimization function can be defined as beam-specific, meaning that the function value is determined by the dose from a single beam. An objective function may include a mix of normal and beam-specific optimization functions. The beam-specific functions can be used to create SFUD plans, but also allow for distributing the responsibility of treating various sub-volumes of a target among the different beams. A typical example would be a head and neck case where a beam from the left (right) should take care of the left (right) side of the head,

thus avoiding crossing of the brain stem. Beam-specific functions also affect the spot selection. When a beam has been associated with a dose driving beam-specific function, the target used for the energy layer and spot selection of that beam will be replaced by the ROI of the beam-specific function (plus any other ROIs that are associated to a dose driving beam-specific function for the same beam). This feature may be used to control the spot placement in an optimization, also when the optimization functions shall only depend on the total dose. To accomplish this, beam-specific functions with a weight equal to zero are used. An example of this is the creation of graded junctions for a cranio-spinal irradiation (CSI) where zero-weighted beam-specific functions can be used to define the overlap of spots from adjacent fields, and thereby the length of the junctions, while total dose robust functions for the entire CTV are used to create the smooth junctions of the beam doses, and a uniform total dose within the CTV.

In RayStation, an optimization function can optionally be defined as a constraint that must be fulfilled in the optimization. The optimal solution to an optimization problem is then the one that gives the best value of the objective function without violating any of the constraints. While constraints sometimes can be useful in an optimization, they should be handled with some caution. If more than one constraint is defined for an optimization, it is of pivotal importance that they are not in mutual conflict so that all constraints in the optimization can be fulfilled. This is in opposition to how optimization functions work, where the optimizer will find a well-defined (Pareto optimal) trade-off between conflicting functions. Another reason to be careful with an extensive use of constraint is that the convergence of the optimization becomes slower compared to an optimization without constraints. A reason for this is that the optimization algorithm requires one gradient calculation for the objective function, and an additional one for each constraint. Thus, it can speed up the treatment planning process to run an initial optimization without constraints to see what can be achieved, followed by adding constraints in a subsequent round of optimization to precisely control, e.g., the max dose. Furthermore, when using constraints, it can be beneficial to include an identical function in the objective, to help steering the optimization in the right direction.

RayStation also provides the possibility of defining optimization functions on the minimum and maximum dose-averaged linear energy transfer (LET_d) in addition to dose. Since LET_d is a property that does not scale with fluence and dose, high LET_d values may be found outside the target in volumes where the dose is so low that the LET_d level has no significant biological impact. To exclude these low dose volumes from the optimization, the max LET_d function is accompanied with a dose threshold that filters out those volumes in the function evaluation. Just like any other function in RayStation, the LET_d functions may be defined as being beam-specific, constraints, and/or robust. Robust max LET_d functions may be a particularly important use-case for patients with risk organs just distal to a target, considering the range uncertainty and the expected elevated biological effective dose at the distal end of a proton field.¹³

Spot and energy layer weight limits

Most proton PBS delivery systems have a lower limit for the spot meterset. In order to create a deliverable plan, a common approach has previously been to apply a spot filtering step after the optimization is completed.¹⁴ However, this post-processing filtering can seriously deteriorate plan quality, especially for plans with several overlapping fields.¹⁴ In RayStation, the minimum spot meterset is instead considered in the optimization, where the filtering is performed after a user-specified number of iterations, typically at iteration 20. The optimization then continues so that plan qual-

ity can be restored. At the spot filtering iteration, the spots in each energy layer are sequentially filtered out and the weight of a removed spot is added to the immediately following spot to avoid an excessive removal of spots with weights just below the minimum limit. To ensure that the spot weights do not fall below the minimum spot weight when the optimization continues, the minimum spot weight is used as a lower bound in the continued optimization. It is also possible to define an upper limit to the spot weights, which is then enforced from the beginning of the optimization. If the spot filtering iteration is set to 0, no spots will be removed, and the lower and upper spot weight bounds are applied from the first iteration. This can for example be useful when *Continue optimization* is used in RayStation.

It is further possible to define a minimum energy layer meter-set in RayStation, which removes all energy layers whose meterset falls below the limit at the spot filtering iteration. The energy layer lower limit is then included as a lower bound in the continued optimization.

The default spot and energy layer weight limits are defined in the beam model. However, those values may be overridden on a plan-by-plan basis by the user. For example, by using a higher minimum spot weight limit than that stipulated by the delivery system, the user has the possibility to explore the balance between delivery speed and plan quality.

After the spot filtration, the spot sequence can optionally be sorted, taking different scan speeds in the x- and y-directions into account, to make the traversal of the scanning beam as fast as possible.

Robust Optimization

In traditional treatment planning, uncertainties during the setup and delivery are handled by expanding the clinical target volume (CTV) with a uniform margin to create a planning target volume (PTV).¹⁵ However, due to the higher dose conformity and the introduction of range uncertainties in proton planning, the PTV concept does not always work so well for this modality.¹⁶ The solution to this problem is robust optimization, where the uncertainties are explicitly included in the optimization, which then is performed on the CTV and original risk organ volumes directly.^{17,18}

RayStation employs a scenario-based robust optimization technique, where the dose in each iteration is computed for a number of “error” scenarios, each representing a specific combination of setup error, density error, and patient image set. In RayStation, each optimization function can optionally be labeled as a robust function and the objective function may contain a mix of non-robust and robust functions.

The uncertainties can be treated as systematic, meaning that the same error occurs throughout the treatment course, or random, meaning that different errors may occur in each fraction. For systematic uncertainties, RayStation employs worst-case (or “minimax”) optimization,¹⁷ where the optimization aims to minimize the function value in the worst scenario. If more than one robust function is included, the optimization strives to minimize the sum of the function values of all robust functions in the worst scenario. The reason for considering the worst-case sum of function values instead of, e.g., the sum of the worst-case value for each function individually is that only physically realizable scenarios should affect the optimization, and not unphysical combination scenarios where, for example, the patient moves to the left when the CTV function is evaluated but to the right when the OAR function is evaluated.¹⁹

When the uncertainties are treated as random (or interfractional), a scenario in the optimization is constructed as the sum of the dose in each fraction of the treatment course, where one error scenario is randomly selected for each fraction.²⁰ The optimization

is performed on the average function value of all simulated treatment course scenarios. To cover the many possible combinations of errors in the fractions, a large number of treatment course scenarios should be included in the optimization (the default number is 3000 for 3 systematic density errors). In the limit of a single fraction treatment course, the random uncertainty option will give a similar result as if systematic uncertainties had been considered,²¹ while in the limit of many fractions, the random optimization will give the same result as optimizing the average dose of all considered single fraction scenarios. When random uncertainties are considered, the resulting plan may underdose the target or overdose the OARs in the nominal scenario, i.e., the case where no uncertainty is considered. To maintain the dose of the nominal scenario, non-robust functions can be applied to the structures that have been associated with random robust functions.

Only setup errors and patient images can be handled as random over the fractions, while the density errors are always considered to be systematic.

Setup and density error scenario sampling

The number of scenarios in a robust optimization depends on the magnitude of the setup and range uncertainties. For setup errors (e_s) up to 5 mm, 2 initial setup scenarios are created for each dimension: ($-e_s$, e_s) giving a total of 7 initial setup scenarios for 3 spatial dimensions, including the non-shifted nominal scenario. For density errors (e_d) up to 5%, 2 initial density scenarios are created giving a total of 3 initial density scenarios ($-e_d$, 0, e_d). Each combination of initial setup and density scenarios is then considered resulting in a total of 21 scenarios for the optimization assuming systematic errors. For larger uncertainties, additional initial scenarios are created to ensure that the scenario sampling is not too sparse. These extra scenarios may be created both along the major setup error axes, as well as in diagonal directions depending on the projected magnitude of the setup errors in these directions.

When the error scenarios are created for the optimization, the patient setup may be chosen to be the same for all beams in a plan, to be different (independent) for beams belonging to different isocenters, or to be different for all beams. This choice will greatly affect the number of evaluated error scenarios in the optimization. In the example above, with 7 setup scenarios, it is assumed that the setup error is the same for all beams in the plan. If we instead assume that the plan has 2 different isocenters with 2 beams associated to each isocenter and we use the independent isocenter option, the number of initial spatial error scenarios becomes $7^2 = 49$, resulting in a total of 147 scenarios when combined with the density scenarios. If we now assume that the setup uncertainties are independent for all 4 beams, the spatial error scenarios become $7^4 = 2401$, resulting in a total of 7203 scenarios. Although optimization with this large number of scenarios is fully possible in RayStation, the optimization time will suffer. To reduce the number of scenarios when using independent isocenters/beams, it is possible to restrict the independence to 1 or 2 directions. A suitable use case for systematic robust optimization with independent setup errors is the creation of graded junctions of a 4-field CSI case, with 2 fields sharing an isocenter in the brain and with 2 additional fields, each with separate isocenters along the spine.²² Here we use the independent isocenter option, but only to be evaluated as independent in the superior-inferior direction along the spine. If we further only consider setup errors (since density errors affect the dose only marginally due to the shallow target), the problem reduces to manageable $3^3 + 4 = 31$ scenarios, which will optimize swiftly. The independent setup error option in RayStation was indeed originally developed with the CSI case in mind, although the function has found a considerable wider use since then.

4D optimization

4D optimization is a technique where multiple images of the patient are considered simultaneously in an optimization. The most wide-spread example is probably the 4D optimization of a lung case using a 4D image dataset, a technique that for protons has been shown to be superior compared to the traditional method of optimizing on an average CT using an ITV.^{23,24}

While robust optimization of a 4D image dataset may be the obvious use case for 4D robust optimization, there are also other usages. One interesting example is the mitigation of unknown air cavities for a pelvic patient. This can be accomplished in RayStation by first making a copy of the planning CT, assigning air as override material for selected volumes of the rectum/intestines in the copied CT, and then include it in a 4D robust optimization.²⁵ When the motion of an internal organ is not known, RayStation can simulate this motion by creating a series of deformed images using the *Simulate organ motion* tool, images that then can be used in a 4D optimization. A typical example here is the daily random motion of the prostate.

4D optimization is a natural extension of the scenario based robust optimization in RayStation, where the additional images simply form additional scenarios. It may be used in combination with setup and density uncertainties, which then create additional scenarios for each included image. As an example, if 5 images are used for the 21-scenario systematic error case described above, a total of $5 \times 21 = 105$ scenarios will be evaluated in each iteration.

When the multiple images are chosen to represent random, interfractional anatomical changes, optimization is performed using treatment course scenarios consisting of image and setup error fraction scenarios as described for random setup errors above. In addition to treating the multiple images as systematic or random, they may also be chosen to represent “free breathing” (intrafractional) motion. Optimization considering intrafractional motion is performed on a dose that is the sum of equally weighted partial doses computed on the individual images and deformed to the planning image. In the absence of setup and density errors, the number of scenarios for a 4D optimized “free breathing” case is therefore always 1, although doses are computed on all included images in each iteration.

Approximate and accurate scenario doses

In scenario based robust optimization, the individual spot doses must be computed for every scenario before the optimization starts. The scenario spot doses are then used to compute the total dose in each iteration, which in turn is used to evaluate the robust function values. In RayStation there are 2 options for computing the scenario doses: the accurate and the approximate methods. When accurate scenario doses are used, the scenario spot doses are explicitly computed in the same way as the nominal spot doses in a non-robust optimization. To save computer memory (RAM), only doses within structures associated to robust optimization functions are stored. This is the most accurate method but may result in long optimization times and consume significant amounts of RAM, especially when many scenarios are evaluated and when large structures (e.g., the External ROI) are associated with a robust function. To avoid excessive optimization times, the approximate option should be used. The scenario spot doses are then determined from the nominal spot doses by means of linear interpolation of the neighboring spots in the same energy layer (setup errors), and interpolation from neighboring spots with higher and lower energy (density errors).²⁶ To minimize the interpolation error, additional spot doses are computed in areas (position and range/energy) where nominal spots are absent or too sparse.²¹ These auxiliary spots are only used in the approximate

dose calculation and are not included in the spot map of the plan. Even though the auxiliary spots are stored in the spot cache, new spot doses must often be computed when robust plans using Approximate scenario doses are optimized a second time. This is because the spot filtering alters the spot pattern so that new auxiliary spots have to be computed in subsequent optimizations.

For RayStation robust plans employing a collimator, accurate scenario dose computation is mandatory. This is motivated by the fact that the influence of the collimator edge for different setup scenarios does not lend itself well to the interpolation technique employed by the approximate method.

Proton Arc Planning

Proton arcs can in general be divided into 2 different types: discrete and dynamic proton arcs. Discrete arcs (also known as static arcs) employ step-and-shoot delivery over a large number of discrete gantry angles with multiple energy layers per angle, while dynamic arcs deliver the protons while rotating the gantry with one energy layer per discretized direction. From a delivery perspective, a discrete arc plan is equivalent to a normal PBS plan, but with more beams than usual. In RayStation, a discrete arc beam is defined by its start and stop gantry angles, rotation direction, as well as the number of discrete directions. As input to the optimizer, the user also specifies the number of initial energy layers to be setup over the arc beam, and the final number of energy layers in the resulting plan. Throughout the optimization, the lowest weighted energy layers are filtered out in several cycles to reach the final number of energy layers at the iteration for spot filtering (iteration ~ 100). This process aims at automatically selecting the most beneficial energy layers over all directions. The main benefit can be seen in reduction of OAR doses and consequently in NTCP values.²⁷ Note that the number of iterations for a proton arc optimization needs to be higher than for a normal PBS plan, since the energy layer selection process needs to be performed over a sufficiently large number of iterations.

Since current proton therapy systems do not yet support the delivery of multiple gantry angles in the same beam, RayStation can convert an arc plan into a conventional PBS plan in a single click. This means that the advantages of proton arc optimization could be introduced at any proton PBS facility even today. However, the multitude of beams may result in long delivery times. This could be remedied by partitioning the discrete arc plan into subplans to be delivered over different fractions.²⁸

RayStation has support for dynamic arc optimization in research versions,²⁹ and it will be available in the clinical system when the treatment machines are capable of delivering protons while rotating the gantry. Such technological development will speed up the delivery, both for dynamic and discrete arcs. With the advent of upright treatments, a natural and cost-efficient alternative for proton arc delivery is to use a fixed proton beam in combination with a rotating patient in a seated position. The algorithms for discrete and dynamic arc optimization in RayStation are equally well-suited for a rotating patient as for a rotating gantry.

Deep Learning Planning

In the traditional planning workflow, it is often necessary to update the optimization functions iteratively to achieve the clinical goals of the plan. This can sometimes be a time-consuming process and sets high demands on the planner for complicated cases. With automatic deep learning (DL) planning, this workflow is replaced by a model that can generate high quality clinical plans considerably faster. Each DL model is associated to a particular treatment site and protocol and needs to be trained on a large set of clinical plans. In addition to a reduction in time, DL planning provides

higher consistency over a patient cohort and is not dependent on individual planners.

Automatic DL planning in RayStation consists of 2 steps. First, a 3D dose distribution is predicted by analyzing selected target and at risk organ structures of the particular patient. The model may subsequently make additional adjustments to the DL predicted dose based on DVH metrics, such as increased target coverage or sparing of certain OARs.³⁰ In the second step, a deliverable plan is created based on the predicted dose. This is accomplished by a so-called dose mimicking optimization, which generates a plan with a dose as close to the predicted dose as possible. It is important to note that the dose mimicking can be robustly optimized, thus ensuring that plan robustness is not lost in the dose prediction. After the DL plan has been created, continued manual optimization is possible as an optional final step.

The dose prediction model in RayStation is based on a 3D U-Net architecture,³¹ and each model is trained on plan dose and selected ROI geometries from a significant number of clinical cases of the treatment site and protocol at hand. CT image data, or any other patient-specific data besides dose and ROIs, is not used. In the training, the model parameters are iteratively updated to minimize the difference of the output 3D doses and the plan dose distributions of the patients. The model training is currently performed at RaySearch in collaboration with partnering clinics. Note that other machine learning models than the 3D U-Net have been previously used in RayStation, e.g., a random forest model.³²

A recent study has shown that RayStation DL-generated proton plans are of similar quality as traditionally optimized clinical treatment plans,³³ and automatic DL planning for protons using RayStation is in regular clinical use for oropharyngeal cancer patients at University Medical Center Groningen in the Netherlands.³² It is approved for clinical use in countries accepting the CE marking (e.g., in Europe), as well as a few additional markets.

Postprocessing

Once a proton plan has been created, either using optimization or DL-planning, RayStation offers a wide range of tools for manual adjustments. The meterset of the plan may be automatically, or manually adjusted to fulfill the prescribed dose. Individual spots can be moved, added, or removed, and spot weights may be adjusted. Spots below a specified weight may be filtered out, although this is better handled by the optimization (see above). Aperture, and static MLC openings may be edited by using a brush tool in the beams eye view (BEV). Even the energies of energy layers can be manually edited.

Dose Computation

In this section we will describe the proton dose calculation in RayStation, including the material interpretation of the CT data, and beam model. We will only describe the MC dose engine, and not the analytical dose engine in RayStation. This is motivated by the fact that MC has proven to be superior to analytical algorithms for protons,^{34–36} and that the MC dose calculation in RayStation now often is faster than the analytical.

CT to stopping power conversion

The MC dose engine in RayStation requires a full material description of the patient in each voxel. The material composition is expressed as mass density, mass fraction of atomic elements, $w(Z)$, and mean ionization energy, I . All cross-sections used in the dose calculation, including stopping power, are then determined from these material properties. The material composition is determined from the CT image data, and from user-defined material overrides

of selected structures. For voxels covered by structures associated with a material override, this data is directly given by the properties of the override material. For the other voxels, the material information is interpreted from the CT data by the user-defined CT-to-mass-density, or CT-to-relative stopping power (RSP) calibration curves. The CT calibration curves are defined in RayPhysics (see below), but the method to determine them is outside the scope of this paper. The elemental composition and mean ionization energy of the voxels are determined by associating one of 75 CT mapping materials to each voxel. The CT mapping materials have been determined by interpolation from 16 established human tissue and metal core materials originating from the ICRU 44 and ICRP 23 reports.^{37,38} When a CT-to-mass-density curve is used the mass density of a voxel is directly given by the calibration curve, while the CT mapping material associated to the voxel is the one that is closest in mass density to the mass density of the voxel. It should be noted that the mass density of the voxel is still the one given by the calibration curve, and not by the CT mapping material. For CT-to-RSP calibration, the associated CT mapping material will be the one that is closest in RSP to that of the voxel, as given by the calibration curve. The mass density for the voxel is then determined so that the voxel RSP from the CT calibration curve is exactly reconstructed in the dose calculation for the associated CT mapping material at a proton energy defined with the CT-to-RSP calibration curve.

RayStation also supports the import of RSP image maps, which can be produced by some dual-energy CT scanners. The RSP images can be used for planning and dose calculation as any other CT image in RayStation, and several studies have shown that the range uncertainty in plans based on RSP maps can be significantly reduced compared to plans based on conventional CTs.^{39,40} These images are exported from the scanners as normal CT images and the RSP values in the images have typically been converted to integer values of similar magnitude as HU-based CT images. For these images, the CT-to-RSP calibration curve in RayStation will simply comprise 2 points that will scale the imported values to absolute RSP, and no calibration of the curve is needed.

Dose calculation on CBCT

In RayStation it is possible to compute the dose based on Cone Beam CTs (CBCTs) by the generation of synthetic CTs (sCT). This enables the evaluation of the daily delivered dose without having to record a re-CT, and may serve as a trigger for potential replanning. For the current version of RayStation, an sCT cannot be used as a planning image in a new plan, but only for evaluation.

RayStation hosts 2 algorithms for the creation of sCTs. The first, the corrected CBCT algorithm (cCBCT), aims at removing artifacts from the CBCT image and converting the CBCT intensity values to correspond to the HUs of the planning image. This is an iterative algorithm where the HU conversion is first established by correlating the intensity values in voxels of the CBCT, to corresponding voxels in the planning CT by means of deformable registration. Low-frequency artifacts in the CBCT are then filtered out and corrected for, after which the process repeats. If the CBCT has a limited field of view (FOV), the missing volumes may be deformably added to the CBCT from the planning CT. The second method is the Virtual CT method (vCT), which is a hybrid between a deformed CT and the cCBCT method. The planning CT is deformed to the CBCT geometry and mismatching low density regions (either in the planning CT or the CBCT), are replaced with values from the cCBCT. One great advantage of these methods is that they are quite general and require no additional HU calibration, or training. Both methods have recently been validated with excellent results by several clinics using different delivery systems and focusing on different body sites.^{41–43}

Beam model

The phase space of the proton beam is modeled by an energy spectrum in combination with a bivariate Gaussian distribution that describes the spatial-angular phase space of the beam.^{44,45} The energy spectrum is discretized in bins with an energy dependent width corresponding to 0.2 mm range in water. In addition, the output of the system is modeled by a factor that relates the number of protons that shall be simulated to the dose monitor signal, where the output optionally can be defined to relate to physical dose, or to 1.1 scaled dose. If physical dose is selected, a constant factor RBE-model must be chosen for plans using that beam model. The phase space and output parameters for several nominal beam energies are determined in the beam modeling and are stored in the beam model. When the dose engine calls for the phase space of an arbitrary energy, the parameters are determined from the stored energies by linear interpolation. The interpolated energy spectrum is determined by a weighted average of neighboring energy spectra in relation to the mean energy of the spectra. It should be noted that the interpolation also works for systems with discrete energies, which means that all energies do not have to be included in the beam model. The beam model also stores the virtual source axis distance (VSAD) that may differ in the X and Y scanning directions but any variation with energy is not accounted for.

The beam modeling that extracts the phase space parameters from measurements is performed in a separate RayStation application named RayPhysics. In the beam modeling, the energy spectra are determined from measured integrated depth dose curves in water (IDDs) by a least-squares fit of precalculated mono energetic IDDs, simulated using a generic spot size and integrated laterally to a radius matching the size of the Bragg peak chamber used in the measurements. This method has been demonstrated⁴⁶ to perfectly compensate for the well-known problem of missing dose due to the limited size of the used detector.^{47,48} The spatial-angular distribution moments of the bivariate Gaussian are determined from lateral profiles measured in air from at least 3 depths at each energy, while the output constants are derived from dose of single energy scanned fields measured at mid-depth in water.

Some delivery systems generate significant non-Gaussian tails in the spot, which give rise to low dose far from the central beam axis. When such systems are modeled assuming single Gaussian spot distributions, the effect is manifested as an incorrect dose level for smaller targets.⁴⁹⁻⁵¹ Although the tails may not be very Gaussian in shape, it has been shown that the output results can be significantly improved by adding a second Gaussian to the spot distributions.⁴⁹⁻⁵¹ The RayStation beam model supports the addition of a secondary Gaussian function. Several such double Gaussian RayStation beam models have been created,^{50,52,53} of which some are in clinical use.

Since the transport of protons through range shifters and apertures is explicitly included in the MC dose engine, only geometrical measures and material properties of those devices are defined in the beam model, without the need for additional measurements.

Monte Carlo transport mechanism

The RayStation Monte Carlo code transports primary protons and secondary ions (protons, deuterons, and alphas). A Class II transport method⁵⁴ is applied for the primary and secondary protons, while nuclear absorption is neglected for the secondary deuterons and alphas.

Neutral reaction products (neutrons and gammas) are not transported, but their fractions of the absorbed energy are included in the energy balance and considered to leak out. Generation and

transport of delta electrons are not considered, since the released electrons have on average a very short range compared to the size of a voxel.⁵⁵

Energy loss by electronic ionization is determined on the fly by numerical integration of the Bethe-Bloch stopping power equation, where the shell and density correction terms have been omitted since they are only of importance for energies well below and above those of interest for therapeutic protons.⁵⁶ Energy loss straggling is handled by the Bohr approximation,⁵⁶ while multiple scattering is determined using the theory of Goudsmit-Saunderson.^{57,58}

The transport mechanics, *i.e.*, the method of propagating ions through the discretized patient and the incorporation of the electromagnetic processes (ionization energy loss, energy loss straggling and multiple coulomb scattering (MCS)), are deeply intertwined. The RayStation MC dose engine employs the so-called random hinge method, originally developed for electron/positron MC codes,⁵⁹ where the transport is divided into short and long steps. Ionization energy loss is evaluated for the short steps, which equals the intersection length of the voxels, while energy loss straggling and lateral deflection through MCS is evaluated for the longer steps. The deflection point (or hinge point) is randomly sampled along the longer hinge steps, whose total length corresponds to 10% of the particle kinetic energy at the beginning of the step. The hinge steps are created until the kinetic energy has fallen below 30.75 MeV, after which MCS is no longer evaluated. This MCS energy threshold is set to 5 MeV for transport in range shifters and other energy absorbing devices.⁶⁰

The modeling of non-elastic nuclear reactions is data-driven and based on a cross-section data library derived from published ICRU63 data.⁶¹ Elastic scattering of protons on hydrogen and on nuclei are included through parameterized models of the absorption probabilities and angular differential cross-sections.

For a patient, the transport grid coincides with the dose scoring grid. Beam modifiers are represented by rectilinear grids where the grid dimensions are chosen to best represent each beam modifier. The lateral resolution of block apertures depends on the size of the aperture opening, starting at 0.5 mm for large apertures and going down to 0.2 mm for aperture with an opening area of 10 cm² and smaller. The gaps between different transport grids (*e.g.*, the air gap between a range shifter and the patient) are treated as vacuum.

To comply with current prescriptions and normal tissue constraints, the RayStation MC dose engine reports physical dose as dose to a small water cavity embedded in the local medium, *i.e.*, as dose-to-water. When comparing dose from RayStation to other MC algorithms, it should be kept in mind that most general-purpose MC codes report dose-to-medium, which for hard bone tissue can be as much as 10% lower than dose-to-water.⁶² In addition to physical dose, the MC dose engine also computes LET_d in water.⁶³ LET from primary and secondary protons are included but LET from heavier fragments are not. This omission is mainly motivated by the fact that most current RBE models have been developed based on LET that only considers primary and secondary protons.⁶⁴ At the end of a proton's path, the LET integral is evaluated in small steps down to an energy of 1 keV. One may note that this proper handling gives a significantly higher LET_d close to end of range compared to a system that simply computes the LET_d as the energy loss divided by the step length in each voxel.

Statistics

If an MC dose is computed with too few simulated primary protons, the dose becomes noisy resulting in, *e.g.*, smeared out DVH curves.⁶⁵ The number of primary protons needed to reach acceptable accuracy is a complex matter that depends on, among other

things, the voxel size and geometrical shape of the target. A better measure than the number of simulated protons is the statistical uncertainty of the dose in the target. RayStation computes the statistical uncertainty by keeping track of the dose variation in each voxel during the calculation and then reporting the uncertainty as the mean variation in voxels with a dose higher than 50% of the maximum dose. The uncertainty is computed per beam, which means that total statistical uncertainty in the target will be smaller when several beams overlap. An MC dose in RayStation is only considered to be clinical if the statistical uncertainty of each beam dose falls below a user-defined threshold.

When a final dose is computed, the statistics can either be controlled by providing the desired uncertainty per beam, or by explicitly stating the number of simulated protons per spot. In the latter case, it is rather the mean number of protons per spot that is used, since the actual number will be proportional to the spot weight. In an optimization, only the number of protons per spot is available for user control. All spots will be computed with the same number of simulated protons, since the weights of the spots are not known prior to optimization. This has 2 consequences: (1) the simulated protons per spot in an optimization must be significantly higher than for a final dose, and (2) it is not possible to compute the statistical uncertainty of the dose for an optimized plan without final dose. Therefore, the dose after an optimization is always labeled as approximate.

GPU implementation and performance

The RayStation MC dose engine is an in-house developed product, tailored to the needs of proton therapy planning. Computational efficiency was a focus from the start, and even though the first version (released in 2016) was implemented to run on a CPU, the computational speed, both for optimization and final dose, was considered fast enough for daily clinical practice.^{35,66} An internal survey revealed that a grand majority of the existing RayStation proton clinics had switched from the analytical to the RayStation MC dose engine within one year from this first clinical release. In 2020, the RayStation MC algorithm was migrated to run on graphics cards (GPUs), which further increased the already fast dose engine by a factor of 10–20.⁶⁷ With these GPU accelerated calculations, most final dose computations have been shown to complete in 3–7 seconds, and it was even noted that a large fraction of this time was spent setting up the calculation rather than simulating the actual proton transport.⁶⁷ Considering these calculation times, it is clear that the less accurate analytical dose algorithm has played out its role, and that all dose calculations (optimization, final dose, and robust evaluation) may be done using the MC dose engine in RayStation.

Validation

The RayStation MC dose engine for proton PBS planning has now been around for almost 7 years and has been clinically implemented at more than 60 clinics. Furthermore, more than 35 peer reviewed articles have been published to date, covering various aspects of the RayStation proton MC dose engine including validating dose distributions of simple plans and QA plans in water,^{35,46,68,69} measurements in various inhomogeneous^{66,70,71} and anthropomorphic phantoms^{35,72,73} including the use of animal tissue,^{74,75} the effects of range shifters,^{36,74,76} apertures,^{60,77–80} and MLCs using both static⁸¹ and dynamic^{3,82} collimation. In addition, dose computations for SRS and ocular treatments using both standard beam lines, as well as specialized PBS beam lines where the range shifter has been positioned as far as 70 cm upstream of the collimating aperture have also been validated.^{60,77,78,80} Most studies have

been conducted by comparing the RayStation dose to measurements, but several publications have also compared the RayStation MC dose to the dose of general-purpose MC algorithms such as Geant4/TOPAS,^{43,83–85} FLUKA,⁶⁸ and MCsquare.⁸⁶ The outcome for a grand majority of these studies is very favorable and no clear trend can be seen which would suggest a significant weakness in the RayStation MC dose engine.

Plan Evaluation

The Plan evaluation module in RayStation includes a flexible workspace to evaluate the resulting plan from the plan generation process. In addition to evaluation of nominal plan doses, it is possible to create and evaluate dose distributions related to the original plan: perturbed doses with shifts in density and patient setup (translations and rotations), doses on additional images, deformed doses, summed doses, and any custom dose that the user can populate via scripting. The latter possibility opens for the creation of, e.g., non-constant RBE doses, using the physical dose and LET_d distribution as input.⁸⁷ All plan and evaluation doses are displayed in a dose tree with the plan doses as root nodes. If the beam model was commissioned with an explicit constant RBE model, physical dose will also be present in the leaf nodes. If LET_d has been computed in conjunction with final dose, it will also appear as a leaf and can be inspected side-by-side with dose for the same plan. In any of the patient planes it is possible to draw lines to visualize 1D dose and LET_d distributions.

The plan evaluation module also hosts a dedicated workspace for comprehensive analysis of the robustness of a treatment plan. The workspace provides batch computation of multiple perturbed scenarios, defined by a set of patient and density shifts, as well as acquired or simulated image sets. The scenario doses are presented in 2D views, DVH clusters and composite clinical goals. The clinical goals list shows the percentage of passed scenarios, as well as values for the current scenario and the worst scenario. Additionally, 2D views and clinical goal evaluation of voxel-wise minimum and maximum aggregate doses are included. These distributions have been shown to provide a useful link to previous experience from photon-based PTV planning.⁸⁸

Scripting

RayStation supports scripting using the CPython programming language. With a few exceptions, almost all information contained in the RayStation database is accessible, and most functions in the RayStation UI may also be executed through the scripting interface. This enables the user to expand the standard functionality of RayStation with new features for extended automation and data analysis. External applications, like Excel or secondary dose computation codes, can also be launched and supplied with RayStation data through a script. By using the .NET framework, or a Python plugin, it is possible to write UI components, allowing for dynamic user input and presentation of data. Scripts can be recorded from interactions with the RayStation UI, and this is often a good way to get started. Scripts can also be written directly in the RayStation script editor, but for more extended scripts the use of a dedicated Python IDE is recommended. Note that a few functions are not accessible through scripting, such as the plan approval, which requires user interaction to be deemed safe.

The possibilities with the RayStation scripting are practically endless. A complete list of implemented RayStation scripts would surely be almost infinitely long, but a few interesting topics and examples include: complete plan generation (including image import, structure definition, optimization and final dose calculation), interplay dose tracking based on machine log files,⁸⁹ customized

plan feasibility and sanity checks, tailored plan reports in Excel/Word/pdf, computation/extraction of IDD and dose profiles, gamma analysis, automatic detection and generation of markers and clips, and automated dose validation.

Summary and Outlook

Since its first clinical use for proton therapy in 2014, RayStation has established itself as a gold standard for proton TPSs, and is, as of 2023, selected by a majority of new proton centers. Driving factors behind this success are support for a multitude of different treatment machines, a rich set of functionalities, ease of use, high computation speed and high accuracy. The RayStation system is further future-proofed by an architecture that facilitates rapid adaptations to new planning techniques and new machine models and features.

Over the past 15 years, proton therapy has evolved from mainly being performed using broad beam techniques to almost exclusively being delivered by PBS, and it is with PBS that proton therapy has grown to what it is today. The technological development has been rapid, not in the least when it comes to TPS advancements. Important examples here include robust optimization and near-instant Monte Carlo dose calculations, features that were introduced to clinical practice through RayStation.

We foresee that technological development, combined with radiobiological advances, will result in a more personalized radiotherapy where protons and other modalities are used in an optimal way, both in terms of treatment outcome and resource management. LET-driven optimization, which can move high LET from OARs into target volumes, has just recently become available in a clinical setting. In the near future, co-optimization of variable RBE-weighted dose in combination with physical dose will have the potential to further widen the therapeutic window. We also expect to see daily online adaptation scenarios becoming part of routine workflows, a necessary development for decreasing the relatively large treatment margins used today, and thereby reaching the full potential of proton therapy. This will be powered by dose tracking and fast replanning techniques, with tight connections to the delivery system. Deep learning auto-planning has already been available for a few years but has not yet become widely spread, partly due to complicated regulatory situations in some markets, and partly due to the limited availability of DL models. However, thanks to the increased treatment planning efficiency and consistency of DL auto-planning, we are convinced that the interest for implementing DL planning in the clinical workflow will grow rapidly.

We further follow the development of more advanced delivery techniques such as dynamic arcs, where high dose conformality and a favorable LET distribution is combined with fast delivery. The support for compact delivery systems using fixed beamlines and seated patients will contribute to make proton therapy cheaper and more widely available. Other emerging techniques are spatial fractionation with or without the use of dedicated collimators and, on longer term, delivery of conformal or transmission FLASH plans. All these topics are focus areas at RaySearch and are available for exploration and clinical trials in research versions of RayStation and will, if clinical efficacy is demonstrated, become available in future clinical releases.

Conflicts of interest

All authors of this paper are employees of RaySearch Laboratories AB, Stockholm, Sweden

Acknowledgments

The authors are grateful to all RaySearch employees, including physicists, software developers, researchers, service and sales personnel, who over the past 16 years have contributed to shape RayStation for proton planning to what it is today. We are of course also eternally grateful to all RayStation clinics, both for the trust they have shown by selecting RayStation for treating their patients, and for the invaluable feedback they have provided over the years.

References

1. The total number of proton centers has been determined from the PTCOG web page (www.ptcog.site) in September 2023.
2. Bodensteiner, D. RayStation: external beam treatment planning system. *Med. Dosim.* **43**(2):168–76; 2018. doi:10.1016/j.meddos.2018.02.013.
3. Kang, M.; Pang, D. Commissioning and beam characterization of the first gantry-mounted accelerator pencil beam scanning proton system. *Med. Phys.* **47**(8):3496–510; 2020. doi:10.1002/mp.13972.
4. O'Grady, F.; Janson, M.; Rao, A.D.; et al. The use of a mini-ridge filter with cyclotron-based pencil beam scanning proton therapy. *Med. Phys.* **50**(4):1999–2008; 2023. doi:10.1002/mp.16254.
5. Fracchiolla, F.; Fellin, F.; Innocenzi, M.; et al. A pre-absorber optimization technique for pencil beam scanning proton therapy treatments. *Phys. Med.* **57**:145–52; 2019. doi:10.1016/j.ejmp.2018.12.014.
6. Engwall, E.; Glimelius, L.; Hynning, E. Effectiveness of different rescanning techniques for scanned proton radiotherapy in lung cancer patients. *Phys. Med. Biol.* **63**(9):095006; 2018. doi:10.1088/1361-6560/aabb7b.
7. Nocedal, J.; Wright, S.J.. *Numerical Optimization*. Springer; 1999.
8. Gill, P.E.; Murray, W.; Saunders, M.A. SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Rev.* **47**(1):99–131; 2005. doi:10.1137/S0036144504446096.
9. Fourkal, E.; Velchev, I.; Fan, J.; et al. Energy optimization procedure for treatment planning with laser-accelerated protons. *Med. Phys.* **34**(2):577–84; 2007. doi:10.1118/1.2431424.
10. Lomax, T.. *Intensity modulated proton therapy. Proton and Charge Particle Radiotherapy*, 2008. Lippincott Williams & Wilkins; 2008.
11. Oelfke, U.; Bortfeld, T. Inverse planning for photon and proton beams. *Med. Dosim.* **26**(2):113–24; 2001. doi:10.1016/S0958-3947(01)00057-7.
12. Fredriksson, A.. *Robust Optimization of Radiation Therapy Accounting for Geometrical Uncertainty*. KTH Royal Institute of Technology; 2013.
13. Underwood, T.S.A.; Grassberger, C.; Bass, R.; et al. Asymptomatic late-phase radiographic changes among chest-wall patients are associated with a proton RBE Exceeding 1.1. *Int. J. Radiat. Oncol. Biol. Phys.* **101**(4):809–19; 2018. doi:10.1016/j.ijrobp.2018.03.037.
14. Zhu, X.R.; Sahoo, N.; Zhang, X.; et al. Intensity modulated proton therapy treatment planning using single-field optimization: the impact of monitor unit constraints on plan quality: Scanning proton beam treatment planning. *Med. Phys.* **37**(3):1210–19; 2010. doi:10.1118/1.3314073.
15. Menzel, H. G. Journal of the ICRU **10**(1); 2010 Report 83. doi:10.1093/jicru/ndq001.
16. Lomax, A.J. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions. *Phys. Med. Biol.* **53**(4):1043–56; 2008. doi:10.1088/0031-9155/53/4/015.
17. Fredriksson, A.; Forsgren, A.; Hårdemark, B. Minimax optimization for handling range and setup uncertainties in proton therapy: minimax optimization for handling uncertainties in proton therapy. *Med. Phys.* **38**(3):1672–84; 2011. doi:10.1118/1.3556559.
18. Stuschke, M.; Kaiser, A.; Pöttgen, C.; et al. Potentials of robust intensity modulated scanning proton plans for locally advanced lung cancer in comparison to intensity modulated photon plans. *Radiother. Oncol.* **104**(1):45–51; 2012. doi:10.1016/j.radonc.2012.03.017.
19. Fredriksson, A.; Bokrantz, R. A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning: critical evaluation of worst case optimization for IMPT. *Med. Phys.* **41**(8Part1):081701; 2014. doi:10.1118/1.4883837.
20. Fredriksson, A.; Engwall, E.; Andersson, B. Robust radiation therapy optimization using simulated treatment courses for handling deformable organ motion. *Phys. Med. Biol.* **66**(4):045010; 2021. doi:10.1088/1361-6560/abd591.
21. Unkelbach, J.; Bortfeld, T.; Martin, B.C.; et al. Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning. *Med. Phys.* **36**(1):149–63; 2009. doi:10.1118/1.3021139.
22. Noufal, M.P.; Sharma, S.D.; Krishnan, G.; et al. Robustly optimized hybrid intensity-modulated proton therapy for craniospinal irradiation. *J. Cancer Res. Ther.* **18**(6):1597–603; 2022. doi:10.4103/jcrt.jcrt_740_20.
23. Ge, S.; Wang, X.; Liao, Z.; et al. Potential for improvements in robustness and optimality of intensity-modulated proton therapy for lung cancer with 4-dimensional robust optimization. *Cancers* **11**(1):35; 2019. doi:10.3390/cancers11010035.
24. Feng, H.; Shan, J.; Ashman, J.B.; et al. Technical NOTE: 4D robust optimization in small spot intensity-modulated proton therapy (IMPT) for distal esophageal carcinoma. *Med. Phys.* **48**(8):4636–47; 2021. doi:10.1002/mp.15003.

25. Zhu, M.; Kaiser, A.; Mishra, M.V.; et al. Multiple computed tomography robust optimization to account for random anatomic density variations during intensity modulated proton therapy. *Adv. Radiat. Oncol.* **5**(5):1022–31; 2019. doi:10.1016/j.adro.2019.12.003.
26. Unkelbach, J.; Alber, M.; Bangert, M.; et al. Robust radiotherapy planning. *Phys. Med. Biol.* **63**(22):22TR02; 2018. doi:10.1088/1361-6560/aac659.
27. De Jong, B.A.; Korevaar, E.W.; Maring, A.; et al. Proton arc therapy increases the benefit of proton therapy for oropharyngeal cancer patients in the model based clinic. *Radiother. Oncol.* **184**:109670; 2023. doi:10.1016/j.radonc.2023.109670.
28. Engwall, E.; Marthin, O.; Wase, V.; et al. Partitioning of discrete proton arcs into interlaced subplans can bring proton arc advances to existing proton facilities. *Med. Phys.* mp.16617 2023. doi:10.1002/mp.16617.
29. Engwall, E.; Battinelli, C.; Wase, V.; et al. Fast robust optimization of proton PBS arc therapy plans using early energy layer selection and spot assignment. *Phys. Med. Biol.* **67**(6):065010; 2022. doi:10.1088/1361-6560/ac55a6.
30. Borderias-Villarreal, E.; Huet Dastarac, M.; Barragán-Montero, A.M.; et al. Machine learning-based automatic proton therapy planning: Impact of post-processing and dose-mimicking in plan robustness. *Med. Phys.* **50**(7):4480–90; 2023. doi:10.1002/mp.16408.
31. Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S.S.; et al. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin S., Joskowicz L., Sabuncu M.R., Unal G., Wells W., editors. *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, 9901. Springer International Publishing; 2016. p. 424–32. Lecture Notes in Computer Science. doi:10.1007/978-3-319-46723-8_49.
32. van Bruggen, I.G.; Huiskes, M.; de Vette, S.P.M.; et al. Automated robust planning for IMPT in oropharyngeal cancer patients using machine learning. *Int. J. Radiat. Oncol.* **115**(5):1283–90; 2023. doi:10.1016/j.ijrobp.2022.12.004.
33. Maes, D.; Holmstrom, M.; Helander, R.; et al. Automated treatment planning for proton pencil beam scanning using deep learning dose prediction and dose-mimicking optimization. *J. Appl. Clin. Med. Phys.* e14065 2023. doi:10.1002/acm2.14065.
34. Taylor, P.A.; Kry, S.F.; Followill, D.S. Pencil beam algorithms are unsuitable for proton dose calculations in lung. *Int. J. Radiat. Oncol.* **99**(3):750–6; 2017. doi:10.1016/j.ijrobp.2017.06.003.
35. Saini, J.; Maes, D.; Egan, A.; et al. Dosimetric evaluation of a commercial proton spot scanning Monte-Carlo dose algorithm: comparisons against measurements and simulations. *Phys. Med. Biol.* **62**(19):7659–81; 2017. doi:10.1088/1361-6560/aa82a5.
36. Shirey, R.J.; Wu, H.T. Quantifying the effect of air gap, depth, and range shifter thickness on TPS dosimetric accuracy in superficial PBS proton therapy. *J. Appl. Clin. Med. Phys.* **19**(1):164–73; 2018. doi:10.1002/acm2.12241.
37. . ICRU 44, *Tissue Substitutes in Radiation Dosimetry and Measurement, Report 44 of the International Commission on Radiation Units and Measurements*; 1989.
38. . ICRP 23, *Report of the Task Group on Reference Man (ICRP 23): Anatomical, Physiological and Metabolic Characteristics*; 1975.
39. Peters, N.; Wohlfahrt, P.; Hofmann, C.; et al. Reduction of clinical safety margins in proton therapy enabled by the clinical implementation of dual-energy CT for direct stopping-power prediction. *Radiother. Oncol.* **166**:71–8; 2022. doi:10.1016/j.radonc.2021.11.002.
40. Sarkar, V.; Paxton, A.; Su, F.; et al. An evaluation of the use of DirectSPR images for proton planning in the RayStation treatment planning software. *J. Appl. Clin. Med. Phys.* **24**(5):e13900; 2023. doi:10.1002/acm2.13900.
41. Taasti, V.T.; Hattu, D.; Peeters, S.; et al. Clinical evaluation of synthetic computed tomography methods in adaptive proton therapy of lung cancer patients. *Phys. Imaging Radiat. Oncol.* **27**:100459; 2023. doi:10.1016/j.phro.2023.100459.
42. Reiners, K.; Dagan, R.; Holtzman, A.; et al. CBCT-based dose monitoring and adaptive planning triggers in head and neck PBS proton therapy. *Cancers* **15**(15):3881; 2023. doi:10.3390/cancers15153881.
43. Chang, C.W.; Nilsson, R.; Andersson, S.; et al. An optimized framework for cone-beam computed tomography-based online evaluation for proton therapy. *Med Phys* **50**(9):5375–86; 2023. doi:10.1002/mp.16625.
44. Schaffner, B.; Pedroni, E.; Lomax, A. Dose calculation models for proton treatment planning using a dynamic beam delivery system: an attempt to include density heterogeneity effects in the analytical dose calculation. *Phys. Med. Biol.* **44**(1):27–41; 1999. doi:10.1088/0031-9155/44/1/004.
45. Grevillot, L.; Bertrand, D.; Dessy, F.; et al. A Monte Carlo pencil beam scanning model for proton treatment plan simulation using GATE/GEANT4. *Phys. Med. Biol.* **56**(16):5203–19; 2011. doi:10.1088/0031-9155/56/16/008.
46. Azcona, J.D.; Aguilar, B.; Perales, Á.; et al. Commissioning of a synchrotron-based proton beam therapy system for use with a Monte Carlo treatment planning system. *Radiat. Phys. Chem.* **204**:110708; 2023. doi:10.1016/j.radphyschem.2022.110708.
47. Sawakuchi, G.O.; Mirkovic, D.; Perles, L.A.; et al. An MCNPX Monte Carlo model of a discrete spot scanning proton beam therapy nozzle. *Med. Phys.* **37**(9):4960–70; 2010. doi:10.1118/1.3476458.
48. Bäumer, C.; Koska, B.; Lambert, J.; et al. Evaluation of detectors for acquisition of pristine depth-dose curves in pencil beam scanning. *J. Appl. Clin. Med. Phys.* **16**(6):151–63; 2015. doi:10.1120/jacmp.v16i6.5577.
49. Zhu, X.R.; Poenisch, F.; Lii, M.; et al. Commissioning dose computation models for spot scanning proton beams in water for a commercially available treatment planning system. *Med. Phys.* **40**(4):041723; 2013. doi:10.1118/1.4798229.
50. Kugel, F.; Wulff, J.; Bäumer, C.; et al. Validating a double Gaussian source model for small proton fields in a commercial Monte-Carlo dose calculation engine. *Z. Für. Med. Phys.*:S0939388922001325 2022. doi:10.1016/j.zemedi.2022.11.011.
51. Harms, J.; Chang, C.; Zhang, R.; et al. Nuclear halo measurements for accurate prediction of field size factor in a Varian ProBeam proton PBS system. *J. Appl. Clin. Med. Phys.* **21**(1):197–204; 2020. doi:10.1002/acm2.12783.
52. Patel, D.; Poenisch, F.; Sahoo, N.; et al. *Evaluation of a pencil beam scanning proton dose model for a commercial treatment planning system*. Poster presentation at the AAPM annual meeting (Nashville, TN); 2018.
53. Ecker, S.; Ellerbrock, M.; Ackermann, B.; et al. *Commissioning of the treatment planning system RayStation at the Heidelberg Ion-Beam Therapy Center*. Poster presentation at the 58th annual PTCOG meeting (Manchester, UK); 2019.
54. Berger, M.J. Monte Carlo calculation of the penetration and diffusion of fast charged particles. *Methods Comput. Phys.* **1**:135–215; 1963.
55. Pfuhl, T.; Horst, F.; Schuy, C.; et al. Dose build-up effects induced by delta electrons and target fragments in proton Bragg curves—measurements and simulations. *Phys. Med. Biol.* **63**(17):175002; 2018. doi:10.1088/1361-6560/aad8fc.
56. ICRU Report 49, *Stopping Power and Ranges for Protons and Alpha Particles*.
57. Goudsmit, S.; Saunderson, J.L. Multiple scattering of electrons. *Phys. Rev.* **57**(1):24–9; 1940. doi:10.1103/PhysRev.57.24.
58. Goudsmit, S.; Saunderson, J.L. Multiple scattering of electrons. II. *Phys. Rev.* **58**(1):36–42; 1940. doi:10.1103/PhysRev.58.36.
59. Kling, A.; Barão, F.; Nakagawa, M.; et al. *Advanced MonteCarlo for radiation physics, particle transport simulation and applications*. In: *Proceedings of the Monte Carlo 2000 Conference*. Springer Science & Business Media; 2014 Lisbon 23–26 October 2000.
60. Knäusl, B.; Langgartner, L.; Stock, M.; et al. Requirements for dose calculation on an active scanned proton beamline for small, shallow fields. *Phys. Med.* **113**:102659; 2023. doi:10.1016/j.ejmp.2023.102659.
61. ICRU Report 63, *Nuclear Data for Neutron and Proton Radiotherapy and for Radiation Protection – ICRU*.
62. Paganetti, H. Dose to water versus dose to medium in proton beam therapy. *Phys. Med. Biol.* **54**(14):4399–421; 2009. doi:10.1088/0031-9155/54/14/004.
63. Wilkens, J.J.; Oelfke, U. Analytical linear energy transfer calculations for proton therapy. *Med. Phys.* **30**(5):806–15; 2003. doi:10.1118/1.1567852.
64. Kaltholm, F.; Grzanka, L.; Traneus, E.; et al. A systematic review on the usage of averaged LET in radiation biology for particle therapy. *Radiother. Oncol.* **161**:211–21; 2021. doi:10.1016/j.radonc.2021.04.007.
65. Kim, H.; Chung, K.; Han, Y.; et al. Searching for optimized selection of Monte Carlo dose calculation parameters for scanned beam proton therapy in RayStation. *J. Korean Phys. Soc.* 2023. doi:10.1007/s40042-023-00897-5.
66. Carlino, A.; Böhlen, T.; Vatnitsky, S.; et al. Commissioning of pencil beam and Monte Carlo dose engines for non-isocentric treatments in scanned proton beam therapy. *Phys. Med. Biol.* **64**(17):17NT01; 2019. doi:10.1088/1361-6560/ab3557.
67. Fracchiolla, F.; Engwall, E.; Janson, M.; et al. Clinical validation of a GPU-based Monte Carlo dose engine of a commercial treatment planning system for pencil beam scanning proton therapy. *Phys. Med.* **88**:226–34; 2021. doi:10.1016/j.ejmp.2021.07.012.
68. Molinelli, S.; Russo, S.; Magro, G.; et al. Impact of TPS calculation algorithms on dose delivered to the patient in proton therapy treatments. *Phys. Med. Biol.* **64**(7):075016; 2019. doi:10.1088/1361-6560/ab044d.
69. Langner, U.W.; Mundis, M.; Strauss, D.; et al. A comparison of two pencil beam scanning treatment planning systems for proton therapy. *J. Appl. Clin. Med. Phys.* **19**(1):156–63; 2018. doi:10.1002/acm2.12235.
70. Ruangchan, S.; Knäusl, B.; Fuchs, H.; et al. Experimental benchmarking of RayStation proton dose calculation algorithms inside and outside the target region in heterogeneous phantom geometries. *Phys. Med.* **76**:182–93; 2020. doi:10.1016/j.ejmp.2020.07.010.
71. Rightetto, R.; Clemens, L.P.; Lorentini, S.; et al. Accurate proton treatment planning for pencil beam crossing titanium fixation implants. *Phys. Med.* **70**:28–38; 2020. doi:10.1016/j.ejmp.2020.01.003.
72. Widesott, L.; Lorentini, S.; Fracchiolla, F.; et al. Improvements in pencil beam scanning proton therapy dose calculation accuracy in brain tumor cases with a commercial Monte Carlo algorithm. *Phys. Med. Biol.* **63**(14):145016; 2018. doi:10.1088/1361-6560/aac279.
73. Kostjukhina, N.; Palmans, H.; Stock, M.; et al. Dynamic lung phantom commissioning for 4D dose assessment in proton therapy. *Phys. Med. Biol.* **64**(23):235001; 2019. doi:10.1088/1361-6560/ab5132.
74. Schreuder, A.N.; Bridges, D.S.; Rigsby, L.; et al. Validation of the RayStation Monte Carlo dose calculation algorithm using a realistic lung phantom. *J. Appl. Clin. Med. Phys.* **20**(12):127–37; 2019. doi:10.1002/acm2.12777.
75. Schreuder, A.N.; Bridges, D.S.; Rigsby, L.; et al. Validation of the RayStation Monte Carlo dose calculation algorithm using realistic animal tissue phantoms. *J. Appl. Clin. Med. Phys.* **20**(10):160–71; 2019. doi:10.1002/acm2.12733.
76. Rana, S.; Samuel, E.J.J. Measurements of in-air spot size of pencil proton beam for various air gaps in conjunction with a range shifter on a ProteusPLUS PBS dedicated machine and comparison to the proton dose calculation algorithms. *Australas. Phys. Eng. Sci. Med.* **42**(3):853–62; 2019. doi:10.1007/s13246-019-00772-3.
77. Bäumer, C.; Janson, M.; Timmermann, B.; et al. Collimated proton pencil-beam scanning for superficial targets: impact of the order of range shifter and aperture. *Phys. Med. Biol.* **63**(8):085020; 2018. doi:10.1088/1361-6560/aab79c.
78. Holmes, J.; Shen, J.; Shan, J.; et al. Technical note: Evaluation and second check of a commercial Monte Carlo dose engine for small-field apertures in pencil beam scanning proton therapy. *Med. Phys.* **49**(5):3497–506; 2022. doi:10.1002/mp.15604.
79. Regmi, R.; Maes, D.; Nevitt, A.; et al. Treatment of ocular tumors through a novel applicator on a conventional proton pencil beam scanning beamline. *Sci. Rep.* **12**(1):4648; 2022. doi:10.1038/s41598-022-08440-5.

80. Saini, J.; Maes, D.; Regmi, R.; et al. Improved lateral penumbra for proton ocular treatments on a general-purpose spot scanning beamline. *Phys Med* **107**:102551; 2023. doi:[10.1016/j.ejmp.2023.102551](https://doi.org/10.1016/j.ejmp.2023.102551).
81. Tominaga, Y.; Sakurai, Y.; Miyata, J.; et al. Validation of pencil beam scanning proton therapy with multi-leaf collimator calculated by a commercial Monte Carlo dose engine. *J. Appl. Clin. Med. Phys.* **23**(12):e13817; 2022. doi:[10.1002/acm2.13817](https://doi.org/10.1002/acm2.13817).
82. Grewal, H.S.; Ahmad, S.; Jin, H. Characterization of penumbra sharpening and scattering by adaptive aperture for a compact pencil beam scanning proton therapy system. *Med. Phys.* **48**(4):1508–19; 2021. doi:[10.1002/mp.14771](https://doi.org/10.1002/mp.14771).
83. Liu, C.; Zhang, Y.; Li, Z.; et al. Commissioning and validation of TOPAS beam model for IBA Proteus-ONE at UFHPTI. *Radiat. Phys. Chem.* **180**:109256; 2021. doi:[10.1016/j.radphyschem.2020.109256](https://doi.org/10.1016/j.radphyschem.2020.109256).
84. Verbeek, N.; Wulff, J.; Bäumer, C.; et al. Single pencil beam benchmark of a module for Monte Carlo simulation of proton transport in the PENELOPE code. *Med. Phys.* **48**(1):456–76; 2021. doi:[10.1002/mp.14598](https://doi.org/10.1002/mp.14598).
85. Verbeek, N.; Wulff, J.; Janson, M.; et al. Experiments and Monte Carlo simulations on multiple Coulomb scattering of protons. *Med. Phys.* **48**(6):3186–99; 2021. doi:[10.1002/mp.14860](https://doi.org/10.1002/mp.14860).
86. Liu, C.; Ho, M.W.; Park, J.; et al. Fast MCsquare-based independent dose verification platform for pencil beam scanning proton therapy. *Technol. Cancer Res. Treat.* **20**:153303382110330; 2021. doi:[10.1177/15330338211033076](https://doi.org/10.1177/15330338211033076).
87. Rørvik, E.; Fjæra, L.F.; Dahle, T.J.; et al. Exploration and application of phenomenological RBE models for proton therapy. *Phys. Med. Biol.* **63**(18):185013; 2018. doi:[10.1088/1361-6560/aad9db](https://doi.org/10.1088/1361-6560/aad9db).
88. Korevaar, E.W.; Habraken, S.J.M.; Scandurra, D.; et al. Practical robustness evaluation in radiotherapy: a photon and proton-proof alternative to PTV-based plan evaluation. *Radiother. Oncol.* **141**:267–74; 2019. doi:[10.1016/j.radonc.2019.08.005](https://doi.org/10.1016/j.radonc.2019.08.005).
89. Meijers, A.; Jakobi, A.; Stützer, K.; et al. Log file-based dose reconstruction and accumulation for 4D adaptive pencil beam scanned proton therapy in a clinical treatment planning system: Implementation and proof-of-concept. *Med. Phys.* **46**(3):1140–9; 2019. doi:[10.1002/mp.13371](https://doi.org/10.1002/mp.13371).