Near real-time adaptive planning for intra-fraction motion in proton therapy: a beamlet-free Monte Carlo approach

Valentine Dormal[1], Danah Pross[1], Ana M. Barragán Montero[1], John A. Lee[1], Edmond Sterpin[1,2,3]

[1] IREC/MIRO, UCLouvain, Brussels, Belgium

[2] Department of Oncology, KU Leuven, Leuven, Belgium

[3] Particle Therapy Interuniversity Center Leuven - PARTICLE, Leuven, Belgium

Abstract

Objective: Intra-fraction anatomical variations can compromise treatment quality, and while real-time adaptation of proton therapy plans during delivery offers a potential solution, its implementation remains challenging, particularly when integrating accurate Monte Carlo dose engines into a sufficiently fast optimization workflow. Conventional optimization methods rely on pre-computed dose influence matrices for each spot (i.e., beamlets), which are time-consuming to calculate and limit the feasibility of real-time adaptation. This study introduces an intra-fraction adaptive framework based on a beamlet-free optimizer that dynamically updates the treatment plan during treatment delivery.

Approach: To eliminate the need for precomputed beamlets, the beamlet-free optimizer integrates the Monte Carlo dose calculation with the optimization process. The proposed adaptive method accounts for delivered dose and anatomical changes and allows for re-delivery of spots when necessary. The optimizer leverages prior plan information from an initial optimization and integrates a plan conformity term to balance plan stability and adaptability. The framework was tested on synthetic phantoms under various motion and deformation scenarios and on three clinical cases (two lung, one cervix) and compared against (1) the initial plan on the planning CT, (2) the initial plan delivered on the evolving anatomy without adaptation, and (3) plans adapted using a conventional beamlet-based optimizer.

Main results: Across all test cases, non-adapted plans failed to meet clinical objectives, while both adaptive approaches restored target coverage and organ-at-risk sparing. The beamlet-free optimizer achieved comparable dosimetric quality to the beamlet-based method while reducing computation times by up to 7.7×. Time gains increased with plan complexity and total number of spots.

Significance: The proposed beamlet-free adaptive workflow enables intra-fraction updates during treatment delivery with Monte Carlo-level accuracy, allowing flexible spot re-delivery while maintaining dosimetric accuracy.

Keywords: near real-time optimization, intra-fraction adaptation, beamlet-free, proton therapy, treatment planning

I. Introduction

Accurately accounting for anatomical variations during treatment delivery is essential for achieving the full potential of intensity-modulated proton therapy (IMPT). Even small intra-fractional changes in patient anatomy can significantly affect the delivered dose distribution, compromising both target coverage and organ-at-risk (OAR) sparing. These include breathing-related motion, but also spontaneous organ motion or shape variations due to bowel movement, bladder or rectal filling, gastrointestinal gas, or gradual target drift over time (Batista et al., 2018; Berger et al., 2021; Heijkoop et al., 2015; Kumagai et al., 2009; Mostafaei et al., 2018; Nejad-Davarani et al., 2019; Rusu et al., 2023).

Online adaptive therapy has demonstrated its ability to compensate for inter-fractional anatomical variations by re-optimizing the plan based on daily imaging. However, extending this concept to real-time intra-fraction adaptation remains largely unexplored, particularly in proton therapy. In conventional radiotherapy, the first works were limited to motion management strategies that tracked the target to deliver the dose at the right place but did not account for how motion affected the actual delivered dose distribution. This includes gating (Keall et al., 2006; Kubo and Hill, 1996), couch repositioning (Keall et al., 2020), MLC tracking (Booth et al., 2021), or robotic linacs (Kilby et al., 2010). Recently, the improvement in intra-fractional imaging (like the MR-linac systems) and computational speed of the dose calculation and optimization, has led to the rise of real-time dose-guided adaptation (Keall et al., 2025). Such an approach does re-optimize plan parameters during delivery and could therefore substantially reduce the dosimetric impact of intra-fractional motion and anatomical deformation. A good example is provided by MR-guided real-time adaptation studies, where one popular strategy is to divide each treatment fraction into multiple sub-fractions with re-optimization between them on the latest anatomy, assuming homogeneous dose delivery within each sub-fraction (Lagerwaard et al., 2018; Snyder et al., 2024; Willigenburg et al., 2022). Other approaches were explored for real-time adaptation, including a motion-adaptive workflow for TomoTherapy (Lu et al., 2009), a constrained re-optimization for non-periodic motion in robotic SBRT(Gerlach et al., 2023), an adaptive pipeline with inter-beam replanning for free-breathing SBRT (Kontaxis et al., 2017), and a re-optimization method with a time-dependent objective function tested on water phantoms (Wiersma and Liu, 2019).

In proton therapy, most research and clinical practice in real-time intrafraction motion management still rely on strategies that do not involve dose re-optimization and are primarily limited to respiratory motion, such as rescanning or gating guided by external markers (Zhang et al., 2023). Only in the last year, a couple of groups started to explore real-time dose-guided adaptation. Gambetta et al. investigated partial adaptation for two-field prostate cancer treatments, showing that deviations detected during delivery of the first field could be

compensated by re-optimizing the second field (Gambetta et al., 2025a). However, only a single adaptation per fraction was possible, and even though only the second field was adapted, the method still required a full re-optimization of both fields followed by a dose mimicking step. Bengtsson et al. explored real-time re-optimization of proton pencil-beam scanning to mitigate interplay effects from breathing motion (Bengtsson et al., 2025). Their framework re-optimized the spot weights of each upcoming energy layer based on the delivered dose and predicted motion, assuming the availability of fast and accurate motion models. The approach was tested only on simulated breathing patterns with variable frequency, without accounting for irregular motion or varying amplitude. Both approaches restrict re-optimization to the preplanned spots and energy layers, assuming that compensation for anatomical variations can be achieved with the remaining fields. However, in the presence of large, unexpected motion and deformation, the ability to re-deliver or introduce new spots becomes essential to maintain a good treatment plan quality. Such flexibility, though, comes at the expense of increased computational burden, as it becomes challenging for the optimizer to converge to a solution with an extended solution space.
In the context of online inter-fraction adaptation, approaches that also avoid a full plan re-optimization have been proposed, combining geometric adaptation by adjusting spot positions and energies to the daily anatomy, with a subsequent re-optimization of spot weights (Botas et al., 2018; Jagt et al., 2018).

A common bottleneck with conventional proton therapy optimization methods is their reliance on pre-computed beamlets, i.e., the dose distribution calculated for each spot position. Generating these beamlets is both time-consuming and memory intensive, particularly with Monte Carlo dose engines, and limits the feasibility of real-time plan adaptation during delivery. Moreover, even in real-time adaptation, a safe treatment would require some level of robust optimization to account for residual uncertainties, such as contouring and proton range calculation errors, which further increases computational requirements.
To address these limitations, a beamlet-free optimization algorithm (Pross et al., 2024) was recently introduced in which the optimization process is directly embedded within the Monte Carlo simulation, removing the need for precomputed beamlets. This beamlet-free approach achieved a reduction in both computation time and memory usage for conventional IMPT planning, even in robust optimization settings (Pross et al., 2025), and demonstrated reduced dependency on the number of spots. For this reason, this method shows great potential to be integrated in a real-time adaptation workflow where both flexibility and speed are essential, while maintaining Monte Carlo-level accuracy.

In this study, we propose an intra-fraction adaptation framework based on this beamlet-free optimizer. The method dynamically updates the treatment plan, accounting for delivered dose

and anatomical deformations, while allowing the flexibility to re-deliver spots if needed. The feasibility of this approach was investigated through a series of phantom scenarios and three patient cases in two treatment sites. Beamlet-free plans were compared with (1) the initial reference plan, optimized and evaluated on the planning CT before treatment, (2) the non-adapted plan, corresponding to the initial plan delivered to the evolving anatomy without re-optimization, and (3) the beamlet-based re-optimized plan, adapted using a state-of-the-art beamlet-based optimizer. Both dosimetric outcome and relative computational efficiency were assessed.

II. Methods

A. Optimization framework

The intra-fraction adaptation approach proposed in this study relies on a previously developed beamlet-free optimizer (Pross et al., 2024). This algorithm iteratively simulates small batches of protons toward single spots, sampled according to a probability distribution that continuously evolves during optimization. After each batch, the impact on the total dose distribution is evaluated, and the probability of picking that spot again is adjusted through a stochastic coordinate descent process (Nesterov, 2012). Only the total dose distribution is stored and not the individual contribution of each spot (i.e., the beamlets), hence the term beamlet-free.
To enable fast replanning during treatment delivery, the optimizer was extended in three ways: (1) to account for the background dose, (2) to leverage prior information from previous optimizations, and (3) to include a stopping criterion suitable for intra-fraction adaptation.

We refer to background dose as the dose already delivered up to time *t*. Since this dose distribution is heterogeneous, the residual dose required to achieve a uniform prescribed dose within a structure varies across voxels. Consequently, structure-specific dose objectives that aim to achieve a homogeneous dose within a Region of Interest (ROI) effectively become voxel-specific when background dose is taken into account, as reflected in the updated formulation of the objective function in equation (*1*).

When performing a re-optimization at time *t* on a new anatomical image, a set of optimized spot weights $\boldsymbol{x}_{t-1}$ from a previous optimization made at time *t-1* is available, either from the initial plan prior to treatment delivery, or from a previously adapted plan if an update already occurred. Although sub-optimal, we hypothesize that these weights might still be a good starting point for the new optimization and should be leveraged as much as possible. Thus, this prior information, denoted as $\tilde{\boldsymbol{x}}_{t-1}$, is used both as an initialization for the optimization algorithm and to define a conformity term in the

objective function. This added term ensures that the newly optimized plan does not deviate unnecessarily from the prior plan, unless required to maintain plan quality in the presence of anatomical changes. The updated objective function with the conformity term and the voxel-wise specific objectives can be expressed as:

$$
\begin{aligned}
F(\boldsymbol{x}_t) = & \sum_k \frac{w_k}{|ROI_k|} \sum_{i \in ROI_k} min(0, d_i(\boldsymbol{x}_t) + D_i^{0\to t} - p_{min\_ROI_k})^2 \\
& + \sum_k \frac{w_k}{|ROI_k|} \sum_{i \in ROI_k} max(0, d_i(\boldsymbol{x}_t) + D_i^{0\to t} - p_{max\_ROI_k})^2 \\
& + \sum_k w_k \, max\left(0, \frac{1}{|ROI_k|} \sum_{i \in ROI_k} (d_i(\boldsymbol{x}_t) + D_i^{0\to t}) - p_{max\_mean\_ROI_k}\right)^2 \\
& + \lambda \|\boldsymbol{x}_t - \tilde{\boldsymbol{x}}_{t-1}\|_1
\end{aligned}
\tag{1}
$$

where $\boldsymbol{x}_t$ is the vector of weights of the candidate spots being currently optimized, $d_i(\boldsymbol{x}_t)$ is the dose at voxel *i* resulting from spot weights $\boldsymbol{x}_t$, w_k is the weight of objective k, $D_i^{0\to t}$ is the background dose at voxel *i* and at time *t*, λ is the conformity parameter controlling compliance with the prior plan, and $p_{min\_ROI_k}$, $p_{max\_ROI_k}$ and $p_{max\_mean\_ROI_k}$ are respectively the minimum, maximum and mean prescribed dose limits for an objective on ROI *k*.

Regarding optimization stopping criteria, standard approaches often rely on monitoring the objective function, but this strategy is not well suited for our workflow. First, frequent evaluation of the objective function is computationally inefficient in the beamlet-free algorithm, as it requires accessing dose values in all voxels of the scoring grid. In contrast, each optimization step only computes finite-difference approximations of the partial derivatives, avoiding the need to parse the entire voxel grid. Second, the objective function value reflects not only the optimization progress but also the statistical uncertainty inherent to the Monte Carlo simulation, since the optimization is integrated within the dose calculation. We therefore adopt a stopping criterion based on the stability of the optimized variables. The optimization is stopped when the mean

absolute difference between the current spot weights at iteration *k* and the previous *n* iterations falls below a predefined threshold ϵ :

$$\frac{1}{N}\left\|\boldsymbol{x}_t^{(k)} - \boldsymbol{x}_t^{(k-n)}\right\|_1 < \epsilon \tag{2}$$

where N is the number of spots and $\boldsymbol{x}_t^{(k)}$ are the spot weights at iteration k of the optimization performed during the adaptation triggered at time *t*. This criterion ensures that optimization terminates once parameter updates have become sufficiently small, indicating convergence. The threshold value ϵ was empirically set to 0.02 and applied consistently across all test cases.

B. Adaptive Workflow

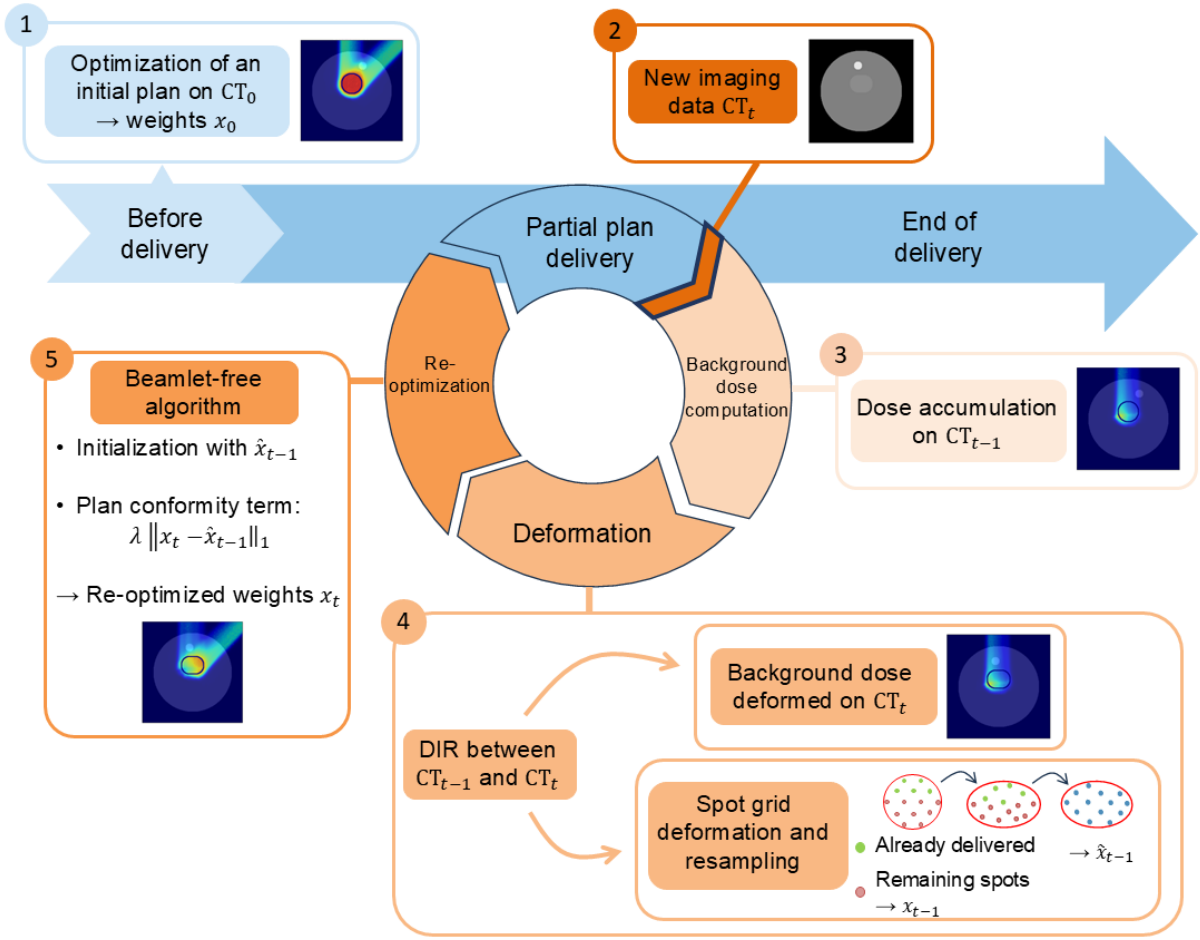

Figure 1 : *Proposed workflow for intra-fraction adaptation illustrated on a water phantom with a deforming spherical target and a moving bone. Notations: CT_t = CT at time t ; x_t = spot*

weights optimized for the anatomy of CT_t ; $\tilde{x}_{t-1}$ = spot weights optimized on the anatomy of CT_{t-1} but displaced on CT_t ; DIR = Deformable Image Registration

The overall adaptive workflow is illustrated in Figure 1. As in conventional workflows, it begins with the optimization of an initial reference plan P_0 on a planning CT scan (CT_0) prior to treatment. This P_0 plan was generated using the beamlet-free optimization algorithm (Pross et al., 2024) in its original, non-adaptive implementation. Based on this plan, treatment delivery starts and is simulated by delivering energy layers in decreasing energy order within each beam.
When a new anatomical image CT_t becomes available at time *t*, the simulated delivery is interrupted, and the intra-fraction adaptation process is triggered.

The first step is to compute the background dose. The accumulated dose from the start of treatment up to the last adaptation at time *t-1* is denoted $\mathbf{D}^{0\to t-1}$. If the current update corresponds to the first adaptation during this fraction, $\mathbf{D}^{0\to t-1}$ is empty. Otherwise, this dose map was already computed on CT_{t-1} during the previous adaptation. The dose $\mathbf{D}^{t-1\to t}$ delivered between *t-1* and *t* is then computed by simulating on CT_{t-1} a total number of particles equal to 10000 times the number of spots delivered during this time interval. This incremental dose $\mathbf{D}^{t-1\to t}$, is added to the previous background dose $\mathbf{D}^{0\to t-1}$, resulting in the updated background dose $\mathbf{D}^{0\to t}$ on CT_{t-1}.

A deformable image registration (DIR) is then performed between CT_{t-1} and CT_t. The resulting deformation vector field is used in two ways. First, it is applied on dose $D_{0\to t}$ to map it onto the current anatomy CT_t. Second, it is used to propagate prior spot information x_{t-1}, to the new geometry. For this purpose, the weights of all spots that have been already delivered are set to zero, and the remaining spot grid from CT_{t-1} is displaced using the deformation vector field. This results in an irregular grid where the hexagonal structure and the energy layers may be disrupted. A new regular hexagonal grid with an updated set of energy layers is then reconstructed to cover the target volume on CT_t using the predefined spot and layer spacings. Spot weights from the displaced grid are assigned to the nearest spot and energy layer of this new grid. These displaced spot weights, denoted $\tilde{x}_{t-1}$, are then provided to the optimizer as prior information for the update of the treatment plan, as described in the previous section.

Then, the re-optimization process begins. The beamlet-free algorithm takes as input the displaced spot weights $\tilde{x}_{t-1}$ and the updated background dose $\mathbf{D}^{0\to t}$ and outputs a set of newly optimized spot weights x_t.

To improve robustness against large anatomical deformations or variations, additional flexibility was introduced into the delivery scheme by allowing spots to be re-delivered if needed. However, once a beam has been fully delivered and the gantry angle changes,

no further irradiation from that direction is possible. By including already used spots as candidates in the optimization, the optimizer gains flexibility, although the conformity term prevents unnecessary changes.

After convergence, the optimized spot weights x_t are filtered to remove values below the minimum deliverable threshold, set to 0.01 MU in this work. The remaining weights define the updated reference plan P_t, and treatment delivery then resumes using this adapted plan.

C. Data and treatment planning

Treatment planning optimization was performed within the open-source research treatment planning system OpenTPS (Wuyckens et al., 2023). This study was conducted using synthetic phantom data and retrospectively collected clinical patient data, including two lung cancer cases and one cervix cancer case. Local ethics committee approval was obtained for the use of all patient data under the agreement "Learning from the past: the MIRO treatment planning data base" which received the ethical approval by the CEHF (Comité d'Ethique hospitalo-facultaire) (2019/16SEP/402).

1. Water phantoms

The proposed framework was initially tested on synthetic phantoms to evaluate a variety of motion scenarios. The phantom consisted of a cylindrical water volume with a density of 1 g/cm^3, containing a spherical target region of the same density positioned slightly off-center. The cylinder was surrounded by air with a density of 0.001 g/cm^3. The phantom geometry, illustrated in Figure 2, had voxels of size 1 x 1 x 1 mm^3.
Two beams with gantry angles at 0° and 45° were used. The prescribed dose was 20 Gy in 10 fractions. To obtain dose distribution conformal to the target region, maximum and minimum dose constraints were set for the target and additional rings with maximal dose objectives were also added around the target. The list of optimization objectives can be found in supplementary material 1. A 10 mm margin around the target was used for spot placement, with a spot and layer spacing of 4 mm.
For all cases, a single adaptation was simulated after ¼ of the estimated delivery time. The cases described below and summarized in Table 1, represent toy models of increasing complexity to test the feasibility and potential of adaptation of our proposed workflow.

Static case

Before evaluating the adaptation potential of the proposed framework, a static case is first tested as the algorithm should be able to maintain the treatment quality on a non-moving anatomy. This allows us to check if the optimizer can account for the background

dose. This may be challenging as it results in an updated target dose objective that is no longer homogeneous.

Moving heterogeneity

To evaluate the adaptation capabilities against moving heterogeneities, a high-density sphere was added in the field of a beam. During treatment delivery, the bone sphere with a 10 mm radius and a 1.5 g/cm^3 density was then shifted to be in the field of both beams.

Moving OAR

Then, we tested the ability to spare an OAR that is moving within the field of a beam. An OAR was defined as a spherical region with a 10 mm radius and its corresponding optimization objective was set to maximum 5 Gy.

Target shift

The framework was also tested on a moving target by simulating a target shift of 10 mm. With this scenario, we can evaluate if the spots weights from the initial plan are properly displaced to be provided to the optimizer as prior information.

Target expansion and shrinkage

Finally, a scenario with a 5 mm target expansion and one with a 5 mm target shrinkage are simulated. With these test cases, we evaluate the ability to adapt the number of spots to maintain a good target coverage by adding or removing spots and layers if needed.

	Scenario	Challenge
Static case	Default water phantom	• Maintain treatment quality on a non-moving anatomy • Consider the background dose which results in an updated target dose objective that is no longer homogeneous
Moving heterogeneity	Moving spherical bone with a 10 mm radius and a density of 1.5 g/cm^3	Adapt against moving heterogeneities in the field of the beams

Moving OAR	Moving spherical OAR with a 10 mm radius and an optimization objective set to maximum 5 Gy	Spare an OAR moving in the field of the beams
Target shift	10 mm target shift	Displace the spots weights from the initial plan to provide them to the optimizer as prior information
Target expansion	5 mm target expansion	Adapt the number of spots to maintain good target coverage by adding spots and layers if needed
Target shrinkage	5 mm target shrinkage	Adapt the number of spots to maintain good target coverage by removing spots and layers if needed

Table 1: *Summary of the different water phantom test cases and the specific challenges associated with each scenario.*

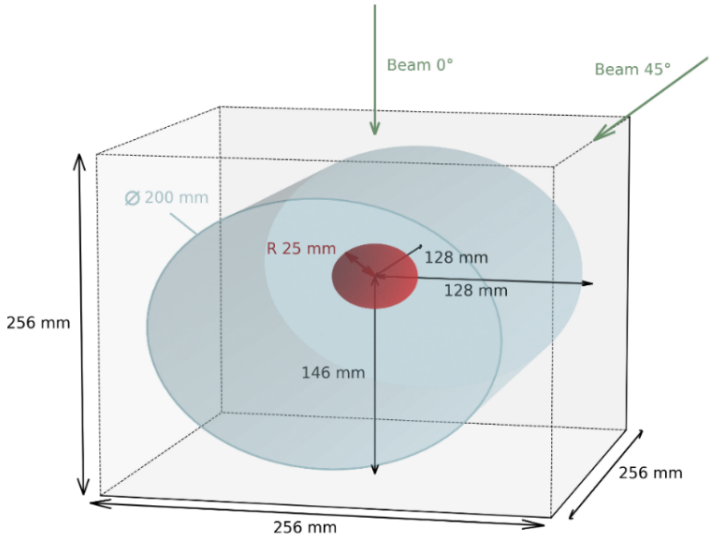

Figure 2 : *Schematic of the phantom used to test the adaptive framework, and showing the two beam directions (green). The geometry consists of a cubic phantom containing a water cylinder (blue) with an off-center spherical target (red) surrounded by air.*

2.Lung cases

To demonstrate the potential of this adaptive approach, we chose to test it on two lung cases for which we had 4D CT scans available with a resolution of 1.17 mm x 1.17 mm x 2 mm together with the target and OARs contours for each phase. Although motion

induced by breathing is not the primary target of the proposed workflow, these cases illustrate the ability of the method to handle complex intra-fraction anatomical changes. The CT of the first phase was used to optimize the initial plan P_0. The CTs of subsequent phases were then used as new images to re-optimize, resulting in 9 simulated re-optimizations in total. The plan adaptations were triggered at regular time intervals through the simulated treatment delivery, and the corresponding CT phase was used to represent the anatomy at that time.

For lung case 1, the tumor was located in the right middle lobe, and for lung case 2, it was located in the right lower lobe with nodal involvement. Additional information about the tumor size and motion amplitude is available in table 2.

The prescribed dose was 60 Gy in 30 fractions. The plans were generated with 3 beams with a fixed couch angle [C] at 0°. For the first patient, the gantry angles [G] were at 190°, 240°, and 290°, and for the second patient, they were at 180°, 270°, and 340°. For both lung cases, the re-optimizations corresponding to the CTs of phases 2 and 3 occurred during delivery of the first beam, phases 4-7 during delivery of the second beam, and phases 8-10 during delivery of the third beam.

A 10 mm margin was applied to the target for spot placement and the spot spacing was set to 5 mm. For lung case 1, the layer spacing was 5 mm and for lung case 2, it was set to 4 mm.

The clinical goals can be found in supplementary material 2.

Patient	**CTV tumor size [cm^3]**	**Tumor motion amplitude [mm]**			**Initial No. spots**
		Right-Left	Anterior-Posterior	Superior-Inferior	
Lung 1	152.6	4.2	2.1	3.1	11359
Lung 2	294.4	2.1	2.5	10.6	27585
Cervix	632.9	5.5	4.2	3.9	98698

Table 2 : *Patient characteristics including the tumor size, the tumor motion amplitude in absolute value and the number of spots in the initial plan P_0*

3. Cervix case

Eventually, the adaptive workflow was tested on a cervix case. Two CT scans with a resolution of 1.074 mm x 1.074 mm x 3 mm, together with corresponding contours for the target and relevant OARs, were obtained for two consecutive fractions that presented significant anatomical changes. These CT scans reflect inter-fractional changes but were assumed to be acquired during delivery to simulate intra-fractional variations. The tumor size and motion magnitude are displayed in table 2.

The initial plan P_0 was optimized on the CT of the first fraction. Then, half of the delivery was simulated on that CT before re-optimizing a new plan on the CT of the second

fraction. Therefore, only a single re-optimization halfway through the delivery was made for this case.

The prescription was 45 Gy in 25 fractions. We designed the plans with four beams (Schoot et al., 2016) including two posterior-oblique beams ([C] = 0°, [G] = 150°, 210°) and two lateral-opposed beams ([C] = 0°, [G] = 90°, 270°). The spot and layer spacing was set to 5 mm with a 10 mm margin around the target. Plans were optimized on the union of the low-risk target volume of the primary tumor (CTV-T LR) and elective nodal target volume (CTV-E). Planning objectives were defined according to the EMBRACE II protocol (image-guided intensity modulated External beam radiochemotherapy and MRIbased adaptive BRAchytherapy in locally advanced CErvical cancer) (Pötter et al., 2018).

D. Benchmarking and plan quality evaluation

Plans obtained with the beamlet-free adaptive workflow described above were evaluated against three benchmarks: the initial plan, the non-adapted plan, and a plan adapted using a state-of-the-art beamlet-based optimization method. Except for the initial plan, all other plans were evaluated on the most recent anatomy by accumulating the incremental dose contributions delivered over the course of treatment onto the latest image acquired during delivery. Dosimetric performance of beamlet-free plans was compared across all three benchmarks, while relative computational efficiency was assessed only between the two adaptive workflows.

1. Initial plans

The initial plan refers to the reference treatment plan P_0, generated prior to the start of treatment delivery and both optimized and evaluated on the planning CT (CT_0). This plan serves as the baseline for all subsequent comparisons.

2. Non-adapted plans

The non-adapted plan corresponds to the delivery of a treatment plan to the evolving anatomy without any re-optimization or adaptation during delivery.

For the water phantom cases, the non-adapted plan corresponds to the same initial plan P_0, delivered and evaluated on the changing phantom geometry.

For the patient cases, the comparison was adjusted to ensure fairness since the initial plans P_0 were not robustly optimized. For the two lung patients, the non-adapted plans were optimized on the ITV contours and on the average CT instead of the CTV contours and the CT of the first respiratory phase. For the cervix case, the PTV was used for the non-adapted plan instead of the union of CTV-T LR and CTV-E.

3. Beamlet-based re-optimized plans

To compare computation times, the full adaptive workflow was also tested with the

traditional beamlet-based optimization method. The beamlet matrix was computed with 50000 particles per beamlet from which the plans were optimized in OpenTPS (Wuyckens et al., 2023) with the SciPy implementation of the L-BFGS algorithm (Liu and Nocedal, 1989). For a fair comparison with the beamlet-free method, the deformed spot weights from previous optimizations were also provided to the optimizer for the initialization, and the same conformity term was added to the objective function. All beamlet-based plans were also non-robustly optimized to ensure a consistent comparison with the beamlet-free workflow. For each plan obtained with the beamlet-free algorithm, the final objective function value was computed and used as a stopping criterion for the beamlet-based optimizations to ensure that similar plan quality was reached for both methods.

4. Metrics

The metrics of interest are the clinical goals as well as the homogeneity index (HI) defined as $(D_2 - D_{98})/D_{50}$ and the Paddick conformity index (CI) (Paddick, 2000), defined as $\frac{V_{95,TV}^2}{V_{TV}V_{95}}$ with $V_{95,TV}$ the volume of the target covered by the 95% isodose, V_{95} the volume of the 95% isodose and V_{TV} the target volume.

The target coverage was considered acceptable if 98% of the target volume was receiving at least 95% of the prescribed dose ($D_{98\%} > 95\%\ D_p$).

For consistency and a fair comparison, all evaluation dose maps were computed with a higher number of particles to reduce statistical uncertainty, corresponding to a total of 50000 particles times the number of spots. For the non-adapted and adapted plans, the evaluation dose map was obtained using the same incremental dose accumulation procedure as for the background dose, while the initial plan was evaluated on the planning CT (CT_0) with a single dose computation. All dose computations, including those of the beamlet matrix for the beamlet-based approach, were performed with the MCsquare dose engine (Souris et al., 2016) which is also employed within the beamlet-free algorithm. The resolution of all dose grids matched that of the corresponding CT data.

To estimate the impact of statistical fluctuations, the full adaptive workflow was repeated five times with both the beamlet-free and the beamlet-based approaches. For the three patient cases, the computation times for each re-optimization were compared between the two approaches. The beamlet-based method involves two distinct components: the time required to compute the beamlets and the optimization time. In contrast, a single time measure is reported for the beamlet-free method, as dose computation and optimization are performed jointly within the Monte Carlo simulation.

All computations were performed on an Intel(R) Xeon(R) Gold 6248 CPU @ 2.50GHz with a total of 80 CPU cores and 512 GB RAM.

III. Results

1. Water phantoms

In the static scenario, the adapted plan achieved the same dose distribution as the initial reference plan, as demonstrated by the overlapping dose-volume histograms (DVHs) shown in Figure 3.

For all cases (Figure 4), the adapted plan maintained good target coverage, with a $D_{98\%}$ between 95.5% and 97.5% of the prescribed dose and it avoided hot spots ($V_{105\%} = 0$). In contrast, the non-adapted plans showed degraded performance.

In the scenario with the moving OAR, adaptation maintained OAR sparing by keeping D_{mean} below 5 Gy, while it increased to 9.9 Gy without adaptation.

In the other scenarios, hot spots were observed with non-adapted plans ($V_{105\%} > 0\%$) and target coverage also decreased, with $D_{98\%}$ below 95% of the prescribed dose.

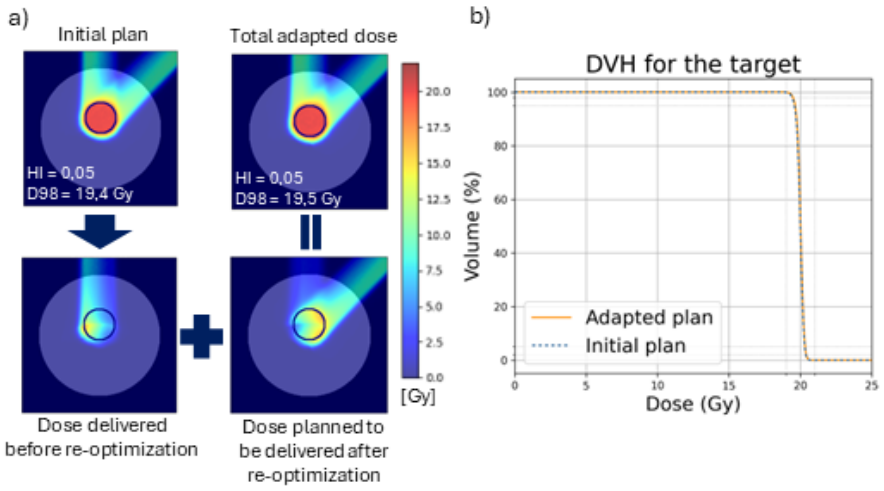

Figure 3 : Static water phantom case a) Dose distributions for the initial plan (upper left), the initial plan partially delivered until re-optimization (lower left), the re-optimized remaining plan (lower right) and the total adapted plan (upper right) which is the sum of the two lower dose distributions b) Dose volume histogram (DVH) comparing the initial total adapted plans

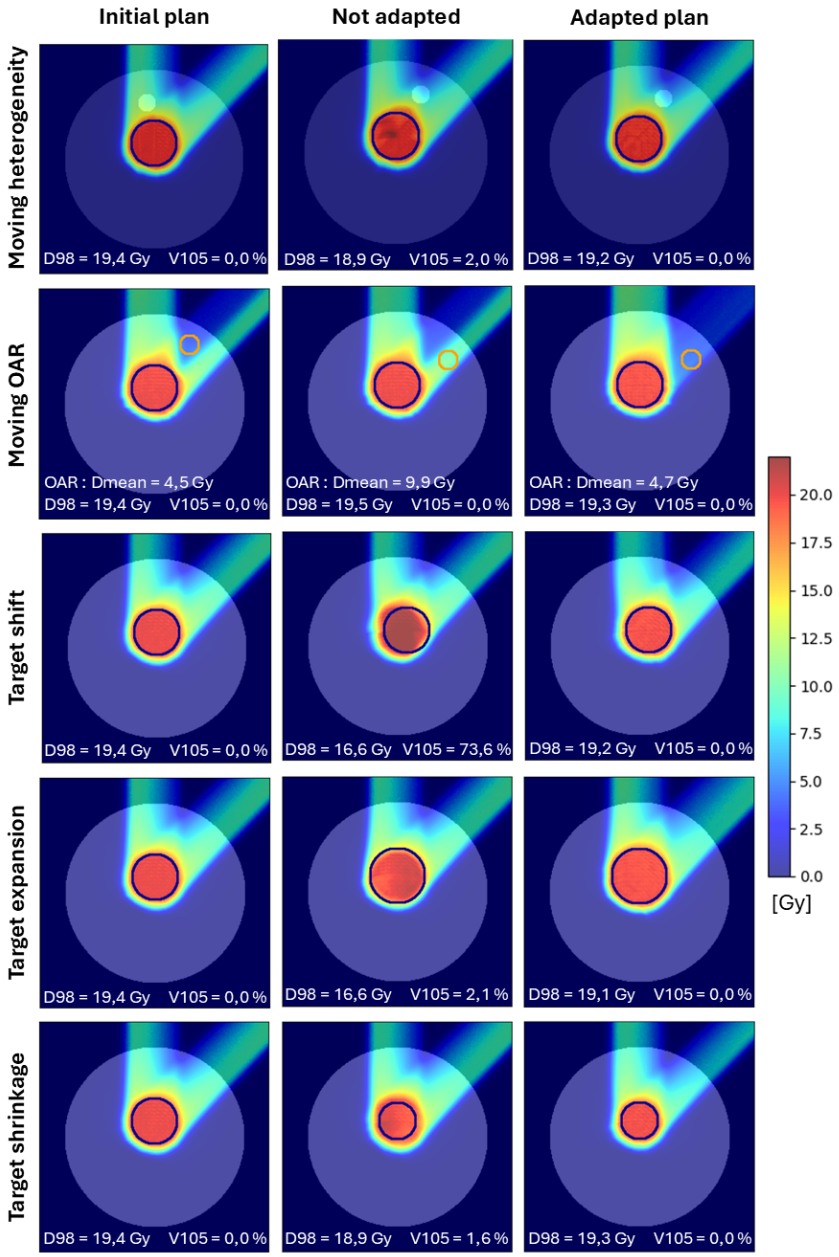

Figure 4 : *Dose distributions comparing the initial, non- adapted, and adapted plans over different scenarios with a motion or deformation simulated after ¼ of the treatment delivery. On each dose map, the target volume is delineated in dark blue and the target metrics $D_{98\%}$ [Gy] and $V_{105\%}$ [%] are displayed for a prescribed dose of 20 Gy.*
From left to right: *dose map of the initial plan optimized on the first CT, dose map obtained after delivery of the initial plan on the changing anatomy without adaptation, and dose map obtained after delivery on the moving anatomy of the adapted plan generated with the proposed beamlet-free approach.*
From top to bottom: *moving heterogeneity (spherical bone), moving spherical organ at risk (OAR) delineated in orange with a 5 Gy maximum dose constraint, 10 mm target shift, 5 mm target expansion, and 5 mm target shrinkage.*

2. Lung cases

Dosimetric evaluation

For both lung patients, the initial plan, the non-adapted plan and the adapted plans using the beamlet-free and beamlet-based approaches were compared in table 3. Figure 5 shows, for each OAR metric, the difference from the initial reference plan P_0.

In lung case 1, although target coverage is still maintained without adaptation ($D_{98\%} > 95\%$ D_p), inspection of the dose distribution revealed hot spots (Figure 6), and most other metrics were degraded compared with the initial and adapted plans. The only exception was the esophageal $D_{0.04cc}$, which improved slightly in comparison with the initial plan but was further reduced with adapted plans. Sparing of the heart and spinal canal also improved with adaptation, and other metrics were maintained within clinical goals for both adaptive methods.

In lung case 2, target coverage, homogeneity, and conformity objectives were not achieved without adaptation but were restored with both adaptative approaches.
Meeting the clinical requirement for the esophagus was particularly challenging for this patient, with values exceeding the clinical threshold by up to 0.7 Gy in both the initial plan and the beamlet-free adapted plan. However, all other clinical goals were satisfied with both adapted plans.

Overall, the beamlet-free and beamlet-based methods achieved comparable performance across most metrics, and neither method consistently outperformed the other. The largest differences for the two methods were mainly observed for the $D_{0.04cc}$ metrics, which are more sensitive to local fluctuations.
These metrics also showed the highest run-to-run variability for the beamlet-free method. In general, the standard deviation across runs was higher for beamlet-free than for beamlet-based optimization, reflecting the stochastic nature of the beamlet-free optimizer.

Computation time comparison

For both lung cases, Figure 7 compares the computation times between the two adaptive approaches for each of the nine re-optimizations. The beamlet-free method was 2.6 to 5.6 times faster than the beamlet-based method and, apart from the last re-optimization in lung case 1, it ran faster than both the beamlet computation and the beamlet-based optimization phases individually. The optimization time for the beamlet-free method was generally decreasing as treatment delivery progressed, except for phase 5, where a slight increase was observed. Similarly, the beamlet-based approach showed an overall trend of decreasing computation times over the course of treatment, but individual re-optimizations within the delivery of the same beam sometimes plateaued or even increased.

Because flexibility was introduced into the delivery scheme, all spots of the currently delivered beam were included in each re-optimization, even if some spots corresponded to energies already delivered. Consequently, re-optimizations performed during the delivery of the same beam involved a similar number of spots. Since beamlet computation time scales with the number of spots, it remained similar for re-optimizations carried out during the delivery of the same beam.

The relative time gain of the beamlet-free method increased with plan size (i.e., with the number of spots). Nevertheless, variations in the time ratio were still observed within the delivery of individual beams, with a tendency for the ratio to increase toward the end of a beam's delivery, sometimes preceded by a slight decrease in the middle.

		Initial plan	Non-adapted plan	Beamlet-free re-optimized plan	Beamlet-based re-optimized plan
Lung case 1	CTV D98 (Gy)	59.12	57.43	58.60 ± 0.04	58.04 ± 0.06
	HI	0.03	0.10	0.04 ± 0.00	0.04 ± 0.00
	CI	0.70	0.64	0.70 ± 0.01	0.74 ± 0.00
	Body D_{mean} (Gy)	2.37	2.55	2.44 ± 0.01	2.34 ± 0.00
	Lungs – GTV D_{mean} (Gy)	4.63	5.28	5.13 ± 0.30	4.90 ± 0.00
	Heart D_{mean} (Gy)	2.13	2.25	1.84 ± 0.01	1.79 ± 0.00
Lung case 2	CTV D98 (Gy)	58.85	52.32	58.41 ± 0.03	58.05 ± 0.02
	HI	0.03	0.21	0.04 ± 0.00	0.04 ± 0.00
	CI	0.65	0.40	0.69 ± 0.00	0.67 ± 0.00
	Body D_{mean} (Gy)	5.14	5.84	5.16 ± 0.01	5.27 ± 0.00
	Lungs – GTV D_{mean} (Gy)	8.65	9.92	9.21 ± 0.02	9.45 ± 0.01
	Heart D_{mean} (Gy)	3.43	3.54	4.19 ± 0.03	4.04 ± 0.01

Table 3 : *Dosimetric evaluation for the* ***two lung cases****. Abbreviations: CTV D98 = dose received at 98% of the CTV; HI = Homogeneity index; CI = Conformity index*

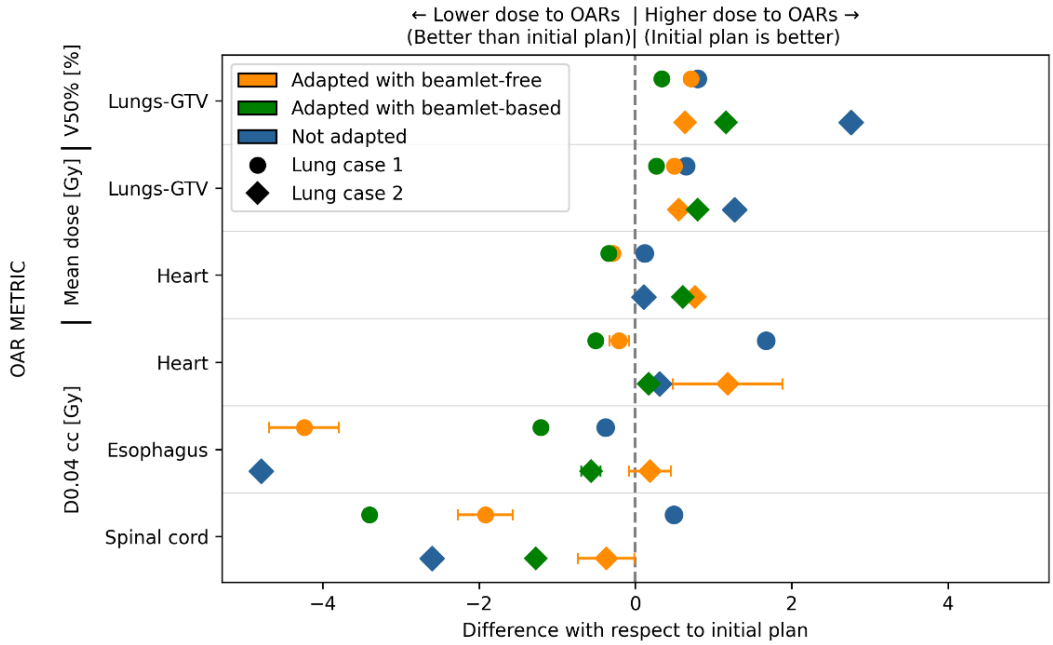

Figure 5 : *For the **two lung cases**, difference in terms of OAR metric between the initial plan and dose distributions after delivery of an adapted plan with the beamlet-free method (orange), an adapted plan with the beamlet-based method (green), and a non-adapted plan (blue). Positive values imply that the initial plan is better for that metric with less dose to the OAR.*

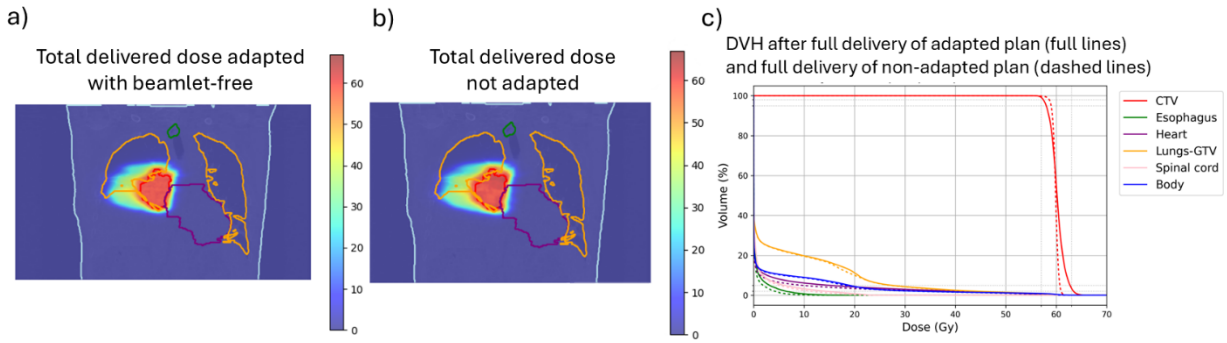

Figure 6 : *Comparison between a plan adapted with the beamlet-free framework and a non-adapted plan for lung case 1. a) Dose distribution after full delivery of the plan adapted 9 times with the beamlet-free optimizer; b) Dose distribution after full delivery of a plan without adaptation; c) Dose volume histogram (DVH) comparing the not adapted and the beamlet-free adapted plans*

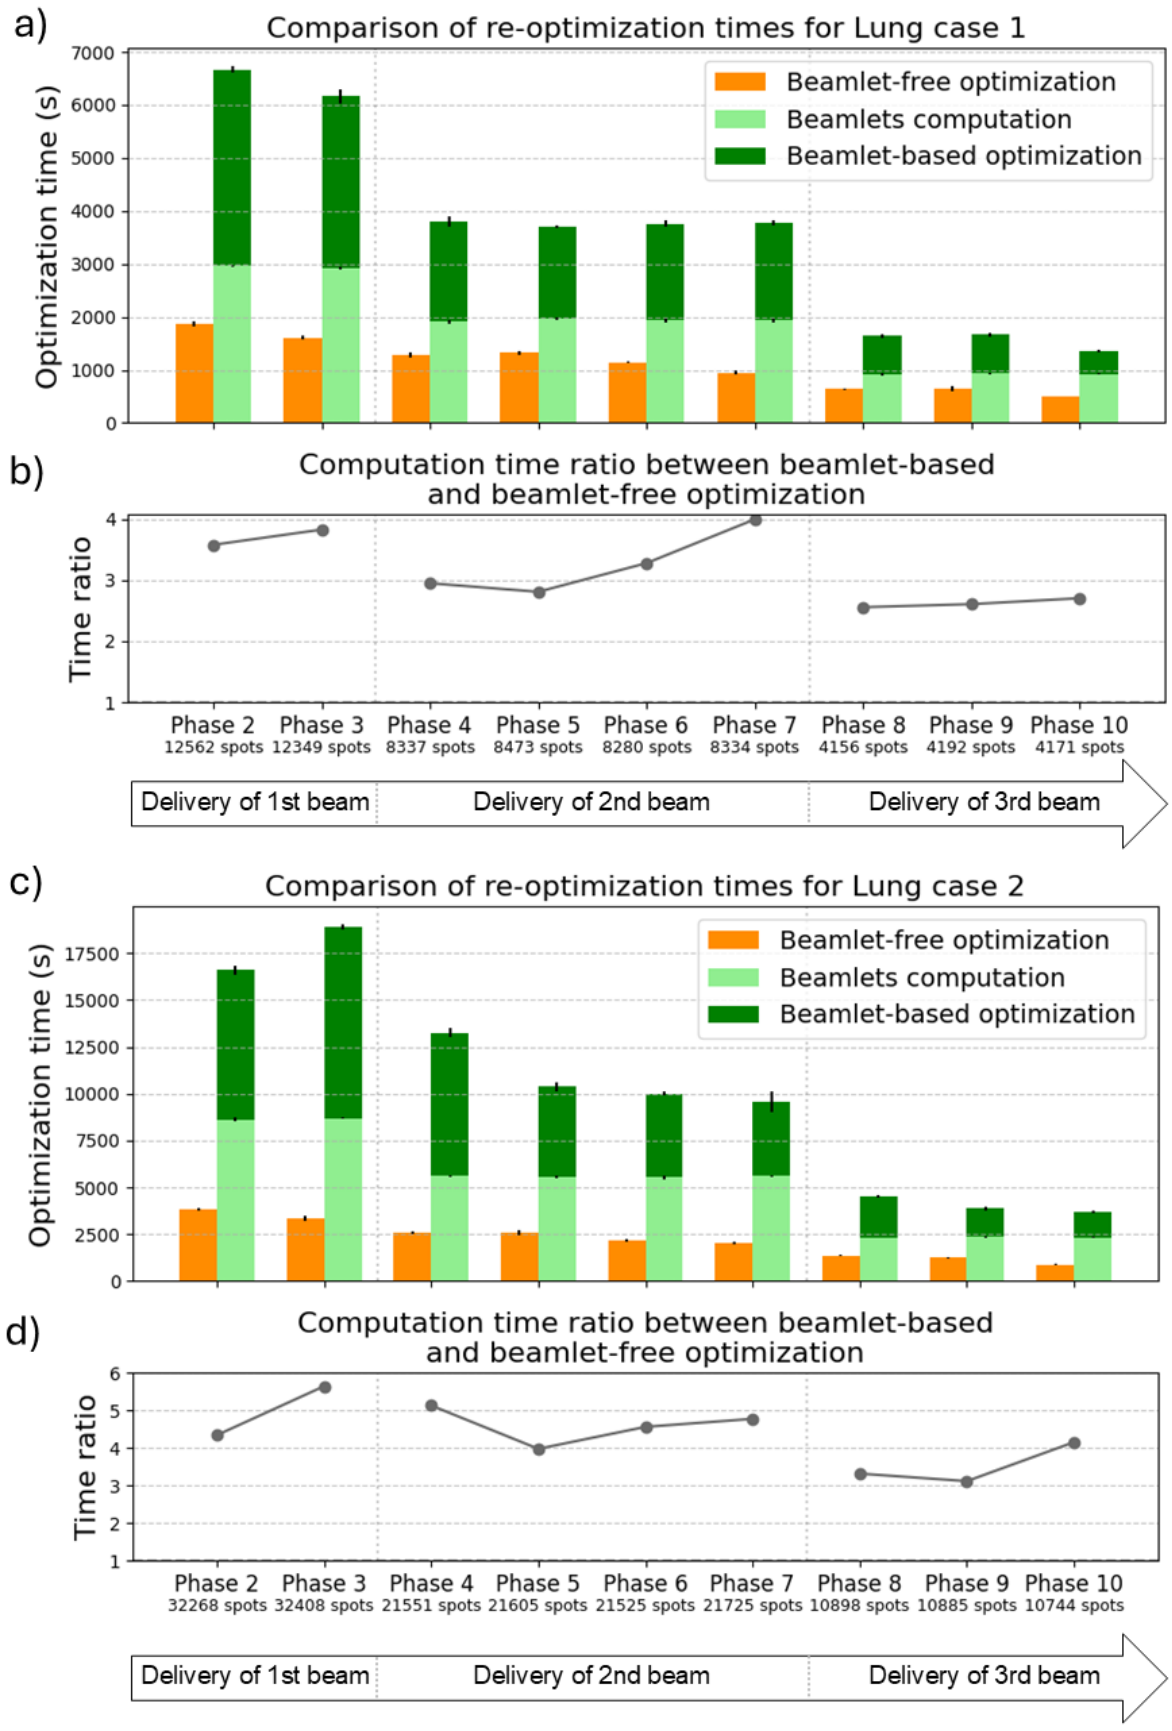
a)
Comparison of re-optimization times for Lung case 1
Optimization time (s)
Beamlet-free optimization
Beamlets computation
Beamlet-based optimization
b)
Computation time ratio between beamlet-based and beamlet-free optimization
Time ratio
Phase 2 12562 spots
Phase 3 12349 spots
Phase 4 8337 spots
Phase 5 8473 spots
Phase 6 8280 spots
Phase 7 8334 spots
Phase 8 4156 spots
Phase 9 4192 spots
Phase 10 4171 spots
Delivery of 1st beam
Delivery of 2nd beam
Delivery of 3rd beam
c)
Comparison of re-optimization times for Lung case 2
Optimization time (s)
Beamlet-free optimization
Beamlets computation
Beamlet-based optimization
d)
Computation time ratio between beamlet-based and beamlet-free optimization
Time ratio
Phase 2 32268 spots
Phase 3 32408 spots
Phase 4 21551 spots
Phase 5 21605 spots
Phase 6 21525 spots
Phase 7 21725 spots
Phase 8 10898 spots
Phase 9 10885 spots
Phase 10 10744 spots
Delivery of 1st beam
Delivery of 2nd beam
Delivery of 3rd beam

Figure 7 : *Computation times for the nine re-optimizations on CT phases 2-10, with the corresponding number of spots re-optimized.* ***(a, b)*** *Lung case 1;* ***(c, d)*** *Lung case 2;* ***(a, c)*** *Times for the beamlet-free optimization (orange), the beamlets computation (light green), and the beamlet-based optimization (dark green) with vertical black lines indicating the standard deviation across five runs;* ***(b, d)*** *Time ratio between the beamlet-free method and the beamlet-based method (beamlets computation + optimization).*

3. Cervix case

For the cervix case, target-related metrics are reported in Table 1, together with the mean body dose. Figure 8 shows the differences in OAR metrics relative to the initial plan, comparing the non-adapted plan, the beamlet-free adapted plan and the beamlet-based adapted plan.

Without adaptation, all evaluated metrics were degraded compared to the initial and adapted plans. In particular, target objectives were not met, showing low target coverage, homogeneity and conformity.

Plans adapted with both methods achieved similar target homogeneity. The beamlet-based approach resulted in slightly higher $D_{98\%,}$ while the beamlet-free approach provided improved target conformity. For the OARs, the beamlet-free method generally provided lower doses to the small bowel and bladder. Results for the anorectum and sigmoid were more variable, with no consistent advantage for either approach. Variability across repeated runs was lower for target metrics with the beamlet-free method, although the standard deviation remained higher for OAR metrics.

Computation times are displayed on Figure 9. This case involved the largest treated volume, leading to longer absolute computation times but also a greater relative speed-up: the beamlet-free approach was 7.7 times faster than the beamlet-based approach.

	Initial plan	Non-adapted plan	Beamlet-free re-optimized plan	Beamlet-based re-optimized plan
CTV D98 (Gy)	43.78	19.75	43.94 ± 0.01	44.27 ± 0.01
CTV V95% (%)	99.98	88.62	99.99 ± 0.00	99.99 ± 0.01
HI	0.03	0.60	0.03 ± 0.00	0.03 ± 0.00
CI	0.64	0.37	0.60 ± 0.00	0.57 ± 0.00
Body D_{mean} (Gy)	6.03	7.36	6.44 ± 0.02	6.40 ± 0.00

Table 1 : *Dosimetric evaluation for the* ***cervix case****.*

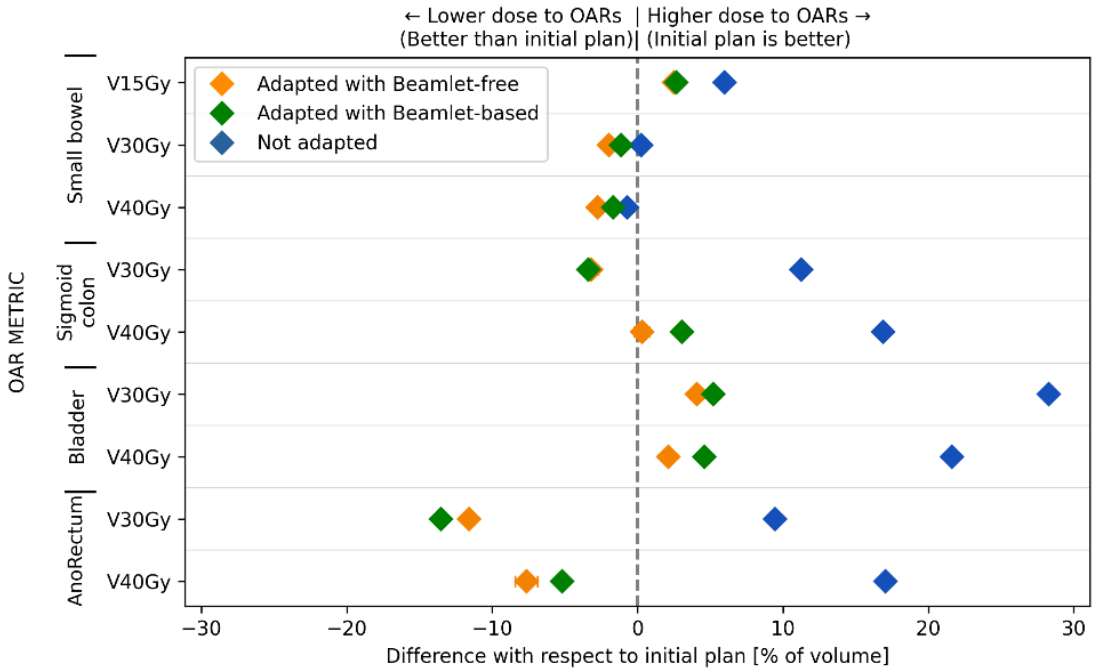

Figure 8 : *For the **cervix case**, difference in terms of OAR metric between the initial plan and dose distributions after delivery of an adapted plan with the beamlet-free method (orange), an adapted plan with the beamlet-based method (green), and a non-adapted plan (blue). Positive values imply that the initial plan is better for that metric with less dose to the OAR.*

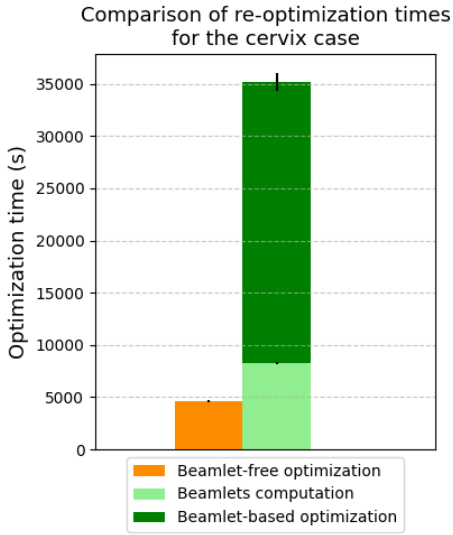

Figure 9 : *Comparison of computation times between the beamlet-free and the beamlet-based adaptive approaches for the cervix case.*

IV. Discussion

This study introduced an intra-fraction adaptive workflow relying on a beamlet-free optimizer. To validate this approach, we applied it to various motion and deformation scenarios in water phantoms and to three clinical patient cases over two treatment sites. The results were compared with the initial reference plans, the non-adapted plans, and plans adapted with state-of-the-art beamlet-based optimization.
For all presented test cases, the non-adapted plans failed to meet clinical requirements, highlighting the need for an intra-fraction adaptive strategy to mitigate motion and deformation. The results demonstrated that the proposed approach achieves equivalent plan quality to the standard beamlet-based approach, under a variety of motion and deformation scenarios, while significantly decreasing the computation times by a factor up to 7.7.

As highlighted by Pross *et al* (Pross et al., 2024), the computation time of the beamlet-free algorithm is less dependent on the total number of spots. This is due to its ability to adaptively direct particle simulations toward the most relevant spots based on prior iterations, thereby reducing the total number of simulated particles compared with conventional beamlet computation. This reduced dependence on spot count results in increasing time gains with plan size, as demonstrated in our results, with the cervix case showing the highest gains. For the lung cases, relative time gains were largest for the first beam angle and decreased for subsequent angles as fewer spots were involved per re-optimization. Within each beam angle, where the spot count per re-optimization remained similar, the beamlet-free method maintained its computational advantage and, in some cases, increased it during beam delivery, reflecting its ability to handle additional candidate spots at a lower computational cost. However, these conclusions were drawn from a limited number of patient cases. To confirm the generality of these observations, future investigations should apply this intra-fraction adaptive framework to a broader range of patient cases and treatment sites.
It should be noted that the intended use case of the proposed method is the compensation of sudden and unpredictable anatomical changes, which are more frequently encountered in treatment sites such as the lower abdomen or pelvis. The lung cases included in this work serve as illustrative examples demonstrating the potential of the approach under continuous motion scenarios.

It is noteworthy that computation time also depends on various factors such as the hardware, the specific beamlet-based optimizer, and the implementation. Over the past decade, substantial efforts have been made to accelerate both the dose calculation process and the optimization process in proton therapy, particularly in the context of online adaptation. GPU-accelerated analytical dose calculation methods have enabled the generation of typical IMPT plans in under ten seconds (Matter et al., 2019). To preserve the Monte Carlo accuracy, several fast Monte Carlo dose engines have been developed (Holmes et al., 2024), and GPU-

based Monte Carlo implementations have accelerated plan generation (Botas et al., 2018; Feng et al., 2022; Ma et al., 2014). GPUs have also been integrated directly in the optimization process to reduce computation times (Ma et al., 2014; Matter et al., 2019). In this study, both adaptive workflows were implemented within a Python research treatment planning system using a Monte Carlo dose engine written in C, which only enables CPU parallelization. Introducing GPU parallelization in both the optimization and the dose engine, and moving the entire framework to a faster and professional programming environment like C or similar, would help test how the proposed beamlet-free workflow can compete with current commercial beamlet-based solutions. Nevertheless, a major computational bottleneck in conventional methods remains the time required to compute the beamlets. By consistently outperforming this critical step in terms of computation time, the beamlet-free strategy demonstrates strong potential to accelerate adaptive workflow, while maintaining the accuracy of a Monte Carlo dose calculation engine.

Moreover, the proposed adaptive framework was able to explore solutions that involve re-delivering some spots, as those may still be useful to ensure optimal dose distributions under intra-fraction anatomical changes. In this work, this flexibility was introduced through a delivery scheme that allowed re-delivery of spots within the same beam. As discussed earlier, increasing the pool of spot candidates was handled more efficiently by the beamlet-free algorithm from a computational perspective. While such flexibility can improve dosimetric outcomes, it may also lead to longer treatment delivery times. Future work could investigate different delivery patterns with varying degrees of flexibility and quantify the trade-off between potential gains in plan quality and increased treatment duration.

Another point of discussion is the conformity term added to the objective function to penalize unnecessary deviations from the previously optimized plan. In this study, the weight of this term was empirically chosen and kept constant. A higher value enforces stronger compliance to the prior plan but limits responsiveness to anatomical variations. However, since this term is optimized jointly with the dosimetric objectives, the optimizer inherently balances plan conformity and dosimetric quality, relaxing the constraint when maintaining conformity would otherwise compromise dose objectives. In practice, the weight of this term could also be adjusted dynamically based on a deformation or motion metric. Alternatively, it could also serve as a user-controlled parameter, allowing clinicians to tune how much or how quickly a plan is allowed to evolve during treatment. While in real-time adaptation this weight cannot be tuned for an individual patient during delivery, it could be set a priori based on retrospective analyses of representative patient cohorts and on expected anatomical variability for the treatment site. Such flexibility for the user could be valuable in balancing plan stability with responsiveness in highly dynamic anatomical regions.

Although intra-fraction adaptation significantly reduces uncertainties related to motion and anatomical changes, residual sources of uncertainty remain such as proton range calculation errors. To address these uncertainties, robust optimization should be integrated into the adaptive workflow although it usually increases computation time. Previous work on the beamlet-free algorithm has explored robust optimization strategies in the context of conventional IMPT planning (Pross et al., 2025) by using a probabilistic approach and has demonstrated the same advantages in robust settings. Similar approaches could be applied to the intra-fraction adaptive setting.

To build a comprehensive real-time adaptive planning framework, three main technological challenges arise. First, fast imaging and segmentation are required to track the anatomy and update the contours of the target and OARs in near real-time. Second, accurate knowledge of the dose already delivered on the current anatomy is needed, either through dose reconstruction or accurate dose accumulation with deformable image registration (DIR). Third, rapid plan re-optimization must be achieved, ideally relying on accurate dose engines such as Monte Carlo simulation. The present work addressed the last challenge by enabling fast and accurate re-optimization that accounts for the accumulated dose previously delivered.

Regarding the challenge of imaging quality and availability, this study assumed access to fast and high-quality volumetric imaging during treatment. In practice, no such imaging modality is currently available for true real-time adaptive proton therapy. Potential solutions to acquire volumetric imaging at the treatment position include CBCT (Ghaznavi et al., 2025), MRI (Pham et al., 2022) and proton CT (Dedes et al., 2019). However, MRI integration in proton therapy and proton CT remain under development and is not yet clinically established. For photon therapy, acquisition of a CBCT using the HyperSight imaging solution (Varian Medical Systems) has been reported to take approximately 6 seconds (Robar et al., 2024). Although continuous or frequent imaging may improve adaptation and potentially reduce setup margins, this must be balanced against the potential increase in the patient's overall radiation exposure, depending on the imaging modality. A more practical strategy may therefore combine intermittent volumetric imaging with continuous anatomical or range monitoring. In this context, range verification techniques such as prompt gamma imaging could be used to trigger volumetric image acquisition when deviations are detected during treatment delivery (Berthold et al., 2023). A comprehensive review of potential real-time imaging modalities and their limitations is beyond the scope of this work and can be found elsewhere (Gambetta et al., 2025b; Lane et al., 2023). However, future studies should assess how image resolution, acquisition frequency and speed affect the accuracy of plan adaptation.
Additionally, it was also assumed that the delineation of anatomical structures was available for each image. In practice, the contour of the target and OARs might come from automated

segmentation methods, such as AI-based tools (Bibault and Giraud, 2024; Isaksson et al., 2023; Rong et al., 2024), or be propagated from previous images using the deformation vector field derived from image registration (Brock et al., 2017; Nenoff et al., 2023). Both approaches can introduce segmentation errors which may affect plan quality (Smolders et al., 2023).

For the second challenge concerning the computation of the background dose, it was estimated in this work by (1) computing the dose delivered on the previous anatomy, (2) mapping it to the current anatomy using deformable image registration (DIR), and (3) accumulating the mapped dose. Each of these steps introduces potential inaccuracies. For dose computation, the accuracy of the simulated dose depends on the number of particles used, and its impact on the final plan quality remains to be quantified. For registration, we assumed the availability of a reliable DIR algorithm, which itself is a source of uncertainty (Brock et al., 2017; Nenoff et al., 2023) and was not evaluated in this study. Ideally, intra-fraction adaptation would rely on true dose reconstruction (Onecha et al., 2025) on the actual anatomy during each delivery segment, rather than depending on dose computation on successive anatomy states and their deformation-based mapping.

Finally, it is important to consider the threshold at which residual uncertainties outweigh the benefits of intra-fraction adaptation. Future work should aim to identify conditions under which adaptation is beneficial, considering factors such as the magnitude of deformation, the accuracy of imaging and dose reconstruction and the time required for the adaptation. A decision framework could be developed to guide when adaptation should be triggered, given the magnitude of anatomical changes and the current technological limitations.

V. Conclusions

This work introduced an intra-fraction adaptation framework based on a beamlet-free optimization algorithm capable of near real-time updates during treatment delivery. The method integrates the accumulated background dose, anatomical changes and prior plan information to generate optimized spot weights with Monte Carlo accuracy, allowing flexible spot re-delivery while maintaining dosimetric accuracy. Future work should focus on computational acceleration, integration with robust optimization, and validation across a wider range of clinical scenarios. Together, these developments could pave the way toward fully adaptive proton therapy workflows capable of compensating intra-fraction motion and deformation in real time.

VI. Acknowledgements

Valentine Dormal is funded by F.R.S-FNRS (Fonds de la Recherche Scientifique) under a Télévie grant (7.4590.23). Ana M. Barragán Montero is a Research Associate and John A. Lee is a Research Director with F.R.S.-FNRS (Fonds de la Recherche Scientifique). We thank David

Dechambre and Dr. Dario Di Perri for providing suitable retrospective data and clinical expertise.

VII. Supplementary material:

1. Supplementary material 1: Optimization objectives for the water phantom cases

Without OAR:

ROI	Objective type	Dose limit [Gy]	Weight
Target volume	D_{max}	20	100
	D_{min}	20	100
Ring 1 (0-2 mm from target)	D_{max}	20	1
Ring 2 (2-4 mm)	D_{max}	18	2
Ring 3 (4-6 mm)	D_{max}	16	3
Ring 4 (6-8 mm)	D_{max}	14	4
Ring 5 (8-10 mm)	D_{max}	12	5

With OAR:

ROI	Objective type	Dose limit [Gy]	Weight
Target volume	D_{max}	20	150
	D_{min}	20	150
Ring 1 (0-2 mm from target)	D_{max}	21	1
Ring 2 (2-4 mm)	D_{max}	19	2
Ring 3 (4-6 mm)	D_{max}	17	3
Ring 4 (6-8 mm)	D_{max}	16	4
Ring 5 (8-10 mm)	D_{max}	15	5
OAR	D_{max}	5	10

2. Supplementary material 2 : Clinical goals for the lung cases

Region of interest (ROI)	Clinical goal
CTV	D95 % > 57 Gy (95% Dp = 57 Gy) V57 Gy > 98 %
Lungs - GTV	Dmean < 20 Gy V30 Gy < 20 %
Heart	Dmean < 20 Gy D0.04 cc < 63 Gy
Esophagus	D0.04 cc < 60 Gy
Spinal canal	D0.04 cc < 50 Gy

3. Supplementary material 3 : Timing evaluation of a complete intra-fraction adaptation framework

3.1 Time structure of an adaptive workflow using AI for contour generation

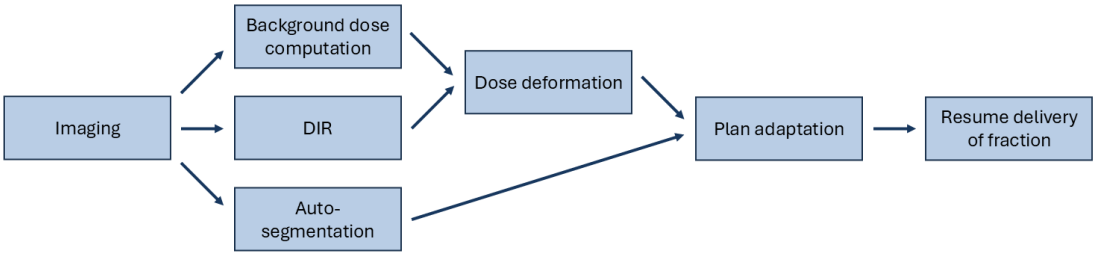

3.2 Time structure of an adaptive workflow using registration to propagate contours

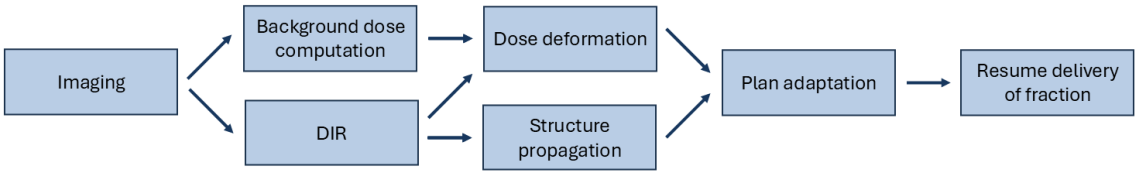

3.3 Table providing time estimations for different steps of the adaptive workflow

Adaptive step	Example of implementation	Timeframe	References
Imaging	CBCT with HyperSight - Photons (Commercial solution)	~ 6 s	(Robar et al., 2024)
Dose computation	GPU Virtual Particle Monte Carlo (Research implementation)	2.84 [0.8 to 6] s	(Shan et al., 2022)
	GPU Monte Carlo in RayStation (Commercial solution)	6.2 [2.2 to 62.2] s	(Fracchiolla et al., 2021)

DIR		ANACONDA in Raystation (Commercial solution)	[8 to 18] s	(Miura et al., 2017) (Weistrand and Svensson, 2015)
Contour Generation	Auto-segmentation (AI-based)	Limbus Contour (Commercial solution)	[24 to 50] s	(Wong et al., 2020) (Smolders et al., 2023)
		3D UNet (Research implementation)	[12.4 to 20.3] s	(Smolders et al., 2023)
	Contour propagation (registration-based)	Elastix with rigid registration (Open-source software)	[15.1 to 16.9] s	(Smolders et al., 2023)
		Elastix with deformable registration (Open-source software)	[3 to 30] s	(Qiao et al., 2019)

Fracchiolla, F., Engwall, E., Janson, M., Tamm, F., Lorentini, S., Fellin, F., Bertolini, M., Algranati, C., Righetto, R., Farace, P., Amichetti, M., Schwarz, M., 2021. Clinical validation of a GPU-based Monte Carlo dose engine of a commercial treatment planning system for pencil beam scanning proton therapy. Physica Medica 88, 226–234. https://doi.org/10.1016/j.ejmp.2021.07.012

Miura, H., Ozawa, S., Nakao, M., Furukawa, K., Doi, Y., Kawabata, H., Kenjou, M., Nagata, Y., 2017. Impact of deformable image registration accuracy on thoracic images with different regularization weight parameter settings. Physica Medica 42, 108–111. https://doi.org/10.1016/j.ejmp.2017.09.122

Qiao, Y., Jagt, T., Hoogeman, M., Lelieveldt, B.P.F., Staring, M., 2019. Evaluation of an Open Source Registration Package for Automatic Contour Propagation in Online Adaptive Intensity-Modulated Proton Therapy of Prostate Cancer. Front. Oncol. 9. https://doi.org/10.3389/fonc.2019.01297

Robar, J.L., Cherpak, A., MacDonald, R.L., Yashayaeva, A., McAloney, D., McMaster, N., Zhan, K., Cwajna, S., Patil, N., Dahn, H., 2024. Novel Technology Allowing Cone Beam Computed Tomography in 6 Seconds: A Patient Study of Comparative Image Quality. Practical Radiation Oncology 14, 277–286. https://doi.org/10.1016/j.prro.2023.10.014

Shan, J., Feng, H., Morales, D.H., Patel, S.H., Wong, W.W., Fatyga, M., Bues, M., Schild, S.E., Foote, R.L., Liu, W., 2022. Virtual Particle Monte Carlo (VPMC), a new concept to avoid simulating secondary particles in proton therapy dose calculation. Med Phys 49, 6666–6683. https://doi.org/10.1002/mp.15913

Smolders, A., Lomax, A., Weber, D., Albertini, F., 2023. Patient-specific neural networks for contour propagation in online adaptive radiotherapy. Phys. Med. Biol. 68, 095010. https://doi.org/10.1088/1361-6560/accaca

Weistrand, O., Svensson, S., 2015. The ANACONDA algorithm for deformable image registration in radiotherapy. Medical Physics 42, 40–53. https://doi.org/10.1118/1.4894702

Wong, J., Fong, A., McVicar, N., Smith, S., Giambattista, Joshua, Wells, D., Kolbeck, C., Giambattista, Jonathan, Gondara, L., Alexander, A., 2020. Comparing deep learning-based auto-segmentation of organs at risk and clinical target volumes to expert inter-observer variability in radiotherapy planning. Radiotherapy and Oncology 144, 152–158. https://doi.org/10.1016/j.radonc.2019.10.019

References:

Batista, V., Richter, D., Chaudhri, N., Naumann, P., Herfarth, K., Jäkel, O., 2018. Significance of intra-fractional motion for pancreatic patients treated with charged particles. Radiat Oncol 13, 120. https://doi.org/10.1186/s13014-018-1060-8

Bengtsson, I., Forsgren, A., Fredriksson, A., 2025. Dosimetric impact of real-time re-optimization of proton pencil-beam scanning for moving targets. https://doi.org/10.48550/arXiv.2501.16840

Berger, T., Godart, J., Jagt, T., Vittrup, A.S., Fokdal, L.U., Lindegaard, J.C., Kibsgaard Jensen, N.B., Zolnay, A., Reijtenbagh, D., Trnkova, P., Tanderup, K., Hoogeman, M., 2021. Dosimetric Impact of Intrafraction Motion in Online-Adaptive Intensity Modulated Proton Therapy for Cervical Cancer. International Journal of Radiation Oncology*Biology*Physics 109, 1580–1587. https://doi.org/10.1016/j.ijrobp.2020.11.037

Berthold, J., Pietsch, J., Piplack, N., Khamfongkhruea, C., Thiele, J., Hölscher, T., Janssens, G., Smeets, J., Traneus, E., Löck, S., Stützer, K., Richter, C., 2023. Detectability of Anatomical Changes With Prompt-Gamma Imaging: First Systematic Evaluation of Clinical Application During Prostate-Cancer Proton Therapy. International Journal of Radiation Oncology*Biology*Physics 117, 718–729. https://doi.org/10.1016/j.ijrobp.2023.05.002

Bibault, J.-E., Giraud, P., 2024. Deep learning for automated segmentation in radiotherapy: a narrative review. Br J Radiol 97, 13–20. https://doi.org/10.1093/bjr/tqad018

Booth, J., Caillet, V., Briggs, A., Hardcastle, N., Angelis, G., Jayamanne, D., Shepherd, M., Podreka, A., Szymura, K., Nguyen, D.T., Poulsen, P., O'Brien, R., Harris, B., Haddad, C., Eade, T., Keall, P., 2021. MLC tracking for lung SABR is feasible, efficient and delivers high-precision target dose and lower normal tissue dose. Radiother Oncol 155, 131–137. https://doi.org/10.1016/j.radonc.2020.10.036

Botas, P., Kim, J., Winey, B., Paganetti, H., 2018. Online adaption approaches for intensity modulated proton therapy for head and neck patients based on cone beam CTs and Monte Carlo simulations. Phys. Med. Biol. 64, 015004. https://doi.org/10.1088/1361-6560/aaf30b

Brock, K.K., Mutic, S., McNutt, T.R., Li, H., Kessler, M.L., 2017. Use of image registration and fusion algorithms and techniques in radiotherapy: Report of the AAPM Radiation

Therapy Committee Task Group No. 132. Med Phys 44, e43–e76. https://doi.org/10.1002/mp.12256

Dedes, G., Dickmann, J., Niepel, K., Wesp, P., Johnson, R.P., Pankuch, M., Bashkirov, V., Rit, S., Volz, L., Schulte, R.W., Landry, G., Parodi, K., 2019. Experimental comparison of proton CT and dual energy x-ray CT for relative stopping power estimation in proton therapy. Phys. Med. Biol. 64, 165002. https://doi.org/10.1088/1361-6560/ab2b72

Feng, H., Patel, S.H., Wong, W.W., Younkin, J.E., Penoncello, G.P., Morales, D.H., Stoker, J.B., Robertson, D.G., Fatyga, M., Bues, M., Schild, S.E., Foote, R.L., Liu, W., 2022. GPU-accelerated Monte Carlo-based online adaptive proton therapy: A feasibility study. Medical Physics 49, 3550–3563. https://doi.org/10.1002/mp.15678

Gambetta, V., Pieta, V., Berthold, J., Hölscher, T., Fredriksson, A., Richter, C., Stützer, K., 2025a. Partial adaptation for online-adaptive proton therapy triggered by during-delivery treatment verification: Feasibility study on prostate cancer treatments. Phys Imaging Radiat Oncol 34, 100755. https://doi.org/10.1016/j.phro.2025.100755

Gambetta, V., Stützer, K., Richter, C., 2025b. Current status and upcoming developments for online adaptive proton therapy enabling a closed feedback loop for near real-time adaptation. Front. Oncol. 15. https://doi.org/10.3389/fonc.2025.1660605

Gerlach, S., Hofmann, T., Fürweger, C., Schlaefer, A., 2023. Towards fast adaptive replanning by constrained reoptimization for intra-fractional non-periodic motion during robotic SBRT. Med Phys 50, 4613–4622. https://doi.org/10.1002/mp.16381

Ghaznavi, H., Maraghechi, B., Zhang, H., Zhu, T., Laugeman, E., Zhang, T., Zhao, T., Mazur, T.R., Darafsheh, A., 2025. Quantitative use of cone-beam computed tomography in proton therapy: challenges and opportunities. Phys. Med. Biol. 70, 09TR01. https://doi.org/10.1088/1361-6560/adc86c

Heijkoop, S.T., Langerak, T.R., Quint, S., Mens, J.W.M., Zolnay, A.G., Heijmen, B.J.M., Hoogeman, M.S., 2015. Quantification of intra-fraction changes during radiotherapy of cervical cancer assessed with pre- and post-fraction Cone Beam CT scans. Radiotherapy and Oncology 117, 536–541. https://doi.org/10.1016/j.radonc.2015.08.034

Holmes, J., Feng, H., Zhang, L., Fix, M.K., Jiang, S.B., Liu, W., 2024. Fast Monte Carlo dose calculation in proton therapy. Phys. Med. Biol. 69, 17TR01. https://doi.org/10.1088/1361-6560/ad67a7

Isaksson, L.J., Summers, P., Mastroleo, F., Marvaso, G., Corrao, G., Vincini, M.G., Zaffaroni, M., Ceci, F., Petralia, G., Orecchia, R., Jereczek-Fossa, B.A., 2023. Automatic Segmentation with Deep Learning in Radiotherapy. Cancers (Basel) 15, 4389. https://doi.org/10.3390/cancers15174389

Jagt, T., Breedveld, S., van Haveren, R., Heijmen, B., Hoogeman, M., 2018. An automated planning strategy for near real-time adaptive proton therapy in prostate cancer. Phys. Med. Biol. 63, 135017. https://doi.org/10.1088/1361-6560/aacaa7

Keall, P., Nguyen, D.T., O'Brien, R., Hewson, E., Ball, H., Poulsen, P., Booth, J., Greer, P., Hunter, P., Wilton, L., Bromley, R., Kipritidis, J., Eade, T., Kneebone, A., Hruby, G., Moodie, T., Hayden, A., Turner, S., Arumugam, S., Sidhom, M., Hardcastle, N., Siva, S., Tai, K.-H., Gebski, V., Martin, J., 2020. Real-Time Image Guided Ablative Prostate Cancer Radiation Therapy: Results From the TROG 15.01 SPARK Trial. Int J Radiat Oncol Biol Phys 107, 530–538. https://doi.org/10.1016/j.ijrobp.2020.03.014

Keall, P., Vedam, S., George, R., Bartee, C., Siebers, J., Lerma, F., Weiss, E., Chung, T., 2006. The clinical implementation of respiratory-gated intensity-modulated radiotherapy. Med Dosim. 31, 152–162. https://doi.org/10.1016/j.meddos.2005.12.002

Keall, P.J., El Naqa, I., Fast, M.F., Hewson, E.A., Hindley, N., Poulsen, P., Sengupta, C., Tyagi, N., Waddington, D.E.J., 2025. Real-Time Dose-Guided Radiation Therapy. Int J Radiat Oncol Biol Phys 122, 787–801. https://doi.org/10.1016/j.ijrobp.2025.04.019

Kilby, W., Dooley, J.R., Kuduvalli, G., Sayeh, S., Maurer, C.R., 2010. The CyberKnife Robotic Radiosurgery System in 2010. Technol Cancer Res Treat 9, 433–452. https://doi.org/10.1177/153303461000900502

Kontaxis, C., Bol, G.H., Stemkens, B., Glitzner, M., Prins, F.M., Kerkmeijer, L.G.W., Lagendijk, J.J.W., Raaymakers, B.W., 2017. Towards fast online intrafraction replanning for free-breathing stereotactic body radiation therapy with the MR-linac. Phys. Med. Biol. 62, 7233. https://doi.org/10.1088/1361-6560/aa82ae

Kubo, H.D., Hill, B.C., 1996. Respiration gated radiotherapy treatment: a technical study. Phys. Med. Biol. 41, 83. https://doi.org/10.1088/0031-9155/41/1/007

Kumagai, M., Hara, R., Mori, S., Yanagi, T., Asakura, H., Kishimoto, R., Kato, H., Yamada, S., Kandatsu, S., Kamada, T., 2009. Impact of Intrafractional Bowel Gas Movement on Carbon Ion Beam Dose Distribution in Pancreatic Radiotherapy. International Journal of Radiation Oncology*Biology*Physics 73, 1276–1281. https://doi.org/10.1016/j.ijrobp.2008.10.055

Lagerwaard, F., Bohoudi, O., Tetar, S., Admiraal, M.A., Rosario, T.S., Bruynzeel, A., 2018. Combined Inter- and Intrafractional Plan Adaptation Using Fraction Partitioning in Magnetic Resonance-guided Radiotherapy Delivery. Cureus 10, e2434. https://doi.org/10.7759/cureus.2434

Lane, S.A., Slater, J.M., Yang, G.Y., 2023. Image-Guided Proton Therapy: A Comprehensive Review. Cancers (Basel) 15, 2555. https://doi.org/10.3390/cancers15092555

Liu, D.C., Nocedal, J., 1989. On the limited memory BFGS method for large scale optimization. Mathematical Programming 45, 503–528. https://doi.org/10.1007/BF01589116

Lu, W., Chen, M., Ruchala, K.J., Chen, Q., Langen, K.M., Kupelian, P.A., Olivera, G.H., 2009. Real-time motion-adaptive-optimization (MAO) in TomoTherapy. Phys. Med. Biol. 54, 4373. https://doi.org/10.1088/0031-9155/54/14/003

Ma, J., Beltran, C., Seum Wan Chan Tseung, H., Herman, M.G., 2014. A GPU-accelerated and Monte Carlo-based intensity modulated proton therapy optimization system. Medical Physics 41, 121707. https://doi.org/10.1118/1.4901522

Matter, M., Nenoff, L., Meier, G., Weber, D.C., Lomax, A.J., Albertini, F., 2019. Intensity modulated proton therapy plan generation in under ten seconds. Acta Oncologica 58, 1435–1439. https://doi.org/10.1080/0284186X.2019.1630753

Mostafaei, F., Tai, A., Omari, E., Song, Y., Christian, J., Paulson, E., Hall, W., Erickson, B., Li, X.A., 2018. Variations of MRI-assessed peristaltic motions during radiation therapy. PLOS ONE 13, e0205917. https://doi.org/10.1371/journal.pone.0205917

Nejad-Davarani, S.P., Sevak, P., Moncion, M., Garbarino, K., Weiss, S., Kim, J., Schultz, L., Elshaikh, M.A., Renisch, S., Glide-Hurst, C., 2019. Geometric and dosimetric impact of anatomical changes for MR-only radiation therapy for the prostate. Journal of Applied Clinical Medical Physics 20, 10–17. https://doi.org/10.1002/acm2.12551

Nenoff, L., Amstutz, F., Murr, M., Archibald-Heeren, B., Fusella, M., Hussein, M., Lechner, W., Zhang, Y., Sharp, G., Vasquez Osorio, E., 2023. Review and recommendations on

deformable image registration uncertainties for radiotherapy applications. Phys Med Biol 68, 24TR01. https://doi.org/10.1088/1361-6560/ad0d8a
Nesterov, Yu., 2012. Efficiency of Coordinate Descent Methods on Huge-Scale Optimization Problems. SIAM J. Optim. 22, 341–362. https://doi.org/10.1137/100802001
Onecha, V.V., Espinosa-Rodriguez, A., Soneira-Landín, C., Arias-Valcayo, F., Gaitán-Dominguez, S., Martinez-Nouvilas, V., García-Díez, M., Ibáñez, P., España, S., Sanchez-Parcerisa, D., Cerrón-Campoo, F., Vera-Sánchez, J.A., Mazal, A., Udias, J.M., Fraile, L.M., 2025. Real-time dose reconstruction in proton therapy from in-beam PET measurements. Phys. Med. Biol. 70, 075008. https://doi.org/10.1088/1361-6560/adbfd9
Paddick, I., 2000. A simple scoring ratio to index the conformity of radiosurgical treatment plans. Technical note. J Neurosurg 93 Suppl 3, 219–222. https://doi.org/10.3171/jns.2000.93.supplement
Pham, T.T., Whelan, B., Oborn, B.M., Delaney, G.P., Vinod, S., Brighi, C., Barton, M., Keall, P., 2022. Magnetic resonance imaging (MRI) guided proton therapy: A review of the clinical challenges, potential benefits and pathway to implementation. Radiotherapy and Oncology 170, 37–47. https://doi.org/10.1016/j.radonc.2022.02.031
Pötter, R., Tanderup, K., Kirisits, C., de Leeuw, A., Kirchheiner, K., Nout, R., Tan, L.T., Haie-Meder, C., Mahantshetty, U., Segedin, B., Hoskin, P., Bruheim, K., Rai, B., Huang, F., Van Limbergen, E., Schmid, M., Nesvacil, N., Sturdza, A., Fokdal, L., Jensen, N.B.K., Georg, D., Assenholt, M., Seppenwoolde, Y., Nomden, C., Fortin, I., Chopra, S., van der Heide, U., Rumpold, T., Lindegaard, J.C., Jürgenliemk-Schulz, I., 2018. The EMBRACE II study: The outcome and prospect of two decades of evolution within the GEC-ESTRO GYN working group and the EMBRACE studies. Clin Transl Radiat Oncol 9, 48–60. https://doi.org/10.1016/j.ctro.2018.01.001
Pross, D., Dormal, V., Souris, K., Montero, A.M.B., Sterpin, E., Lee, J.A., 2025. Scalable, memory-efficient robust proton therapy optimization through beamlet-free treatment planning. Med Phys 52, e70177. https://doi.org/10.1002/mp.70177
Pross, D., Wuyckens, S., Deffet, S., Sterpin, E., Lee, J.A., Souris, K., 2024. Technical note: Beamlet-free optimization for Monte-Carlo-based treatment planning in proton therapy. Med Phys 51, 485–493. https://doi.org/10.1002/mp.16813
Robar, J.L., Cherpak, A., MacDonald, R.L., Yashayaeva, A., McAloney, D., McMaster, N., Zhan, K., Cwajna, S., Patil, N., Dahn, H., 2024. Novel Technology Allowing Cone Beam Computed Tomography in 6 Seconds: A Patient Study of Comparative Image Quality. Practical Radiation Oncology 14, 277–286. https://doi.org/10.1016/j.prro.2023.10.014
Rong, Y., Chen, Q., Fu, Y., Yang, X., Al-Hallaq, H.A., Wu, Q.J., Yuan, L., Xiao, Y., Cai, B., Latifi, K., Benedict, S.H., Buchsbaum, J.C., Qi, X.S., 2024. NRG Oncology Assessment of Artificial Intelligence Deep Learning-Based Auto-segmentation for Radiation Therapy: Current Developments, Clinical Considerations, and Future Directions. Int J Radiat Oncol Biol Phys 119, 261–280. https://doi.org/10.1016/j.ijrobp.2023.10.033
Rusu, D.N., Cunningham, J.M., Arch, J.V., Chetty, I.J., Parikh, P.J., Dolan, J.L., 2023. Impact of intrafraction motion in pancreatic cancer treatments with MR-guided adaptive radiation therapy. Front Oncol 13, 1298099. https://doi.org/10.3389/fonc.2023.1298099
Schoot, A.J.A.J. van de, Visser, J., Kesteren, Z. van, Janssen, T.M., Rasch, C.R.N., Bel, A., 2016. Beam configuration selection for robust intensity-modulated proton therapy in cervical cancer using Pareto front comparison. Phys. Med. Biol. 61, 1780. https://doi.org/10.1088/0031-9155/61/4/1780

Smolders, A., Choulilitsa, E., Czerska, K., Bizzocchi, N., Krcek, R., Lomax, A., Weber, D.C., Albertini, F., 2023. Dosimetric comparison of autocontouring techniques for online adaptive proton therapy. Phys. Med. Biol. 68, 175006. https://doi.org/10.1088/1361-6560/ace307

Snyder, J., Smith, B., Aubin, J.S., Shepard, A., Hyer, D., 2024. Simulating an intra-fraction adaptive workflow to enable PTV margin reduction in MRIgART volumetric modulated arc therapy for prostate SBRT. Front. Oncol. 13. https://doi.org/10.3389/fonc.2023.1325105

Souris, K., Lee, J.A., Sterpin, E., 2016. Fast multipurpose Monte Carlo simulation for proton therapy using multi- and many-core CPU architectures. Med Phys 43, 1700. https://doi.org/10.1118/1.4943377

Wiersma, R.D., Liu, X., 2019. A conceptual study on real-time adaptive radiation therapy optimization through ultra-fast beamlet control. Biomed Phys Eng Express 5, 055016. https://doi.org/10.1088/2057-1976/ab3ba9

Willigenburg, T., Zachiu, C., Bol, G.H., de Groot-van Beugel, E.N., Lagendijk, J.J.W., van der Voort van Zyp, J.R.N., Raaymakers, B.W., de Boer, J.C.J., 2022. Clinical application of a sub-fractionation workflow for intrafraction re-planning during prostate radiotherapy treatment on a 1.5 Tesla MR-Linac: A practical method to mitigate intrafraction motion. Radiother Oncol. 176, 25–30. https://doi.org/10.1016/j.radonc.2022.09.004

Wuyckens, S., Dasnoy, D., Janssens, G., Hamaide, V., Huet, M., Loÿen, E., de Hertaing, G.R., Macq, B., Sterpin, E., Lee, J.A., Souris, K., Deffet, S., 2023. OpenTPS -- Open-source treatment planning system for research in proton therapy. https://doi.org/10.48550/arXiv.2303.00365

Zhang, Y., Trnkova, P., Toshito, T., Heijmen, B., Richter, C., Aznar, M., Albertini, F., Bolsi, A., Daartz, J., Bertholet, J., Knopf, A., 2023. A survey of practice patterns for real-time intrafractional motion-management in particle therapy. Phys Imaging Radiat Oncol 26, 100439. https://doi.org/10.1016/j.phro.2023.100439