

ACCEPTED MANUSCRIPT

Dose-based optimisation for multi-leaf collimator tracking during radiation therapy

To cite this article before publication: Lars Mejnertsen *et al* 2021 *Phys. Med. Biol.* in press <https://doi.org/10.1088/1361-6560/abe836>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2021 Institute of Physics and Engineering in Medicine.

During the embargo period (the 12 month period from the publication of the Version of Record of this article), the Accepted Manuscript is fully protected by copyright and cannot be reused or reposted elsewhere.

As the Version of Record of this article is going to be / has been published on a subscription basis, this Accepted Manuscript is available for reuse under a CC BY-NC-ND 3.0 licence after the 12 month embargo period.

After the embargo period, everyone is permitted to use copy and redistribute this article for non-commercial purposes only, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by-nc-nd/3.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions will likely be required. All third party content is fully copyright protected, unless specifically stated otherwise in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Dose-Based Optimisation for Multi-Leaf Collimator Tracking during Radiation Therapy

Lars Mejnertsen^{*,1}, Emily Hewson¹, Doan Trang Nguyen^{1,2},
Jeremy Booth^{3,4}, Paul Keall¹

¹ACRF Image X Institute, Faculty of Medicine and Health, University of Sydney, NSW, Australia

²School of Biomedical Engineering, University of Technology Sydney, NSW, Australia

³Northern Sydney Cancer Centre, Royal North Shore Hospital, Sydney, NSW, Australia

⁴School of Physics, University of Sydney, Sydney, NSW, Australia

E-mail: *lars.mejnertsen@sydney.edu.au

October 2020

Abstract. Motion in the patient anatomy causes a reduction in dose delivered to the target, while increasing dose to healthy tissue. Multi-Leaf Collimator (MLC) tracking has been clinically implemented to adapt dose delivery to account for intrafraction motion. Current methods shift the planned MLC aperture in the direction of motion, then optimise the new aperture based on the difference in fluence. The drawback of these methods is that 3D dose, a function of patient anatomy and MLC aperture sequence, is not properly accounted for. To overcome the drawback of current fluence-based methods, we have developed and investigated real-time adaptive MLC tracking based on dose optimisation. A novel MLC tracking algorithm, dose optimisation, has been developed which accounts for the moving patient anatomy by optimising the MLC based on the dose delivered during treatment, simulated using a simplified dose calculation algorithm. The MLC tracking with dose optimisation method was applied *in silico* to a prostate cancer VMAT treatment dataset with observed intrafraction motion. Its performance was compared to MLC tracking with fluence optimisation and, as a baseline, without MLC tracking. To quantitatively assess performance, we computed the dose error and 3D γ failure rate (2 mm/2%) for each fraction and method. Dose optimisation achieved a γ failure rate of $(4.7 \pm 1.2)\%$ (mean and standard deviation) over all fractions, which was significantly lower than fluence optimisation $(7.5 \pm 2.9)\%$ (Wilcoxon sign-rank test $p < 0.01$). Without MLC tracking, a γ failure rate of $(15.3 \pm 12.9)\%$ was achieved. By considering the accumulation of dose in the moving anatomy during treatment, dose optimisation is able to optimise the aperture to actively target regions of underdose while avoiding overdose.

Dose Optimisation for MLC Tracking

1. Introduction

Intrafraction motion in the patient anatomy causes a reduced efficacy of radiotherapy, reducing dose delivered to the target while increasing dose to healthy tissue. The goal of adaptive radiotherapy is to adapt the treatment to the anatomy at the time of treatment. Accounting for intrafraction motion can be accomplished in a number of ways, including gating when the target position exceeds pre-defined margins (Keall et al. 2015), or actively adapting the treatment through either couch tracking (D'Souza et al. 2005), Multi-Leaf Collimator (MLC) tracking or a dedicated gimbaled/robotic tracking hardware (Kilby et al. 2010, Kamino et al. 2006). In MLC tracking, the MLC actively adapts to account for intrafraction motion (Sawant et al. 2008) by re-optimising the aperture in accordance with observed target motion in real-time. It has been implemented on a variety of commercial linear accelerators (*e.g.* Tacke et al. (2010), Fast et al. (2014), Hansen et al. (2016)), where it has been demonstrated to potentially reduce planning target volumes (PTV) and dose to organs at risk (OAR) (Menten et al. 2016, Caillet et al. 2017). Furthermore, MLC tracking has been used in a number of clinical trials for prostate (Keall et al. 2014, Keall et al. 2017) and lung (Booth et al. 2016) cancers, where it has been shown to improve dose conformity (Colvill et al. 2015, Keall et al. 2018).

Current MLC tracking methods use an intrafraction motion trace, typically obtained in real-time through tracking implanted markers, to shift the pre-planned MLC aperture in the direction of the observed motion. However, due to geometric and physical constraints of the MLC, an optimisation algorithm is required to determine the position of each leaf. Early optimisation methods took a geometrical approach, placing the MLC leaves based on the shifted position of the leaves (Sawant et al. 2008, Tacke et al. 2010). In current fluence based methods, such as direct (Ruan & Keall 2011) and piecewise (Moore et al. 2015) optimisation, the aperture position is determined by minimising the regions of under and over exposure. Minimising based on fluence allows for weighting to, for example, favour under-exposure over over-exposure, and account for general transformations (*e.g.* translation of the target) and tissue weighting.

Fluence based optimisation has a number of limitations. While it aims to improve dose conformity, it does not directly optimise the MLC based on the dose, which can lead to unintended consequences such as overdosing OAR in cases where the intrafraction motion causes overlap of the target and the OAR. It also does not accumulate such dose errors over time, which accumulate due to system lag and physical constraints of the MLC. Direct optimisation was extended to include previous error accumulation in the Beam's Eye View (BEV) plane (Wisotzky et al. 2015), though it only provided a modest improvement over the original method.

An optimisation based on the three-dimensional dose distribution, rather than the two-dimensional fluence distribution, has the potential to greatly improve the performance of MLC tracking. Recent work has focused on 4D real-time dose calculations, where the dose is calculated and accumulated during treatment, allowing

Dose Optimisation for MLC Tracking

3

for real-time monitoring of the impact of intrafraction motion. Using an in-house developed code, DoseTracker uses a simplified pencil beam convolution algorithm to calculate the dose during Volumetric Arc Therapy (VMAT) treatments (Ravkilde et al. 2014). Though DoseTracker's use of a simplified dose calculation creates uncertainty in the absolute dose, it can accurately calculate the difference between planned and accumulated dose (dose error) as compared to a treatment planning system (Ravkilde et al. 2014), providing real-time dose reconstructions for liver Stereotactic Body Radiation Therapy (SBRT) with phantom experiments (Ravkilde et al. 2018) and *in silico* with previous clinical treatments (Skouboe et al. 2019). Similarly, Fast et al. (2016) provided real-time dose reconstruction using pre-calculated dose influence data generated by a Monte Carlo dose calculation engine on prostate step-and-shoot Intensity Modulated Radiation Therapy (IMRT) treatments, which was further extended to lung SBRT (Kamerling et al. 2016) and prostate VMAT (Kamerling et al. 2017). A new online motion management workflow for MRI-Linac has been proposed in which IMRT treatments are adapted on the fly to account for inter- and intrafraction motion (Kontaxis et al. 2015a, Kontaxis et al. 2015b). Rather than adapt in real-time during treatment, on the fly treatment replanning is performed in between IMRT segments. The discrepancy between accumulated dose from previous segments and the planned dose is used to calculate new segments to best account for motion induced differences. Their methodology has since been applied to assess the feasibility of such a workflow on free-breathing SBRT treatment (Kontaxis et al. 2017).

The above work shows that the use of dose rather than fluence for treatment adaptation is emerging. The motivation for the use of dose optimisation for MLC tracking in this study is based on the residual target and OAR dose variations across multiple cancer sites reported in clinical trials employing MLC tracking with fluence optimisation. The dose reconstruction results from a prostate cancer trial using Kilovoltage Intrafraction Monitoring (KIM) guided MLC tracking showed that with MLC tracking, the clinical target volume (CTV) maximum dose (D2) was up to 5% higher than planned and the minimum dose (D98) was up to 3% lower than planned. Larger variations were seen in the bladder and rectum doses (Hewson et al. 2020). Similar variations were seen with electromagnetic transponder-guided (Calypso, Varian Medical Systems) MLC tracking for prostate cancer treatment (Colvill et al. 2015). In a lung cancer clinical trial employing electromagnetic transponder-guided MLC tracking, 3% fraction to fraction variations in the delivered gross tumour volume (GTV) dose were observed, with 2% underdose and 3% overdose of the GTV shown (Booth et al. 2016). In simulations on a liver cancer dataset, KIM-guided MLC tracking resulted in an average 2.4% overdose to the centre of the CTV (Poulsen et al. 2020). Although these variations are admittedly small, they do contribute to the total treatment uncertainty, making the treatment delivered deviate further from the planned prescription and add uncertainty to our understanding of the target and normal tissue dose-volume relationships and toxicity thresholds.

In this work, we present a novel MLC tracking algorithm which optimises the MLC

Dose Optimisation for MLC Tracking

aperture based on dose error accumulated in the moving patient anatomy. By optimising the aperture based on accumulated dose, errors due to physical MLC constraints are automatically captured and taken into account. Using a simplified dose accumulation calculation, this algorithm achieves real time performance allowing it to be used during a radiotherapy treatment. Here, we describe the new method and apply it to a prostate VMAT dataset to analyse its performance *in silico*.

2. Methodology

The methodology for dose-based MLC tracking is summarised as follows. The following steps occur within a single iteration of the optimisation routine, which determine the MLC aperture that best delivers dose to match the planned dose:

- (i) The planned dose is calculated using the patient treatment plan in the unshifted dose volume (fig. 1.1).
- (ii) The target position is acquired through motion tracking and the dose volume is shifted accordingly (fig. 1.2).
- (iii) The 3D dose is integrated onto the 2D BEV (fig. 1.3).
- (iv) For each leaf track, the MLC aperture is fitted by minimising the cost function, constrained by finite leaf velocities (fig. 1.4).
- (v) The delivered dose is calculated using the new aperture (fig. 1.5).
- (vi) The process repeats from step 1 with new gantry and target positions, until treatment finishes.

Each iteration of dose optimisation corresponds to a fixed timestep, $\Delta t = 50$ ms, stepping the iteration forward in time. This timestep corresponds to the operational frequency of the treatment machine. Between each iteration of dose optimisation, the gantry angle and target position are retrieved and used for the next iteration. In real experimental or treatment scenarios, the gantry angle and target position would be retrieved directly from the treatment machine and target positioning method. In this *in silico* study, these have been simulated to provide a similar level of information as would be provided in an experimental or treatment scenario.

Dose is accumulated on a three-dimensional point cloud, referred to as the dose points, which span a sub-volume of the patient's body that is predominantly in the line-of-sight of the radiation beam. The dose points are placed with uniform spacing of 2 mm in a staggered grid. Each repeating cell consists of a cube with dose points at the corners, and a single dose point in the centre. All dose points are confined to within a cylinder whose axis is symmetric about gantry rotation axis, as illustrated by the cylinder in fig. 1. The cylinder's dimensions are set using the size of the jaws: this ensures the dose points are predominantly in the aperture of the MLC. However, these points have no underlying topology or connectivity, thus can be moulded into any shape required, and can move independently of each other if required.

Dose Optimisation for MLC Tracking

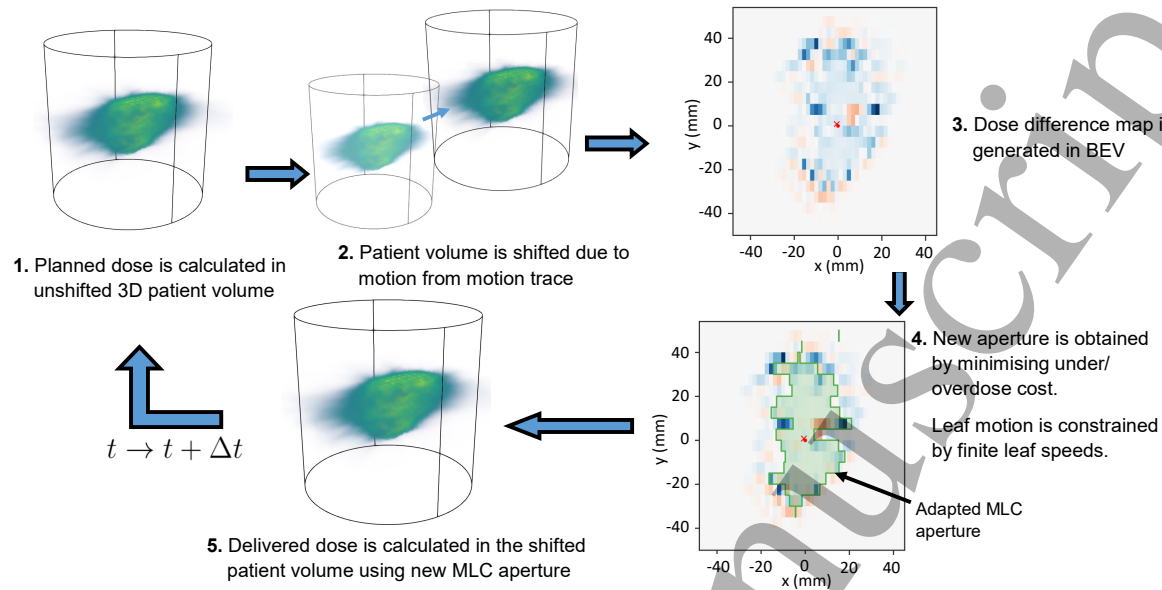


Figure 1. The steps within an iteration of the dose optimisation algorithm. Starting at step 1, dose optimisation calculates the planned dose (1), shifts the dose volume according to the intrafraction motion (2), optimises based on the beam's eye view (BEV) dose difference (3-4) then calculates the delivered dose using the optimised aperture (5), then repeats the cycle until treatment finished.

Additionally, a two-dimensional BEV grid is defined on which the MLC aperture is optimised. The BEV grid is in the IEC Beam Limiting Device (BLD) coordinate system (IEC 61217:2011), with the x axis in the direction of the MLC leaf motion, the z axis in the direction from the isocentre to the radiation source, and the y axis completing the right handed set. In the y direction, the grid is aligned with the MLC itself, each leaf track corresponding to a set of pixels on the grid along the y direction. Along each leaf track (x direction), the grid has a uniform spacing of 2 mm. Each of these pixels correspond to an integral of the 3D dose points, along the direction normal to the BEV grid, the z direction.

2.1. The 3D Dose Calculation

Approximations have been made to render the dose calculation feasible for real time operation on the set of dose points. Equation 1 describes a simple line of sight dose calculation, where the gantry, collimator and leaf positions are considered constant within a set timestep, Δt . The dose calculation does not account for beam divergence or attenuation, or tissue composition. VMAT and IMRT functionality is enabled by many of these dose calculations, then stepping forward in time by Δt , where the MLC leaves and gantry are moved.

$$d(x, y, z) = \begin{cases} d\Delta t, & \text{where } x_{j,B} < x < x_{j,A}, y_j < y < y_{j+1} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

Dose Optimisation for MLC Tracking

6

where d is the dose at position (x, y, z) , in the BLD coordinate system, \dot{d} is the dose rate, y_j is the lower bound of leaf track j , and Δt is the timestep. $x_{j,B}$ and $x_{j,A}$ are the two leaf positions of the j^{th} MLC leaf track. The dependence of the dose calculation on the gantry angle is present in the transformation of the patient coordinate system to the BLD coordinate system. Ideally, the dose rate would be queried from the treatment machine. However, for this work it is set to a constant value, acting as a scaling factor.

The dose points are shifted according to the interfraction motion trace. The planned dose is calculated on the un-shifted dose points, which is then moved as indicated by the motion trace. The delivered dose is calculated on the shifted dose points with the dose-optimised MLC aperture.

In determining the value of the timestep, Δt , the fastest speeds and typical distances in the system need to be considered. In this work, Δt is set to 50 ms: in this timestep, the leaves can move at most 1.8 mm (assuming a maximum MLC leaf speed of 3.6 cm/s), approximately equal to the dose point spacing.

2.2. Reduction from 3D to 2D

While dose is accumulated in 3D, the MLC motion is confined to the two-dimensional BEV plane. Therefore, a method to transform the 3D dose information to BEV is required in order to determine the aperture position.

Dose optimisation aims to minimise the difference between the planned (d_P) and delivered doses (d_D) squared, accumulated up until the current treatment time,

$$C(\vec{x}_B, \vec{x}_A) = \int \int \int (d_D^t(\vec{x}, \vec{x}_m, \vec{x}_B, \vec{x}_A) - d_P^t(\vec{x}))^2 dV, \quad (2)$$

where C is the cost, the integration is over the patient volume (V). The delivered dose (d_D) is a function of position in the patient volume ($\vec{x} = (x, y, z)$), observed target position (\vec{x}_m), and MLC leaf positions \vec{x}_B and \vec{x}_A .

Due to the dose calculation used, equation 2 is substantially simplified. The use of the line-of-sight stationary dose calculation (equation 1) allows each dose point to map to a unique leaf track, *i.e.* those points are only ‘dosed’ by a given leaf track. Hence, the optimisation problem is reduced to a set of independent cost functions for each leaf track, with each given leaf track (with index j) having its own associated cost,

$$C_j(x_{j,B}, x_{j,A}) = \int (D_D^t(x, x_{j,B}, x_{j,A}) - D_P^t(x))^2 dx, \quad (3)$$

where D^t is the dose integrated along the z direction at timestep t .

With some rearrangement, equation 3 leads to the final version of the dose optimisation,

$$C_j(x_{j,B}, x_{j,A}) = \int (\Delta D_j(x, x_{j,B}, x_{j,A}) + E_j(x))^2 dx, \quad (4)$$

where C_j is the cost function for a given leaf track with index j , ΔD_j is the dose to be delivered in timestep t , and E_j is the dose difference: the difference between delivered dose up to the previous timestep ($t - 1$) and the planned dose at this timestep (t), $D_D^{t-1} - D_P^t$.

Dose Optimisation for MLC Tracking

7

Up until now, the equations have dealt with continuous integrals of the dose volume. However, the dose volume is made up of discrete points in three-dimensional space. The quantities ΔD_j and E_j are integrated from the three-dimensional dose points to the two-dimensional BEV grid, illustrated in fig. 2 for a single leaf track.

The BEV grid is uniformly spaced along the x axis. The y axis spacing follows the spacing of the MLC leaf tracks, *e.g.* a Varian Millennium 120 MLC (Varian Medical Systems) contains forty pairs of leaves of width 5 mm, with two sets of ten leaves of width 10 mm on either side.

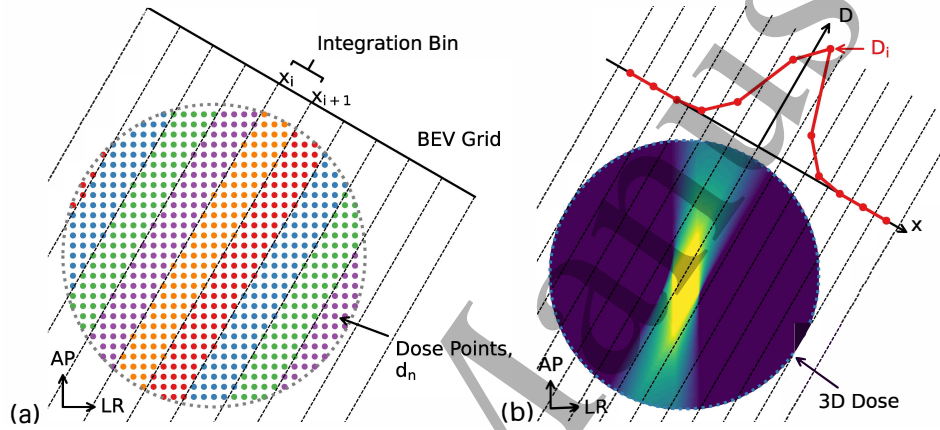


Figure 2. Integration of the 3D dose volume onto the BEV plane for a single leaf track. Each subplot shows 2D slices of the 3D dose volume corresponding to a given leaf track. The solid straight line denotes the BEV for a single leaf track. (a) shows the binning of dose points, d_n , into each BEV pixel. *e.g.* the BEV pixel from x_i to x_{i+1} integrates each of the orange dose points in its view. (b) shows the dose distribution in that slice. The line plot shows the integral of that dose distribution onto the BEV for a single leaf track.

Since the 3D dose is a set of points, the integration is a sum of all dose points inside those bins. For the BEV pixel, i , bounded by positions x_i and x_{i+1} , the integral of the dose points d_n is given by the sum of all the dose points within those bounds. As illustrated in fig. 2a, the sum of the orange dose points map to the BEV pixel bounded by x_i and x_{i+1} . An example of a 3D dose distribution corresponding to a given leaf track, and its corresponding BEV integrated dose (the red line plotted on the x and D axes) is shown in fig. 2b.

Hence, ΔD and E in equation 4, are expressed as follows,

$$E_{i,j} = \sum_n d_{D,n}^{t-1} - d_{P,n}^t, \text{ where } x_i \leq x_n < x_{i+1}, y_j \leq y_n < y_{j+1} \quad (5)$$

$$\Delta D_{i,j} = \sum_n \dot{d} \Delta t, \text{ where } x_i \leq x_n < x_{i+1}, y_j \leq y_n < y_{j+1} \quad (6)$$

where the summation is over all dose points and i, j is the index of the BEV pixel.

The discrete version of equation 4 is given as follows,

$$C_{i,j}(x_B, x_A) = \sum_i (\Delta D_{i,j}(x_B, x_A) + E_{i,j})^2 \Delta x, \quad (7)$$

Dose Optimisation for MLC Tracking

where the summation along a leaf track (over pixels in the x direction) and Δx is the pixel size.

By differentiating the discrete cost function (equation 7), and assuming the dose applied to a pixel varies linearly with leaf position inside that pixel,

$$\Delta D_i(x_B, x_A) = \Delta D_i \frac{\min(x_A, x_{i+1}) - \max(x_B, x_i)}{\Delta x_i} \quad (8)$$

the following expressions are obtained for the position of the leaf,

$$x_B = x_{i_B+1} + \Delta x \frac{E_{i_B}}{\Delta D_{i_B}} \quad (9)$$

$$x_A = x_{i_A} - \Delta x \frac{E_{i_A}}{\Delta D_{i_A}} \quad (10)$$

where x_B and x_A are the leaf positions, x_i is the lower bound of the BEV pixel, and i_B is the index of the BEV pixel containing the position x_B (*e.g.* $x_{i_B} \leq x_B < x_{i_B+1}$) and similarly for x_A . The difference between the indices in positions x_{i_B+1} and x_{i_A} is due to the impact of moving a leaf in bank B is opposite to bank A. Moving x_B in the positive x direction decreases the aperture size and hence the dose, whereas moving x_A in the same direction opens the aperture and increases the dose.

Equations 9 and 10 specify the optimal position of the leaves for a given BEV pixel (*i.e.* for pixel i , the optimal position for x_B is given by equation 9 where $i_B = i$, and similarly for x_A). However, not all indices correspond to the global minimum. In order to obtain the global minimum, the set of optimal positions are computed by applying equations 9 and 10 for each pixel. The cost is computed for each x_B and x_A pair: the x_B and x_A positions with the lowest cost are chosen as the optimised aperture positions.

2.2.1. Dose Optimisation With the cost function established, a method is outlined by which to fit the leaf, update the dose, and perform multiple iterations of the optimisation. Consider a 3D dose volume that is mid-fraction. Dose has already been delivered to the dose points, and the plan provides the target dose distribution. However, due to intrafraction motion, the delivered dose and the planned dose are mismatched, as is illustrated in fig. 3. Fig. 3a shows the planned dose, which has more dose than the delivered dose (fig. 3b): this is due to the planned dose being one timestep ahead. Dose optimisation calculates the aperture that brings the delivered dose closest to the planned. The final panel, fig. 3c, shows the dose difference, with underdose in blue and overdose in red.

These doses are integrated along the z direction in order to generate the BEV dose distributions shown in fig. 4a. Both the planned dose (in green) and the accumulated dose (shown as purple dashed), look similar with a peak in dose at $x_B \approx 0$. The planned dose is larger in magnitude, as it indicates the target dose for this timestep (D_P^t), whereas the accumulated dose indicates the dose at the previous timestep (D_D^{t-1}).

Fig. 4b shows the difference between delivered and planned. It indicates where the distribution is under- and overdosed, suggesting where dose can be recouped. Areas of

Dose Optimisation for MLC Tracking

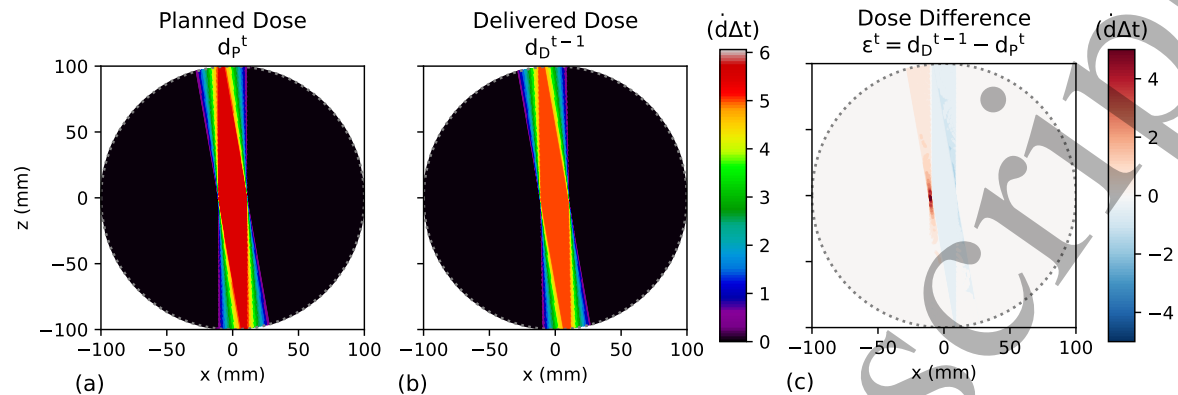


Figure 3. Slices of the 3D dose points mapped to a single leaf track, at the start of the fitting procedure. The colour represents the dose interpolated between dose points of: (a) the planned dose at timestep t , (b) the dose that has been delivered up until this timestep ($t - 1$), and (c) the difference between delivered dose and planned dose, which makes up the cost function. Regions of red show areas of overdose, while regions of blue show areas of underdose. The colour map for (a) and (b) has been chosen to accentuate difference between planned and delivered.

underdose are shown in blue, and overdose in red. The leaf fitting algorithm should place the leaf aperture in the regions of blue so as to reduce the underdose, while avoiding areas of overdose as much as possible, as this would only exacerbate the overdose.

To see where the aperture is best placed, the cost function is plotted (equation 7) as a function of the leaf positions, x_B and x_A , as shown in fig. 4d. The cost function is split into two quadrants, separated by the $x_B = x_A$ line. All points below this line are not considered, as they correspond to the overlapping of MLC leaves which is not physically possible. Above the $x_B = x_A$ line, the cost generally increases with distance from the line, which corresponds to the aperture opening. As the aperture opens ($x_B \rightarrow -\infty, x_A \rightarrow \infty$), the cost increases, indicating overdose.

Closing the aperture ($x_B = x_A$) generally has a lower cost than keeping it open. However, there is a global minimum at $x_B \approx -7, x_A \approx 13$, indicating this is the best position for the leaves. This is illustrated in fig. 4.a-c by purple vertical bars in each panel. With these leaf positions, the fitted dose (the delivered dose at timestep t), conforms well to the planned dose, as shown by the orange line in fig. 4a. By subtracting the fitted dose from the planned dose to create the dose error, fig. 4c shows that the error has been reduced, mainly by applying dose in the region of underdose. However, it has created a small region of overdose in the region $-7 < x < 0$. In regions where the change in dose, ΔD , is greater than the dose difference, there is a trade-off between reducing the underdose and creating overdose.

With the new aperture obtained through dose optimisation, the dose is updated in the three-dimensional patient volume. Fig. 5 shows the dose distributions at the end of the dose optimisation timestep, with the delivered dose in fig. 5b and dose error in fig. 5c. The planned dose, fig. 3a, is repeated here for easy comparison in fig. 5a.

Dose Optimisation for MLC Tracking

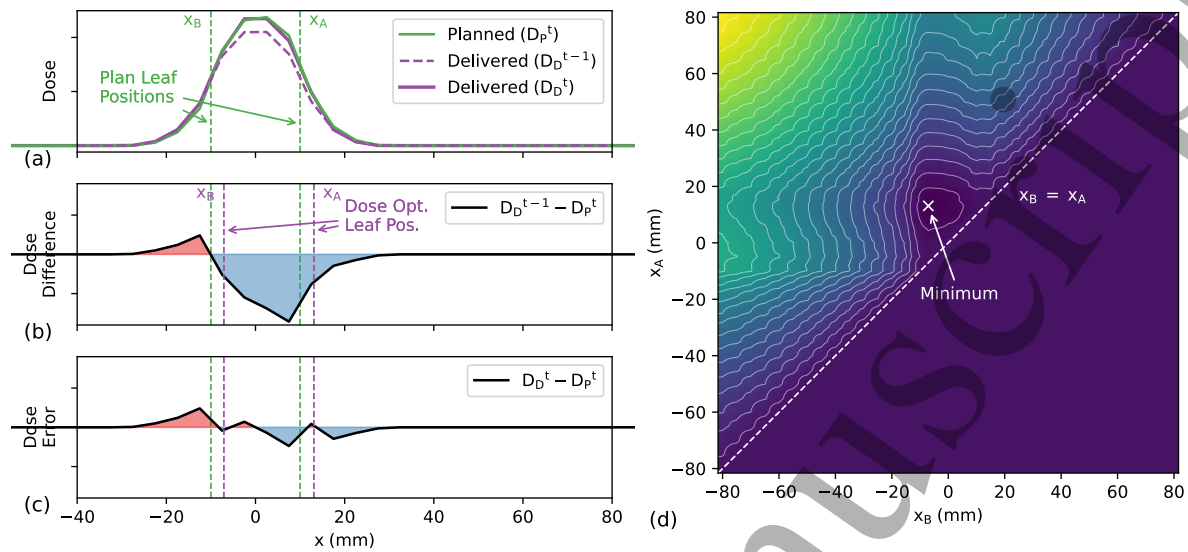


Figure 4. The integral of the dose (as illustrated in fig. 3) along a leaf track. (a) shows the planned dose at timestep t in green, the delivered dose accumulated until timestep $t - 1$ in purple dashed, and the delivered dose at the current timestep t in purple. (b) The dose difference: delivered minus planned, with underdose shown in blue and overdose in red. This is used for determining the positions of the aperture, x_B and x_A . (c) After dose optimisation determines the aperture positions, the delivered dose is updated with the new aperture to timestep t . The planned dose is then subtracted from the fitted dose to show the dose error. (d) Shows the cost as a function of the leaf positions (x_B and x_A), with yellow indicating high cost and blue low cost. The optimal leaf positions are at the point where the cost is at a minimum.

Comparing the planned dose (fig. 5a) and delivered dose now (fig. 5b), they now look much more similar, with the central region in dark red aligning well. There are still differences, especially at the edge of this central region. This is accentuated by the dose error, fig. 5c, showing a region of high overdose (in red) at $x \approx -10$, $y \approx 0$, and a large region of low underdose (in blue). These errors are carried forward into the next timestep, where the leaf fitting procedure begins anew with the updated delivered and planned doses.

2.2.2. Finite Leaf Velocities In the current formulation, the fitting procedure does not take finite leaf velocities into account: *i.e.* the fitted leaf positions can ‘jump’ to the position that minimises the cost function (equation 7), regardless of whether the MLC leaves can reach that position in the allotted timestep. The MLC leaf speeds, typically of up to 3.6 cm/s, are considered slow enough to adversely impact performance (Wijesooriya et al. 2005).

Rather than constraining the range of fitted apertures to only those within reach, dose optimisation sets the target aperture and constrains the motion of the leaves. If the algorithm returns leaf positions out of reach, the leaf positions will move towards the target fitted leaf positions. For target apertures completely outside the current aperture

Dose Optimisation for MLC Tracking

11

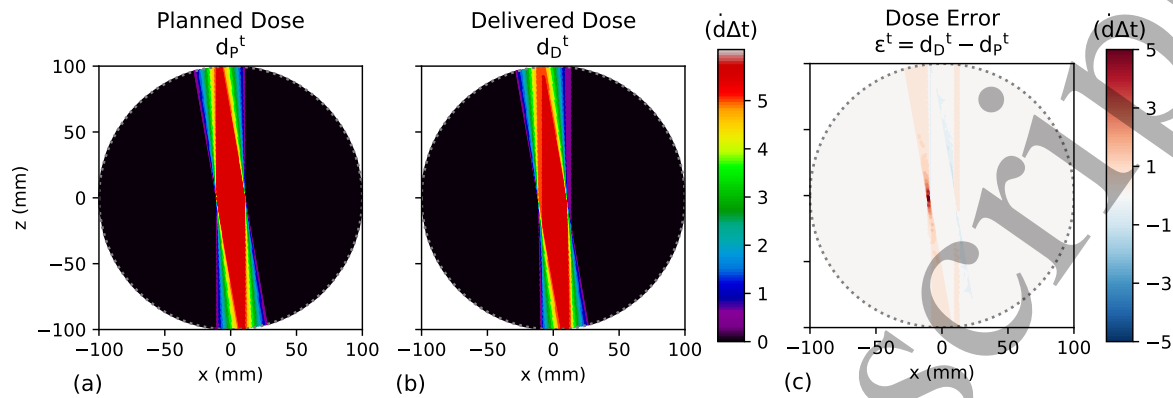


Figure 5. Doses at the end of the fitting procedure, similar to fig. 3. (b) instead shows the fitted dose, which is now similar in dose to the planned dose (a). (c) shows the dose error, which highlights the differences between planned and fitted doses.

positions, the aperture will close but move towards the target aperture: this was found to avoid unnecessary overdosing. The aperture is allowed to move to new regions, though will not dose optimally while the leaves move toward this target. However, this dose error is accumulated and hence will factor into the leaf fitting at later timesteps.

2.3. Treatment Simulation

The method was applied to and benchmarked against prostate cancer VMAT treatments from the TROG 15.01 SPARK trial (Keall et al. 2020), where intrafraction motion was observed. Fifteen fractions, with two arcs each, were used to comparing the three methods: dose optimisation, fluence optimisation (Ruan & Keall 2011), and without MLC tracking. The dose calculation method allows for fast dose calculation with an average computation time of 14.6 ± 3.6 ms per aperture for two dose calculations (planned and actual), without any significant performance enhancements such as multithreading or GPU implementations. MLC aperture optimisation, which includes the integration of the 3D dose distribution along the BEV, took 24.5 ± 5.3 ms, taking the total computation time to 39.1 ± 8.5 per timestep.

At the start of each fraction, the dose points and BEV grid are generated based on the size of the jaws of the given patient treatment plan. The patient treatment plan is also loaded to provide the planned MLC leaf positions and gantry angles for the planned dose calculation. During a treatment, the intrafraction motion would be obtained in real time using a tracking method. However, in this *in silico* study, a motion trace is imported from file, and simulated to provide a similar level of information as would be provided by a motion trace obtained from a tracking method. Similarly, the motion of the gantry, which would usually be queried directly from the treatment machine during treatment, was determined from treatment machine log files.

The method of accumulating dose *in silico*, used for the analysis, is the same for each of the optimisation methods used: the only difference is the aperture used,

Dose Optimisation for MLC Tracking

12

which is obtained through each respective optimisation method. For MLC tracking with dose optimisation, the aperture positions were obtained using the method described in section 2.2.1. The TROG 15.01 SPARK trial used fluence optimisation, specifically the direct optimisation method (Ruan & Keall 2011), for MLC tracking. Hence, for fluence optimisation, the apertures were taken from the log files generated by the treatment machine. Without MLC tracking, the apertures were directly interpolated from the patient treatment plan. In both the fluence optimisation and without tracking cases, their respective apertures were used to calculate the dose in the patient volume for comparison purposes.

2.4. Fraction Analysis

In order to quantitatively assess the performance of dose optimisation and compare with other methods, the dose error and the γ failure rate (Low et al. 1998) is used. For the γ analysis, a threshold of 2 mm/2% was used. The doses used in these calculations are calculated using the dose calculation given in equation 1.

The dose error, ϵ_T , is the sum over the entire dose volume of the absolute difference between the delivered and planned doses, normalised over the total planned dose, as given by equation 11,

$$\epsilon_T = \frac{\sum_i |d_D^i - d_P^i|}{\sum_i d_P^i}, \text{ where } d_P^i \geq d_{\min}, \quad (11)$$

where d_{\min} is the dose threshold. The dose threshold removes points with low dose values, and is set at 10% of the maximum planned dose.

To assess the statistical significance of the performance of each method, we perform a Wilcoxon signed-rank test ($p < 0.01$) to determine whether dose optimisation produces lower dose errors and γ failure rates than fluence optimisation for each fraction.

3. Results

The performance of the three methods, MLC tracking with dose optimisation, with fluence optimisation and without MLC tracking, was compared. Firstly, the rate of increase in dose error was calculated over the course of treatment for two fractions, one with representative motion for the dataset and one with little motion. Fig. 6 shows the target position (6.1a and 6.2a) used to shift the dose points, and the dose error over time (6.1b and 6.2b).

Fraction P06-F02 (fig. 6.1) is a fraction with typical motion for the dataset. Fig. 6.1.1 shows the motion trace, which ranges from -2 mm to 2 mm. The target position was offset from its original position by superior-inferior (SI) ~ 2 mm, left-right (LR) ~ 0.2 mm, anterior-posterior (AP) ~ -1 mm, with variation in target position especially later in the treatment. As shown in fig. 6.1b, the dose error started at zero, as no dose has been delivered to the patient. However, as treatment progressed, dose was continually delivered, and the dose error increased for all methods as shown. The

Dose Optimisation for MLC Tracking

13

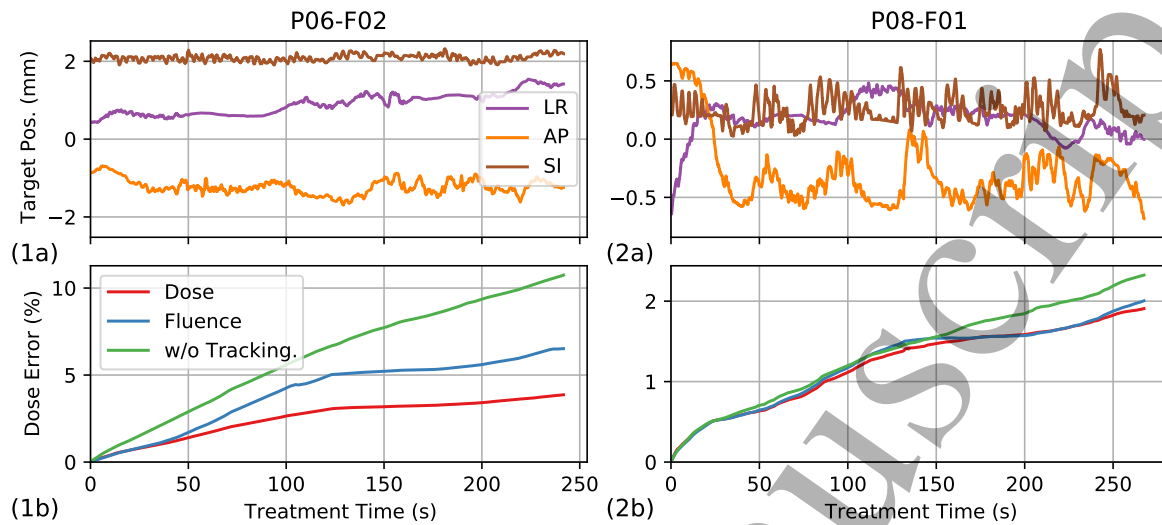


Figure 6. The target position trace in the left-right (LR), anterior-posterior (AP), and superior-inferior (SI) directions (1) and the total dose error (2) over treatment time for patient 6, fraction 4 (a) and patient 8, fraction 1 (b) showing the three methods used in this analysis. The dose error is normalised to the total planned dose.

dose error increased most steeply without MLC tracking (green), leading to the largest dose error at the end of the fraction (10.7% of the planned dose). The large dose error is expected as no motion mitigation is used and is considered the baseline case.

The dose error for both the fluence (blue) and dose (red) optimisation methods both rose at the same rate until $t \sim 40$ s, after which they separated, with fluence optimisation causing more dose error than dose optimisation. The similarity of the two methods for early treatment times suggests that dose optimisation method tends towards fluence optimisation for low dose accumulation. At early treatment times, the amount of dose delivered is small and the gantry has not rotated far from its starting position, hence the BEV integrated dose is approximately proportional to the delivered fluence. At later times, dose optimisation is able to use the 3D dose error information to better inform the aperture position, as shown by the slower rise in error. Dose optimisation achieves a final error of 3.9%, which was lower than with fluence optimisation, 6.5%.

Fraction P08-F01 (fig. 6.2), on the other hand, shows little motion, with a maximum motion of ~ 1 mm in each direction. While the target position remained offset at ~ 0.25 mm in the SI direction, it moved in the LR and SI directions in the first ~ 40 s of treatment. The dose errors in this fraction for each method rose at the same rate before separating at $t \sim 40$ s. Overall, the dose error due to dose optimisation (red) remained the lowest of all methods. However, fluence optimisation (blue) was comparable, agreeing with the dose optimisation curve at times during the treatment. At the end of the treatment, the dose error for all methods was low, with dose optimisation at 1.9%, fluence optimisation at 2.0% and without MLC tracking at 2.4%.

Fig. 7 shows axial slices of dose in the patient volume at the end of the same fraction. The slice is positioned half way along the PTV in the SI direction. Fig. 7a

Dose Optimisation for MLC Tracking




Figure 7. Axial slices of the patient volume, showing (a) planned dose distribution, and dose errors using dose optimisation (b), fluence optimisation (c) and without MLC tracking (d). The colour scale of each slice have been normalised to the maximum planned dose. Additionally, slices of the PTV, rectum and bladder are superimposed.

shows the planned dose calculated on the set of dose points. The regions of high dose (yellow regions) align well with the superimposed slice of the PTV surface, and avoid the rectum and bladder as expected. This suggests that, though the dose calculation is simplified, it captures the delivered dose adequately. Fig. 7b-d show the dose errors for each of the three methods. For this fraction, without MLC tracking (fig. 7d) shows the highest dose error, with pronounced regions of under- (blue) and overdose (red). Fluence optimisation (fig. 7c) shows an improvement over without MLC tracking, as the regions of dose error are fainter in colour. Dose optimisation performed the best (fig. 7b) with dose errors closest to zero.

3.1. Fraction Comparison

To validate the method against all fractions, the dose error and γ failure rate (Low et al. 1998) were computed and compared. This was done for all three methods: dose optimisation, fluence optimisation and without MLC tracking.

The results for each fraction are shown in fig. 8a. In almost all fractions, dose optimisation performed best, achieving a lower dose error than both fluence optimisation and the baseline without MLC tracking cases.

Fig. 8b shows the aggregate data of fig. 8a, presenting the overall performance of each method. Overall, dose optimisation performed best in this analysis, indicated by the lower mean dose error (shown the be middle dashed line). Fluence and without MLC tracking had a much larger spread of dose errors than dose optimisation. On average, dose optimisation achieves a dose error of $3.4\% \pm 0.6\%$, improving over fluence optimisation ($4.4\% \pm 1.1\%$) and without MLC tracking ($7.2\% \pm 3.4\%$).

Fig. 9 shows the γ 2mm/2% failure rates for these fractions. It shows a similar picture to that of the dose error results: dose optimisation performed better for every case, as is shown by the data for the individuals cases (fig. 9a) and the aggregate data (fig. 9b). Dose optimisation achieved γ failure rates of $4.7\% \pm 1.2\%$, improving over fluence optimisation ($7.5\% \pm 2.9\%$) and without MLC tracking ($15.3\% \pm 12.9\%$).

Dose Optimisation for MLC Tracking

15

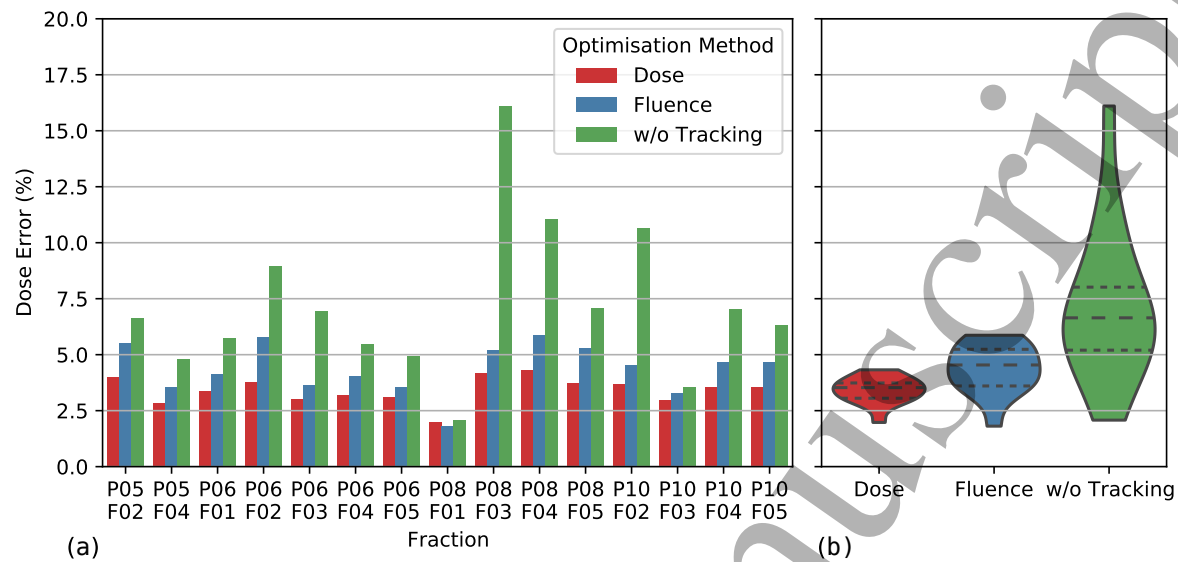


Figure 8. Dose errors, as given by equation 11. (a) shows the results for each fraction in this study. (b) shows violin plots of the aggregate data in (a). The width of the plot shows the distribution of points, with thickness indicating the number of points at that value. The dashed lines indicate the quartiles of the dataset.

The γ failure rates and dose errors were both significantly lower with the dose optimisation method compared with fluence optimisation, as assessed with a Wilcoxon signed-ranks test ($p < 0.01, T = 1$; $p < 0.01, T = 0$ respectively).

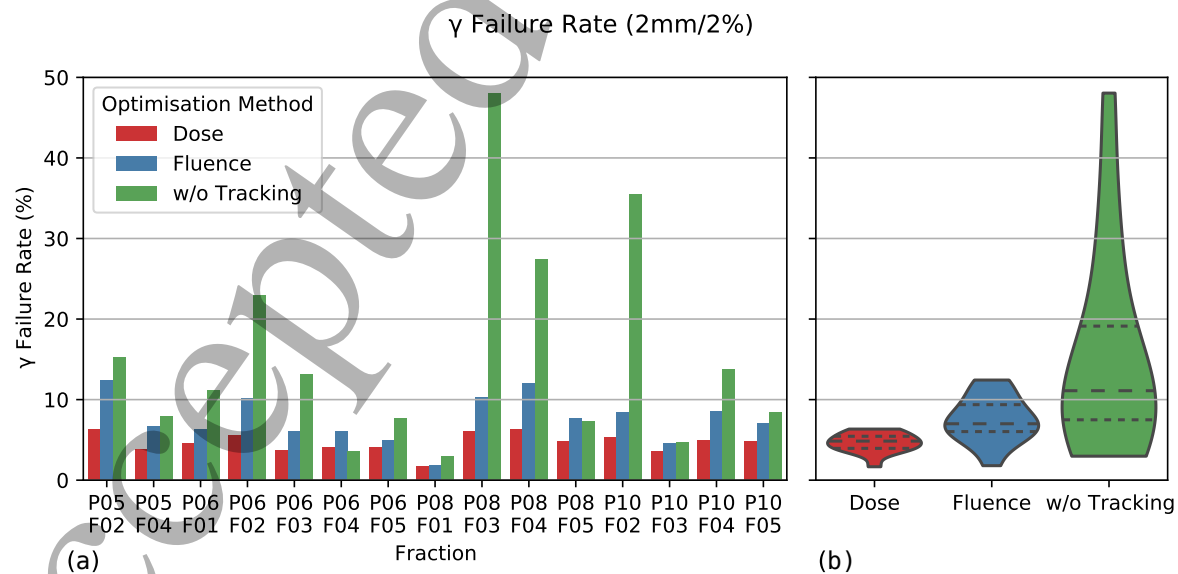


Figure 9. γ (2mm/2%) failure rates. (a) shows the results for each fraction in this study. (b) shows violin plots of the aggregate data in (1). The width of the plot shows the distribution of points, with thickness indicating the number of points at that value. The dashed lines indicate the quartiles of the dataset.

4. Discussion

In this study, we have detailed the dose optimisation method, and applied it to a prostate VMAT dataset. Dose optimisation is similar to fluence optimisation for MLC tracking (*e.g.* direct (Ruan & Keall 2011) and piecewise (Moore et al. 2015) optimisation) in that it optimises the MLC aperture to account for intrafraction motion. Rather than using difference between planned fluence and motion shifted fluence, dose optimisation optimises based on the difference between planned dose and the delivered dose, accounting for accumulated dose errors due to intrafraction motion. The planned dose is calculated, then shifted according to motion. Integrating the 3D dose volume onto the 2D BEV plane, allows for the MLC aperture position to be determined. This new aperture is then used to update the delivered dose. These steps repeat until the treatment completes.

When applied to a prostate VMAT dataset *in silico*, dose optimisation is able to target regions of underdose while avoiding regions overdose. These errors are not captured as well in the fluence optimisation and in the baseline without MLC tracking. Both metrics of total dose error (difference between planned and delivered doses) and γ failure rates at the end of each fraction show that dose optimisation enabled more accurate dose delivery to the planned dose compared with fluence optimisation.

In this work, the dose calculation and aperture optimisation step is assumed to have completed within this 50 ms timestep. However, in cases where the process takes considerably longer, such as when using a more complex dose calculation, prediction algorithms would be required to estimate the position of the target ahead of time, using similar methods that account for MLC latencies in the system (Poulsen et al. 2009, Ruan & Keall 2010, Caillet et al. 2020). This estimation would be used to calculate the dose error and hence optimise the aperture.

By explicitly optimising the MLC aperture based on the dose delivered, dose optimisation can provide valuable insight into and improvements to other tumour locations and multi-target applications. Complex motion can occur in lung and liver tumour sites, where the motion can have a large range and be deformable (Bertholet et al. 2016, Caillet et al. 2017). For large motions, the dose optimisation would accumulate the induced dose errors and aim to optimise the aperture to account for these. However, limitations in the MLC design such as finite leaf widths and velocities would potentially limit the ability to account for these errors. In cases where motion deforms the underlying structures, the underlying 3D dose distribution would need to deform as observed (Kamerling et al. 2016). Dose calculations algorithms which calculate dose on a point cloud, such as in this work, can move those points independently to account for deformation. Such motion would require a deformation vector field in order to deform the dose grid, which would require a more complex motion tracking methods (Glitzner et al. 2015, Hunt et al. 2018). In treatments with multiple targets, such as in patients with locally advanced cancers (Machtay et al. 2012), such targets can move or deform independently to one another. In locally advanced prostate

cancer treatment, where treatment of both the prostate and the nearby lymph nodes is required, the prostate exhibits significant intrafraction motion while the lymph nodes are relatively stationary. By optimising based on dose, this difference in intrafraction motion can be accounted for by moving and weighting dose points accordingly and implementing these into the cost function (equation 3). Furthermore, by using the excellent soft tissue contrast of MRI-Linac machines (Hunt et al. 2018), dose optimisation is well placed to adapt for complex intrafraction motion and deformation.

In order to fully optimise for dose delivered, a more advanced dose calculation algorithm (such as those found in treatment planning systems) could be used. This means that, due to the simplified dose calculation on which dose optimisation is based, the optimised MLC aperture does not necessarily optimise for actual dose delivered to the target. However, Ravkilde et al. (2014) observe dose errors from their simplified dose reconstruction algorithm agree with dose errors as calculated from a treatment planning system. The method of calculating dose in this work is similar to that of DoseTracker's (Ravkilde et al. 2014), albeit further simplified. DoseTracker uses a simplified pencil beam algorithm that includes beam divergence and attenuation, allowing a temporal resolution of 500 ms. This work's dose calculation method also differs to DoseTracker in that it calculates the dose in healthy tissue, rather than only the target and organs at risk. The complexity of the dose calculation algorithm becomes more important in regions of complex tissue composition, such as the lung, where density variations affect the dose distribution (Kamerling et al. 2016). The dose from such a simplified dose calculation would likely differ considerably from the actual dose in these cases, and may benefit from a more accurate dose calculation algorithm. Increasing the accuracy and realism of the dose calculation would likely increase the computation time, either introducing errors due to the increased latency or making real-time performance unattainable. Hence, there is a balance to be struck between dose accuracy and speed of computation. Further analysis is required to assess the true dosimetric impact of dose optimisation and to what extent a more complex dose calculation algorithm is required, especially in other tumour sites such as lung and liver.

While dose optimisation reduces dose error due to intrafraction motion over fluence based methods, it does not completely eliminate the error. Physical limitations of the MLC's design causes inherent sources of error, such as the finite size of the MLC leaf widths (5 mm in the case of this work), and finite leaf velocities. One way of managing finite leaf velocities is to only consider MLC apertures that are attainable in the timestep, as discussed as a potential improvement over the previous optimisation methods (Moore et al. 2015). This method proved ineffective for dose optimisation as it made the aperture position stick to its current location, ultimately generating excessive overdose. Consider a case similar to that in fig. 4, but with two separated regions of underdose. At the first instance, the algorithm would select whichever underdose trough is best to minimise (lowest cost), and start to dose that region. Eventually it may be able to dose that region enough such that it is no longer underdosed. However, the other region, which has as yet not received any dose, is still underdosed and cannot be reached

Dose Optimisation for MLC Tracking 18

by the algorithm as it is outside of the region that is attainable due to the finite leaf velocity. By limiting the range of the leaf fitting algorithm to only positions the leaves can reach, the leaves can never reach this second region of underdose. Moore et al. (2015) acknowledged this situation, suggesting to apply weighting to solutions that could not be attained. Constraining the motion of the aperture, rather than the available positions, was found to improve dose conformance. The current form of dose optimisation only considers dose accumulated up until the current timestep. The Kontaxis et al. (2015a) approach re-plans the next IMRT segment, generating an MLC sequence that converges well to deliver the originally planned dose. This highlights the need for MLC tracking algorithms to take into account future dose and re-planning, and would potentially overcome the issue of finite leaf velocities and widths.

5. Conclusion

In this work, dose optimisation for MLC tracking has been implemented and applied to a dataset containing prostate VMAT treatments where motion was observed. Dose is accumulated in the three-dimensional patient volume, then integrated along the BEV. With an integrated dose map in the BEV plane, the aperture is then fitted based on minimising under- and overdose. Using a dataset of prostate VMAT treatments, the dose optimisation method was applied *in silico* and compared to existing MLC tracking methods. Dose optimisation successfully targets regions of underdose as determined by the simplified dose algorithm, causing lower dose errors than the existing clinically implemented fluence optimisation and the baseline case without MLC tracking. While the dose calculation in this work has been greatly simplified, the method for dose optimisation can be used with more accurate dose calculations algorithms.

6. Acknowledgements

This work was funded by a Cancer Council NSW grant (RG 19-10) and, in addition, PK and DTN are supported by fellowships from the NHMRC (1112096, 1138807) and Cancer Institute NSW (2018/ECF007). The authors thank Dr. Helen Ball for reviewing the manuscript.

Dose Optimisation for MLC Tracking

19

- Bertholet J, Worm E S, Fledelius W, Høyer M & Poulsen P R 2016 *International Journal of Radiation Oncology*Biological*Physics* **95**(2), 802–809.
- Booth J T, Caillet V, Hardcastle N, O'Brien R, Szymura K, Crasta C, Harris B, Haddad C, Eade T & Keall P J 2016 *Radiotherapy and Oncology* **121**(1), 19–25.
- Caillet V, Keall P J, Colvill E, Hardcastle N, O'Brien R, Szymura K & Booth J T 2017 *Radiotherapy and Oncology* **124**(1), 18–24.
- Caillet V, Zwan B, Briggs A, Hardcastle N, Szymura K, Prodreka A, O'Brien R, Harris B E, Greer P, Haddad C, Jayamanne D, Eade T, Booth J & Keall P 2020 *Physics in Medicine & Biology* **65**(23), 235040.
- Colvill E, Booth J T, O'Brien R T, Eade T N, Kneebone A B, Poulsen P R & Keall P J 2015 *International Journal of Radiation Oncology*Biological*Physics* **92**(5), 1141–1147.
- D'Souza W D, Naqvi S A & Yu C X 2005 *Physics in Medicine and Biology* **50**(17), 4021–4033.
- Fast M F, Kamerling C P, Ziegenhein P, Menten M J, Bedford J L, Nill S & Oelfke U 2016 *Physics in Medicine and Biology* **61**(4), 1546–1562.
- Fast M F, Nill S, Bedford J L & Oelfke U 2014 *Medical Physics* **41**(11), 1117–1119.
- Glitzner M, Crijns S P M, de Senneville B D, Kontaxis C, Prins F M, Lagendijk J J W & Raaymakers B W 2015 *Physics in Medicine and Biology* **60**(22), 8869–8883.
- Hansen R, Ravkilde T, Worm E S, Toftegaard J, Grau C, Macek K & Poulsen P R 2016 *Medical Physics* **43**(5), 2387–2398.
- Hewson E A, Nguyen D T, O'Brien R, Poulsen P R, Booth J T, Greer P, Eade T, Kneebone A, Hruba G, Moodie T, Hayden A J, Turner S L, Hardcastle N, Siva S, Tai K H, Martin J & Keall P J 2020 *Radiotherapy and Oncology* **0**(0).
- Hunt A, Hansen V N, Oelfke U, Nill S & Hafeez S 2018 *Clinical Oncology (Royal College of Radiologists (Great Britain))* **30**(11), 711–719.
- IEC 2011 Radiotherapy equipment - Coordinates, movements and scales Technical Report IEC 61217:2011 International Electrotechnical Commission Geneva, CH.
- Kamerling C P, Fast M F, Ziegenhein P, Menten M J, Nill S & Oelfke U 2016 *Medical Physics* **43**(11), 6072–6081.
- Kamerling C P, Fast M F, Ziegenhein P, Menten M J, Nill S & Oelfke U 2017 *Medical Physics* **44**(11), 5997–6007.
- Kamino Y, Takayama K, Kokubo M, Narita Y, Hirai E, Kawawada N, Mizowaki T, Nagata Y, Nishidai T & Hiraoka M 2006 *International Journal of Radiation Oncology, Biology, Physics* **66**(1), 271–278.
- Keall P J, Colvill E, O'Brien R, Caillet V, Eade T, Kneebone A, Hruba G, Poulsen P R, Zwan B, Greer P B & Booth J 2018 *International Journal of Radiation Oncology, Biology, Physics* **101**(2), 387–395.
- Keall P J, Colvill E, O'Brien R, Ng J A, Poulsen P R, Eade T, Kneebone A & Booth J T 2014 *Medical Physics* **41**(2), 020702.
- Keall P J, Ng J A, O'Brien R, Colvill E, Huang C Y, Poulsen P R, Fledelius W, Juneja P, Simpson E, Bell L, Alfieri F, Eade T, Kneebone A & Booth J T 2015 *Medical Physics* **42**(1), 354–358.
- Keall P, Nguyen D T, O'Brien R, Booth J, Greer P, Poulsen P, Gebiski V, Kneebone A & Martin J 2017 *BMC Cancer* **17**(1), 180.
- 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, Hruba G, Moodie T, Hayden A, Turner S, Arumugam S, Sidhom M, Hardcastle N, Siva S, Tai K H, Gebiski V & Martin J 2020 *International Journal of Radiation Oncology*Biological*Physics* **107**(3), 530–538.
- Kilby W, Dooley J R, Kuduvalli G, Sayeh S & Maurer C R 2010 *Technology in Cancer Research & Treatment* **9**(5), 433–452.
- Kontaxis C, Bol G H, Lagendijk J J W & Raaymakers B W 2015a *Physics in Medicine and Biology* **60**(19), 7485–7497.
- Kontaxis C, Bol G H, Lagendijk J J W & Raaymakers B W 2015b *Physics in Medicine and Biology*

Dose Optimisation for MLC Tracking

20

- 60**(6), 2493–2509.
- Kontaxis C, Bol G H, Stemkens B, Glitzner M, Prins F M, Kerkmeijer L G W, Lagendijk J J W & Raaymakers B W 2017 *Physics in Medicine & Biology* **62**(18), 7233–7248.
- Low D A, Harms W B, Mutic S & Purdy J A 1998 *Medical Physics* **25**(5), 656–661.
- Machtay M, Bae K, Movsas B, Paulus R, Gore E M, Komaki R, Albain K, Sause W T & Curran W J 2012 *International Journal of Radiation Oncology, Biology, Physics* **82**(1), 425–434.
- Menten M J, Fast M F, Nill S, Kamerling C P, McDonald F & Oelfke U 2016 *Radiotherapy and Oncology* **119**(3), 461–466.
- Moore D, Ruan D & Sawant A 2015 *Medical Physics* **43**(1), 465–474.
- Poulsen P R, Cho B & Keall P J 2009 *Physics in Medicine and Biology* **54**(13), 4019–4035.
- Poulsen P R, Murtaza G, Worm E S, Ravkilde T, O'Brien R, Grau C, Høyer M & Keall P 2020 *Radiotherapy and Oncology* **144**, 93–100.
- Ravkilde T, Keall P J, Grau C, Høyer M & Poulsen P R 2014 *Physics in Medicine and Biology* **59**(23), 7279–7296.
- Ravkilde T, Skouboe S, Hansen R, Worm E & Poulsen P R 2018 *Medical Physics* **45**(8), 3893–3903.
- Ruan D & Keall P 2010 *Physics in Medicine and Biology* **55**(11), 3011–3025.
- Ruan D & Keall P 2011 in '2011 IEEE Power Engineering and Automation Conference' IEEE Wuhan, China pp. 100–103.
- Sawant A, Venkat R, Srivastava V, Carlson D, Povzner S, Cattell H & Keall P 2008 *Medical Physics* **35**(5), 2050–2061.
- Skouboe S, Poulsen P R, Muurholm C G, Worm E, Hansen R, Høyer M & Ravkilde T 2019 *Medical Physics* **46**(11), 4738–4748.
- Tacke M B, Nill S, Krauss A & Oelfke U 2010 *Medical Physics* **37**(2), 753–761.
- Wijesooriya K, Bartee C, Siebers J V, Vedam S S & Keall P J 2005 *Medical Physics* **32**(4), 932–941.
- Wisotzky E, O'Brien R & Keall P J 2015 *Medical Physics* **43**(1), 132–136.