Planning for Intensity Modulated Radiation Therapy

Matthew Davison

1 Introduction

One in two people in the United Kingdom will get cancer at some point in their lives according to forecasting by Ahmad et al. (2015). Currently more than 166,000 people die of cancer every year in the UK according to Cancer Research UK (2017).

Cancer is a disease that is more common in people over the age of 65 (Cancer Research UK, 2017) and with people benefitting from improved healthcare many more people are surpassing this age. It is said that if you live long enough, you will likely get cancer at some point. The above figures show that there is a significant and increasing demand for effective and fast treatment for cancer.

In the United States, more than half of the cancer patients in the country are treated using some form of radiation therapy (Romeijn et al., 2006). In radiation therapy, beams of radiation are passed through the patient killing cancerous and healthy cells. Multiple low powered beams can be angled towards a tumour to kill the cancerous cells but leave surrounding tissue relatively unharmed.

Careful planning is required to ensure minimal damage is done to healthy tissue, whilst managing to irradiate cancers (target structures). Radiation treatment plans often deliver too little radiation to the targets, too much radiation to healthy tissue, or both (Romeijn et al., 2006). This can cause potentially life threatening problems for the patient after treatment.

There are currently no completely effective methods of planning radiation therapy. Several planning techniques have been looked at but nearly all have limitations either in the quality of the plan or the computational power needed to produce the plan. These limitations in the planning process produce bottlenecks at clinics, delays in treatment and increases costs for labour and resources (Preciado-Walters et al., 2004).

2 Intensity Modulated Radiation Therapy

Intensity Modulated Radiation Therapy (IMRT) is a form of radiation therapy that delivers radiation to the patient from several angles by using a linear accelerator and a collimator (Gören and Taşkin, 2015). The linear accelerator provides the radiation, whilst the collimator controls the dose distribution.

Patients receiving IMRT are immobilized with restraints on a couch and the accelerator head can move a full 360 degrees about a point 1 metre from the source. Radiation is often delivered from multiple beam

angles, although usually only 3 to 9 different angles are used to simplify the treatment plan. Similarly, the couch is able to move but in practice it is easier to leave the couch in a fixed location (Romeijn et al., 2006).

An IMRT capable accelerator is equipped with a multi-leaf collimator (MLC) system which is used to generate a desired dose distribution at a given angle (Romeijn et al., 2006). There is only one physical beam emitted from the accelerator but it can be thought of as a grid of beamlets. The multi-leaf collimator uses metal leaves that can restrict the beam opening or aperture to specified subsets of beamlets (Preciado-Walters et al., 2004).

The apertures formed by the collimator system depend on the physical capabilities of that machine (Gören and Taşkin, 2015). Regular MLC is a collimator system that has a single leaf pairing; left and right sets of leaves which are the same size. Rotating MLC is identical to Regular MLC except it can rotate the aperture by 90 degrees. Dual MLC has two pairs of leaves; a left-right pair as in Regular MLC but also an orthogonal top-bottom pair. There may also be further constraints on the positions of the leaves depending on the machine.

IMRT can be delivered either statically or dynamically (Gören and Taşkin, 2015). In static treatment, the leaves of the collimator stay still whilst the radiation beam is on. On the other hand, in dynamic treatment the leaves of the collimator form new shapes whilst the beam is on. In this report we only consider static treatment.

Static IMRT is typically composed of three phases (Gören and Taşkin, 2015). The first phase is to determine the optimal angles to deliver the radiation from and is called beam angle optimization (BAO). The second phase is to find a non-negative integer matrix of values called a fluence map (also called an intensity profile) for each angle. This fluence map describes the required dose of radiation needed from each beamlet to kill target structures (called "treatment dose") whilst making sure that other structures (healthy tissue, easily damaged structures) do not exceed a specific "tolerance dose" (Romeijn et al., 2006). This phase is called fluence map optimization (FMO). The last phase is called leaf sequencing optimization (LSO). In this phase the fluence matrices are decomposed into deliverable apertures.

Effectively optimising treatment plans for IMRT is a difficult problem. Trade-offs need to be made between killing the cancerous cells and sparing the healthy cells but the trade-off is hard to quantify in practice (Romeijn et al., 2006). As each phase uses the solution to the previous phase as an input, this makes optimising the treatment plan even harder. Often, an iterative series of plans is produced with repeated discussion between the planner and the clinician to produce a final plan (Preciado-Walters et al., 2004).

There is a lot of literature outlining different methods that can be used to plan IMRT as this technology has been around for over a decade. The planning process is no easier now than it was then. Smart formulations and methods are needed to provide quality plans in a realistic time, even with recent advances in computing. This report explains and compares a selection of these methods, starting with the FMO phase.

3 Fluence Map Optimization (FMO)

3.1 Summary of FMO

Like temperature, radiation dosage is defined at individual points of the body and determining dosage distribution is an essential part of IMRT treatment (Preciado-Walters et al., 2004). It is within the FMO phase that this dose distribution is calculated and the output of this phase is a fluence map for each angle. The FMO problem typically has multiple local optima.

The shape and location of tissue is modelled by a collection of cubes called "voxels" (Tuncel et al., 2012). The smallest squares a multi-leaf collimator can make with its leaves are often referred to as a "bixels", let J be the set of bixels. Throughout this section assume a set of angles from which radiation will be fired, L, is given and assume also that each bixel can be controlled independently. Denote x_{lj} as the intensity of bixel j from angle $l \in L$ ($x_{lj} \in \mathbb{Z}, x_{lj} \geq 0$) and that the dose delivered to voxel i from bixel j per unit intensity of x_{lj} as a_{lij} ($a_{lij} \geq 0$). The a_{lij} are known as "dose coefficients" (Preciado-Walters et al., 2004).

The distinguishing feature of different FMO models lies in the choice of evaluation criteria that are employed (Romeijn et al., 2004). Most formulations have different criteria and within some multi-criteria formulations there can be treatment goals that conflict. For example, a tumour may be very close to the spinal cord (a structure we wish to not irradiate) therefore an individual voxel may contain both. This means it could be necessary to deliver a high dosage to a voxel to kill the cancer and deliver a low dosage to the same voxel to spare the spine, but these are actions that conflict.

As in Tuncel et al. (2012) the tissue index sets T, S, H, and D are for the primary target, secondary targets, regular healthy tissue, and dose-volume healthy tissue respectively. Note that $D \subseteq H$. Let $I(t)$ denote the set of voxels in the target T. The set of voxels in any other tissue structure k is indicated with $I(k)$ for $k \in S \cup H \cup D$.

3.2 Mixed integer formulation

To tackle the FMO problem both Preciado-Walters et al. (2004) and Tuncel et al. (2012) use a Mixed Integer Program (MIP) and the formulations they present are very similar. The weakest link in tumour control is the "minimum tumour dose" (Preciado-Walters et al., 2004). If this number, denoted as θ, is too low then there is a risk of not completely killing the cancer. Note that θ is a decision variable.

It is also common to introduce a "homogeneity factor", denoted as α, and homogeneity constraints to enforce a "close-to-uniform" dose distribution through the primary target (Tuncel et al., 2012). These constraints are as follows:

$$\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \geq \theta, \qquad \forall i \in I(t), \tag{1}$$

$$\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \leq \frac{\theta}{\alpha}, \qquad \forall i \in I(t). \tag{2}$$

It is important also to ensure secondary targets are free from cancer (secondary targets are often those where the clinician may believe small amounts of cancer cells could have spread to) so we need to provide them with a minimum dosage. On the other hand, it is important that healthy tissue is not irradiated so we need to set an upper bound on the dosage. Let l_k and b_k be the minimum dosage and maximum dosage for $k \in S$ and $k \in H$ respectively. Mathematically these constraints are as follows:

$$\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \geq l_k, \qquad \forall i \in I(k), \quad \forall k \in S, \tag{3}$$

$$\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \leq b_k, \qquad \forall i \in I(k), \quad \forall k \in H. \tag{4}$$

Some tissue is further restricted by dose-volume requirement that forces a fraction f_k of their points not to exceed a given threshold level $d_k < b_k$ (Preciado-Walters et al., 2004). Auxiliary binary variables y_i are introduced to decide if a point is allowed to exceed threshold d_k ($y_i = 1$), or not ($y_i = 0$). To satisfy the dose-volume constraint y_i must be 1 in no more than $(1 - f_k)$ of the $|I(k)|$ points for the tissue $k \in D$. This adds the final constraints for this formulation

$$\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \leq d_k + (b_k - d_k) y_i, \qquad \forall i \in I(k), \quad \forall k \in D, \tag{5}$$

$$\sum_{j \in J} y_i \leq \lfloor (1 - f_k) |I(t)| \rfloor, \qquad \forall k \in D, \tag{6}$$

$$y_i \in \{0, 1\}, \qquad \forall i \in I(k), \quad \forall k \in D. \tag{7}$$

There is currently no objective to go with these constraints (a set of constraints with no objective is called a "feasibility problem") and this is because despite using the same constraints, Preciado-Walters et al. (2004) and Tuncel et al. (2012) use differing objective functions. In the latter, the objective is to maximise the average primary target dose. This is formulated as

$$\max \quad \frac{1}{|I(t)|} \sum_{i \in I(t)} \left(\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \right), \tag{8}$$

where the maximization is taken over $\theta \geq 0$ and all x_{lj} for $l \in L, j \in J$. In the former, the objective is to maximise the minimum tumour dose which is formulated as

$$\max \quad \theta, \tag{9}$$

where the maximization is taken over $\theta \geq 0$ and all x_{lj} for $l \in L, j \in J$.

The objective given by Equation 9 prioritises the minimum tumour dose, which will increase the average dosage for the primary target. The objective given by Equation 8 prioritises having a large average dose. This will have some effect on the minimum tumour dose but only because of the homogeneity constraint given by Equation 2, and it may not have any effect if α is chosen to be too small.

In the context of trying to irradiate a tumour, it could be argued that ensuring that all of the tumour gets at least enough radiation required to kill it completely is more critical than a high average dose. Alternatively,

relying on the homogeneity constraints to "pull up" the minimum tumour dose provides more flexibility and a more uniform dose.

It may look like both these objectives can be pushed to infinity, but this can only be done if for all $k \in H$ we either have $b_k = \infty$ or $a_{lij} = 0$ for all $i \in I(k)$. The former option meaning it would be considered acceptable to irradiate a volume of healthy tissue, and the latter option meaning that tissue k is unaffected by an infinitely intense bixel. The clinician would likely not allow the first option, and the second option doesn't happen in practice as usually all voxels are affected by an individual bixel to some extent ($1 \gg a_{lij} > 0$). It is also impossible for the accelerator to output an infinite intensity.

3.3 Beamlets versus Apertures

Preciado-Walters et al. (2004) initially consider individual bixels but reformulate the problem from individual bixels, to full shapes or apertures. An aperture can be described as a 0-1 matrix with the $(i,j)^{\text{th}}$ entry representing the whether the bixel on the i^{th} row, j^{th} column of the collimator is covered or not covered. A value of 1 corresponds with being uncovered and a value of 0 corresponds with being covered.

In the above MIP formulation, let P be the set of all apertures and x_{lp} be the intensity of aperture p at angle $l \in L$ ($x_{lp} \in \mathbb{Z}, x_{lp} \geq 0$). Denote the dose delivered to voxel i from aperture p per unit intensity of x_{lp} as a_{lip}. Replace J, x_{lj} and a_{lij} with P, x_{lp} and a_{lip} respectively to get the same problem but in terms of apertures instead of beamlets. Note that the problem is just as hard as the original due to the dose-volume constraints (these constraints make the problem NP-Hard as proved by Tuncel et al. (2012)).

3.4 Column generation approach

In Preciado-Walters et al. (2004) they opt to solve this problem using a coupled column generation approach. In this context, a column is an aperture that can contribute to creating the fluence map. The first step is to relax the problem by ignoring the integrality restrictions and generate columns (apertures) for this LP-relaxed version of the problem, also known as the Relaxed MIP (RMIP).

To construct an aperture, they calculate a number associated with an individual bixel called a "beamlet potential". This is roughly speaking the reduced cost of adding the aperture where only that bixel is uncovered. Given all of these beamlet potentials for a given angle, an aperture is constructed such that the bixels with positive potential are uncovered, and the rest are covered.

This aperture is then added to the current RMIP, also known as the Working RMIP (WRMIP) and the whole process is repeated until no more apertures can be created. At this point the WRMIP optimum is an optimal solution to the full RMIP (Preciado-Walters et al., 2004). Despite each aperture being binary, due to the positive continuous x_{lp} values in the optimal solution it gives the effect of a more complex intensity.

There is a recognised issue with this approach in that the columns generated for the RMIP do not necessarily lead to a near optimal MIP solution, which is the actual problem that needs to be solved. This is due to the incredibly large number of binary y_i variables (often 1000s of these).

3.5 Strong valid inequality approach

In Tuncel et al. (2012) they try to solve the original beamlet MIP formulation using strong valid inequalities. In general, these are used to strengthen formulations of difficult optimisation problems by "cutting off" areas of the feasible region.

They consider a simpler formulation which involves a single dose-healthy tissue and a primary target with no homogeneity constraints, but still containing the difficult dose-volume restrictions. The objective still maximises the average dose provided to primary target. Mathematically this formulation is described by Equation 5 through to Equation 8.

This formulation can be relaxed by only considering a subset of the constraints from all available constraints and ignoring integrality, known as the "Linear Program with Constraint Selection" or LPCS. The aim is to derive valid inequalities for the LPCS and add these inequalities to our original MIP formulation. This can be done as the LPCS is a relaxation of the original problem.

From here on disjunctive programming theory is used to produce valid inequalities for the LPCS. Since there are exponentially many valid inequalities they identify two sub-families of inequalities that remove a significant amount of the optimality gap. These are the "Main disjunctive inequality" and "Disjuntive support inequalities".

Finally, they present a separation algorithm. This algorithm solves the linear relaxation of the original problem, identifies disjunctive cuts that can tighten the formulation, and then solves this new relaxation. The algorithm repeats this until it no longer identifies new cuts. It is hoped that at this point the solution is a near-optimal solution for the original MIP.

3.6 Linear programming formulation

In contrast to the mixed integer program used above, Romeijn et al. (2006) use a (mostly) linear program formulation. The linearity of the formulation means that computationally it is significantly faster to find a solution (Romeijn et al., 2006).

The formulation given has both Equation 3 and Equation 4, and doesn't have any hard homogeneity constraints (these are linear in any case so this doesn't impact the discussion of this method). The main difference is the formulation of the dose-volume constraints. Instead of forcing a fraction of voxels to not exceed a threshold dose, they ensure that the average radiation that a fraction of voxels receiving the highest amounts of radiation doesn't exceed an upper limit.

This is done using "conditional value at risk" or CVaR, common place in situations involving risk management (finance, engineering). At a high level, the CVaR lets us avoid dealing with constraints of a probabilistic nature and instead use what are known as tail averages. The details of this in Romeijn et al. (2006) and full the formulation of the linear constraint is given as

$$\zeta_k^u - \frac{1}{(1-f_k)|I(k)|}\left(\sum_{i\in I(k)}\sum_{l\in L}\sum_{j\in J} a_{lij}x_{lj} - \zeta_k^u\right)_+ \leq U_k, \qquad k \in D, \tag{10}$$

where ζ_k^u is a free variable and U_k is some upper limit for $k \in D$ and $(\cdot)_+ = \max\{0, \cdot\}$. Romeijn et al. (2006) also include a lower bound constraint in a similar way, introducing another free variable ζ_k^l and a lower limit L_k. The main takeaway of this formulation is that the dose-volume constraints can be modelled linearly.

Once again, this is currently a feasibility problem. By denoting the set of all structures as V so that $V = T \cup S \cup H \cup D$, and by letting F_v be the penalty function of the dose received by voxels in structure $v \in V$, a general objective to be minimised can be written as

$$\sum_{v \in V} \sum_{i \in I(v)} F_v \left(\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \right). \tag{11}$$

3.7 Projected gradient approach

Romeijn et al. (2006) choose an objective such that it uses piecewise linear convex penalty functions F_v. They also recognise that the hard and soft bound constraints can often conflict, meaning no feasible solution can be found. To counter this they allow some or all of the constraints to be violated and add further penalty functions into the objective.

If all of these penalty functions are piecewise linear and convex, then this problem is a linear program. There are many popular algorithms to solve linear programs, but it is confirmed by Aleman et al. (2010) that the technique used in Romeijn et al. (2006) was a projected gradient algorithm.

These gradient descent algorithms work by finding the gradient (or an approximation to it) of the objective at the current solution and take a "step" of a certain size in the opposite direction. This hopefully converges to a local minima of the objective. Conversely, steps can be made in the same direction as the gradient and arrive at a local maxima.

3.8 Interior point algorithm approach

In Aleman et al. (2010) they opt for an objective that is not linear and remove all constraints from the linear formulation except the non-negativity constraints for the decision variables. The structure of the objective is exactly the same as in Equation 11. For each voxel i in structure v the penalty function is given as

$$F_v \left(\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \right) = \frac{1}{|I(v)|} \left[\underline{w}_i \left(T_i - \sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} \right)_+^2 + \overline{w}_i \left(\sum_{l \in L} \sum_{j \in J} a_{lij} x_{lj} - T_i \right)_+^2 \right], \tag{12}$$

where T_i is the desired dose delivered to the voxel i, and $\underline{w}_i$ and $\overline{w}_i$ are the weights penalizing underdose and overdose to voxel i respectively. Despite the flexibility in choosing the weights, Aleman et al. (2010) assign weights so that voxels in the same structure have the same weight. The important property of this model is that it is convex.

A primal-dual interior point method is presented. The main benefit of using an interior point method is that at each point we can calculate the difference between the primal objective and the dual objective, called the duality gap (Aleman et al., 2010). This provides an upper bound on the distance from the optimal solution for the FMO problem. At optimality the duality gap is zero, although optimality is not necessary

for an acceptable treatment plan. Therefore, the interior point algorithm can be stopped when the gap is sufficiently small.

Primal-dual interior point methods move through the interior of the feasible region along a path defined by perturbing the Karush-Kuhn-Tucker (KKT) conditions (Aleman et al., 2010). By starting at an initial point in the feasible region, the current solution is moved along this path until we have a small enough duality gap.

The most expensive part of this process is calculating the Hessian in every iteration. The Hessian is a square matrix with dimension equal to the number of bixels and each entry is some second derivative of the objective (used to calculate "direction" of path). Aleman et al. (2010) present two methods of approximating this matrix to speed up the process.

The algorithm is made faster by starting at a solution close to the optimal solution, a process called a "warm start". On the other hand, requiring a solution of any quality so this can be done could be considered a drawback.

Aleman et al. (2010) take pride in providing a way to obtain a truly optimal solution. Despite this, they do suggest that if beamlets that have little contribution to the quality of the plan could be identified then these could be removed from the model, reducing the dimension of the problem. There is then a trade-off between computation time and true optimality.

3.9 Comparison of FMO methods

This section looked at four different methods that were chosen for both their differing approaches. In the first approach, Preciado-Walters et al. (2004), an MIP was considered and the problem was transformed to consider apertures instead of beamlets. Different apertures are constructed then added based on how much they are expected to benefit the solution. This is computationally slow in difficult FMO instances and it is not provable that it produces a quality plan in a reasonable time (Tuncel et al., 2012).

The second method, Tuncel et al. (2012), uses strong valid inequalities to cut off parts of the feasible region for the MIP. This method is better computationally than the last one, but still could be improved. One improvement could be to somehow strengthen the inequalities not used in the LPCS but are in the original MIP.

The third method, Romeijn et al. (2006), switches to a linear formulation and it achieves this linear formulation using a technique borrowed from finance, CVaR. This linear formulation is very fast computationally when using projected gradient procedures. Unfortunately, using a projected gradient procedure means there is not a guarantee we are at an optimal solution or if the current solution is close to the optimal solution (Aleman et al., 2010).

The fourth method, Aleman et al. (2010), tries to resolve this lack of guarantee by employing a primal-dual interior point method. They also use a quadratic function which represents the idea in clinical decision making that increasing deviations from the prescription dose is increasingly bad. The formulation they use

is quite minimal and doesn't incorporate all of the constraints that the other three methods do (dose-volume constraints, for example).

4 Beam Angle Optimization (BAO)

4.1 Summary of BAO

In the FMO section it was assumed throughout the set of optimal angles to fire beams of radiation from were given. The problem therefore was how much power we need delivered from each beamlet. Beam angle optimization is the process of determining the optimal set of beam angles.

Beam angle optimization is a difficult problem to solve because of the huge space of possible beam configurations (Wang et al., 2004). It is also a difficult problem since it is non-convex and has many minima (Craft, 2007). When considering using an angle, at some point dose calculations are needed at that particular angle to determine its role in the plan (Craft, 2007). Only once this is known it can be decided whether to keep that angle in the plan or not.

It is common for a planner to select from a set of fixed angles based on experience or by trial-and-error approach (Wang et al., 2004). This is an adequate and well-established approach in the area. Typically only 3 to 9 angles are actually used (Romeijn et al., 2006), mainly to try and reduce treatment time and increase patient comfort.

Despite the difficulties, BAO has some benefits. In particular, the optimisation may find optimal beam angles that are counter-intuitive and maybe dismissed by a clinician developing a plan (Wang et al., 2004).

4.2 Exhaustive search approach

Wang et al. (2004) propose what they call an "exhaustive search". This means that they will evaluate IMRT plans for all combinations of m angles out of a pool of n equispaced beam angles. By evaluate they mean they will find the objective value given by an FMO method and this value, also known as a "score", will be compared with other plans. Depending on the values of the integers m and n, this can be a very time consuming process.

It is expected that many near-optimal plans will not vary much in score. Therefore groups of beam angles below (or above depending on the FMO method) a threshold objective value are considered equally good.

The key to this method is speed. When considering a group of angles, the FMO method used uses a variant of Newton's optimisation algorithm to quickly find an approximate fluence map for each angle. The fluence maps found at this stage will not be delivered to the patient as they are not accurate enough, but the objective value from the optimisation can be compared with other groups of angles. Once the best group of angles is found then the FMO can be solved using a time-consuming but accurate method. It is then these fluence maps and angles that are delivered to the patient.

There are immediate drawbacks to this method in that it requires an initial decision on how many angles will be used and the degeneracy of having plans with differing angles but similar scores needs to be dealt with. On the otherhand, Wang et al. (2004) suggest that plans with a small number of angles but better optimised are better than plans with lots of angles that are not fully optimised, and that the best plans avoid parallel-opposing pairs (reducing the size of search space).

4.3 Gradient search approach

In Craft (2007), the proposed approach is to use a gradient search approach. Let $L = \{L_1, \ldots, L_n\}$ denote a set of n fixed angles and $f(L)$ denote the objective cost of the optimised IMRT plan using the angles in L.

Using notation from earlier, recall that a_{lij} is the dose delivered to voxel i from bixel j per unit intensity of x_{lj}. Given an $l \in L$, let A_l be the matrix such that $(i, j)^{\text{th}}$ entry is equal to a_{lij}. Let A be the matrix that is a concatenation of $A_{L_1}, \ldots, A_{L_n}$

The method uses duality theory to get a change in f when A is perturbed. The slopes of the matrix A with respect to the elements of L along with the chain rule can be used to determine how f changes in L. By using the optimal solution from the FMO method (and their corresponding dual values) with the slope of A the gradient of the function $f(L)$ can be found.

This gradient is then used for gradient-based optimisation of the $f(L)$ function. Note that this will only arrive at a local solution and it is possible to have many local minima in $f(L)$. Craft (2007) suggests using some heuristic to sample the beam angle domain (global search), and perform the local search on one or several candidate angle sets.

4.4 Hybrid approach

Bertsimas et al. (2013) use a hybrid approach. It combines a procedure called simulated annealing with a gradient based procedure like in the previous approach. Essentially, this is an implementation of the suggestion given by Craft (2007) to perform a global search, then perform a local search.

Simulated annealing is a procedure that has been used for a long time in the BAO process. Given a set of angles at a local minima, L, the simulated annealing procedure generates angles nearby with the hope it escapes the local minima. Bertsimas et al. (2013) do this by generating a set of random numbers (using a Gaussian distribution) that is the same size as L and add the first random number to L_1, second to L_2 and so on.

The FMO problem is then solved for the new angle set, and if there is improvement keep that angle set and repeat. If there isn't improvement then there is a probability of keeping the angles anyway, otherwise keep the old angle set.

The method performs a fixed number of simulated annealing steps as described above, to escape local minima, and then performs a number of gradient descent steps to refine the solution. The whole process

swaps back and forth until the algorithm reaches a certain number of iterations or some time limit has been reached.

Attempts were made to introduce binary variables to correspond with a particular angle being included in the solution, and to add constraints that limit the number of angles, but the solutions found were worse than simply using the hybrid method.

4.5 Comparison of BAO methods

This section looked at three methods for solving the BAO problem. These methods were chosen to illustrate how there is a trade-off between looking at a large portion of the solution space completely to guarantee a good solution, and "scanning" the solution space to quickly find a solution that is acceptable.

The first method in Wang et al. (2004) tries to perform an exhaustive search. It could be claimed that the pool of angles used in their tests was not large enough for the method to be effectively "exhaustive", but the argument they make is that the rest of the planning process can correct for the sub-optimal beams.

The second method, Craft (2007), is a local search method that uses duality and gradient descent to effectively find a local optimal. The obvious drawback of this is that it only finds a local optimal. This may come up with suitable angles for a treatment plan, but the lack of exploration means it is hard to know if a better set of angles were available.

The final method, Bertsimas et al. (2013), improves on this method by introducing a simulated annealing step to perform a global search, and then do a local search. It is demonstrated that this hybrid approach is better than simulated annealing on its own. As simulated annealing is still a heuristic algorithm, a better hybrid algorithm may wish to incorporate the "exhaustive search" seen in Wang et al. (2004) as the global search step (although this is computationally more expensive).

5 Leaf Sequencing Optimization (LSO)

5.1 Summary of LSO

The multi-leaf collimator (MLC) is an essential part of IMRT. To create apertures in the past, a clinician would commission bespoke apertures to be made for each patient out of large pieces of metal (Hogstrom et al., 2004). This process is time consuming and expensive, although thanks to the 3D printing of tungsten apertures this is a slightly quicker (Skinner et al., 2019). In IMRT, the leaves can produce a huge range of apertures in a matter of seconds, the exact number depending on the collimator system used (Gören and Taşkin, 2015).

The LSO phase is where the fluence maps, given by FMO, are decomposed into shapes that can be delivered by the MLC system. As discussed earlier there are multiple types of multi-leaf collimators with differing physical capabilities and this manifests mathematically as constraints in the LSO problem.

There are often two main objectives of LSO. These are to minimise what is known as "beam-on time" (BOT), which is the total time that the linear accelerator is turned on for, and to minimise the number of apertures used. These have the overall effect of reducing the time of the treatment and therefore increase patient comfort. Small amounts of radiation can leak between leaves of the collimator and the best way of controlling this is to reduce the BOT and the number of apertures used, improving the quality of the treatment.

Minimising beam-on time is not the same as minimising delivery time, which is the problem of minimising the beam-on time and the total time spent switching between apertures. This is a harder problem and is proven to be NP-hard (Ahuja and Hamacher, 2005).

5.2 Unconstrained minimum beam-on time problem

To discuss the methods used to solve LSO the unconstrained minimum beam-on time (MBT) problem is introduced. What is meant by “unconstrained” is that there is no focus on a particular collimator system right now and as default assume that bixels can be covered or uncovered independently, also known as a free-form collimator (Gören and Taşkin, 2015).

Let the set of feasible apertures be denoted as C (remember an aperture can be described as a 0-1 matrix) and the size of this set be denoted as κ. For each of the shape matrices, C^k where $k = 1,\dots,\kappa$, if the bixel in the i^{th} row and j^{th} column of the collimator is uncovered by shape C^k, then $C^k_{ij} = 1$, and zero otherwise. Let x_k be the beam-on time for shape k which is a non-negative decision variable in our problem. Therefore given an $m \times n$ integer fluence map F, the unconstrained MBT problem is given as

$$\begin{aligned} \min \quad & z = \sum_{k=1}^{\kappa} x_k, \\ \text{s.t.} \quad & \sum_{k=1}^{\kappa} C^k_{ij} x_k = F_{ij} \quad \forall\, i = 1,\dots,m, \quad \forall\, j = 1,\dots,n. \end{aligned} \tag{13}$$

It cannot be understated that the complexity to this problem depends heavily on the number of feasible apertures. This number, κ, can be huge (Gören and Taşkin, 2015). There are at least $m \times n$ possible shapes, and κ can be as large as $2^{m\times n}$. This number comes from considering the number of combinations of covering or uncovering the bixels in the free-form collimator case.

5.3 Mixed integer formulation

In Langer et al. (2002) they avoid the need for a set of apertures but instead require binary variables and linking constraints. The collimator considered here is a collimator with a row of left leaves and a row of right leaves. Let $T = \max_{i,j}(F_{ij})$. For $t = 1,\dots,T$, p^t_{ij} and l^t_{ij} take the value 1 if the bixel in row i, column j is covered by the the right leaf or left leaf respectively. d^t_{ij} takes the value 1 if the bixel in row i, column j is uncovered in time t and 0 otherwise. The relationship between these binary variables is given by

$$p_{ij}^t + l_{ij}^t = 1 - d_{ij}^t, \quad \forall i = 1,\dots,m, \quad \forall j = 1,\dots,n, \quad \forall t = 1,\dots,T. \tag{14}$$

This relationship means that a bixel may only be covered by one or none of the leaves. A further relationship between the leaf variables is referred to as the "consecutive ones" property (Ahuja and Hamacher, 2005) and says that if a bixel is covered by a left leaf, then all bixels to the left of that are also covered. Similarly for bixels covered by a right leaf. Mathematically, this is given as

$$p_{ij}^t \leq p_{i(j+1)}^t, \quad \forall i = 1,\dots,m, \quad \forall j = 1,\dots,n, \quad \forall t = 1,\dots,T, \tag{15}$$

$$l_{ij}^t \geq l_{i(j+1)}^t, \quad \forall i = 1,\dots,m, \quad \forall j = 1,\dots,n, \quad \forall t = 1,\dots,T. \tag{16}$$

The sum of the d_{ij}^t over t should be equal to the intensity of the $(i,j)^{\text{th}}$ entry of F, that is,

$$\sum_{t=1}^{T} d_{ij}^t = F_{ij}, \quad \forall i = 1,\dots,m, \quad \forall j = 1,\dots,n. \tag{17}$$

To finish off the formulation so that wit resembles Equation 13 Langer et al. (2002) introduce a variable z^t that is 1 if any beamlet in time t is uncovered and 0 otherwise. This is modelled as

$$\sum_{i=1}^{m}\sum_{j=1}^{n} d_{ij}^t \leq z^t mn, \quad \forall t = 1,\dots,T, \tag{18}$$

and the objective defined as

$$\min \quad z = \sum_{t=1}^{T} z^t. \tag{19}$$

where z is the total beam-on time as in the unconstrained MBT problem. This is a clever use of binary variables to model unconstrained MBT with two common collimator constraints. The downsides of this formulation is that many of the variables will be zero, so there is potential for wasting computational power by considering them, and it is not clear how many unique shapes are being used.

Langer et al. (2002) do in fact provide a way of working out the minimum number of shapes needed to deliver the minimum beam-on time. It involves solving the problem above to get a minimum beam-on time, denoted as $\tilde{Z}$, and then turning Equation 19 into a constraint. That is,

$$\sum_{t=1}^{T} z^t \leq \tilde{Z}, \quad \forall i = 1,\dots,m, \quad \forall j = 1,\dots,n. \tag{20}$$

More binary variables and linking constraints are used to identify if any of the leaves change position between times t and $t+1$ for $t = 1,\dots,T-1$. A decision variable g^t takes the value 1 if any number of leaf changes occur at time t (modelled using a constraint similar to Equation 18) and the objective of this modified model is

$$\min \quad \sum_{t=1}^{T} g^t. \tag{21}$$

This seemingly minor adjustment is mathematically sound, but testing of this formulation in Python 3 using the Gurobi Optimizer showed that even small fluence maps with as little as 5 rows and 8 columns can take hours to solve on a moderately powerful computer. This is likely due to the large number of binary variables. For comparison, finding only the minimum beam-on time took less than 30 seconds for fluence maps as big as 22 rows and 23 columns.

5.4 Column generation approach

Gören and Taşkin (2015) consider the problem described in Equation 13. They suggest that a solution to the problem will not use most of the available shapes our collimator can produce so instead of considering all of C, try to solve a restricted problem with a subset of these shapes, $\tilde{C}$.

A feasible subset of shapes is needed to begin the column generation approach. The subset that is taken is simply all the shape matrices that have exactly one entry with value 1. The minimum beam-on time for this set of apertures is the sum of the entries of the fluence map.

By using duality theory and the optimal solution using the restricted set $\tilde{C}$, dual variables associated with the constraint in Equation 13 can be found. These reduced costs can be used to find a shape that shows promise of improving our current solution. Denote λ_{ij} as the dual variable corresponding with the constraint enforcing the correct intensity of the bixel in row i, column j.

These variables are used to define a subproblem that aims to find a shape matrix that has the minimum reduced cost. It is as follows:

$$\begin{aligned} \min \quad & 1 - \sum_{i=1}^{m} \sum_{j=1}^{n} \lambda_{ij} C_{ij}^{k}, \\ \text{s.t.} \quad & C^k \text{ is a feasible shape matrix} \end{aligned} \tag{22}$$

If the optimal value of the subproblem is non-negative, then the current solution to the restricted problem is optimal for the original formulation in Equation 13. If the optimal value of the subproblem is negative, then add the shape matrix that solves it to optimality to $\tilde{C}$ and solve the restricted problem again. This gives a new set of dual variables and we solve the subproblem again. This process repeats until optimality or until the plan is of satisfactory quality.

5.5 Network flow approach

Like in the column generation approach, Ahuja and Hamacher (2005) start off with the problem as stated in Equation 13. The important observation made here is that in the unconstrained problem the rows are independent of each other and the problem can be split into m subproblems. The subproblem is to find a decomposition of an intensity row into a postive linear combination of 0-1 rows with the consecutive ones property.

It is shown by Ahuja and Hamacher (2005) that these subproblems can be reformulated as minimum cost flow problems. The directed networks used to solve these problems are acyclic and complete, and every arc has cost 1 and are uncapacitated. These properties allow efficient solving of these subproblems, and in fact mean that the whole problem can be solved in $O(nm)$ operations.

Since this approach only works for the unconstrained MBT problem it appears this method may not have much use, but this approach could be used to find a starting point for a problem with more constraints. Kalinowski (2010) have explored a similar network flow technique that finds only an approximate decomposition of the intensity matrix, but manages to consider some collimator constraints.

5.6 Comparison of LSO methods

This section looked at three methods for solving variants of the LSO problem. These methods were chosen to illustrate how important the formulation is when trying to solve this problem. It is surprising how some formulations, whilst mathematically sound, can be problematic in practice.

The first method, Langer et al. (2002), was an MIP approach that looks promising on the surface. It uses a clever combination of binary variables and linking constraints to formulate the problem clearly. The problems appear when trying to implement this method. The large number of binary variables mean that the problem has a high dimension and this slows down the commercial MIP solver an unreasonable amount.

Gören and Taşkin (2015) used a column generation approach, a technique already seen used in the context of FMO. The idea was to start with a set of apertures and slowly introduce better apertures based on reduced costs. This solves some of the dimension problems, but a lot apertures may still have to be considered until a one worth introducing into the solution is found.

The third approach, Ahuja and Hamacher (2005), exploited the structure of the unconstrained MBT problem to find an efficient way of solving it. This approach is limited since it relied on the independence of the rows and in practice the leaves of an MLC system do not operate independently.

All of the approaches in this section couldn't provide an effective way of reducing the beam-on time and the number of shapes used. The first approach was the only one to give a formulation of this problem.

6 Conclusion

The conclusion to this report will discuss the planning process as a whole and point out the most interesting approaches used within it. Hopefully this report has made it clear that planning for IMRT treatment is not an easy task. Thankfully, due to the experience of clinicians and practitioners, truly optimal solutions are not necessary for a treatment plan to be able to cure a patient of cancer. However it cannot be ignored that mathematical optimisation is a valuable tool used to assist the clinician, providing great insight into improving treatment quality and delivery time.

One interesting approach is the use of an MIP formulation, like the ones used in FMO and LSO. Using a mixed integer formulation is an effective way of describing a problem mathematically. It is usually quite easy to read a mixed integer formulation and understand what it is trying to achieve. The drawback is that the computational power required to solve them to optimality is impractical, unless algorithms like column generation are applied.

Other interesting approaches are the heuristic global searches combined with local gradient searches. These approaches focus on speed and can find solutions quickly. These procedures are best used within subproblems and subroutines of algorithms. It is known that a sub-optimal treatment plan can still be an acceptable plan so despite not being optimal these procedures can be used provide a feasible solution to kick-start a more involved optimisation method. Due to a spike in popularity for genetic algorithms and

machine learning these approaches could be updated.

7 Further research

Throughout this report the focus has been on the three-step approach to planning. That is BAO, FMO and LSO. There is a large body of literature about each of these problems but something that has caught the attention of researchers recently is a so-called "combination approach". The aim is to integrate two or more steps of the planning process in the hope that doing so:

1. Streamlines the planning process, meaning plans can be produced faster.
2. Produces plans that are of higher quality, meaning patients are better treated.

7.1 Motivation for combined approaches

Dividing the IMRT treatment plan into phases and dealing with them separately may cause loss in treatment quality (Gören and Taşkin, 2015). Discretization steps to make problems tractable, such as using a set of evenly spaced angles in BAO, can also reduce quality of plan. The main problem with separating the planning into three phases is that there is the potential to lose information between steps.

Not only is there the potential to lose information but the quality of information we have at the start of each phase is only as good as the output of the previous phase. For example, a poor choice of angles in BAO could lead to sub-optimal fluence maps being output from the FMO phase.

7.2 Fluence Angle Optimization (FAO)

The output of the BAO phase is a set of beam angles to input into the FMO process. As discussed in the BAO section, the approaches earlier in this report require solving the FMO problem to some extent to determine if the chosen set of beam angles are any good.

In particular, the methods looked at in the report chose a group of angles, then solved the FMO problem and repeated this process with many groups of angles. The group of angles with the best FMO objective value are the angles that are taken forward for the treatment. This process can take several hours if it needs to be repeated a lot before a clinically acceptable plan is found (Cabrera-Guerrero et al., 2016).

One solution is to optimise the beam angles and the fluence maps simultaneously in one larger problem called Fluence Angle Optimization or FAO. Fiege et al. (2011) tackle this problem by formulating it as a multi-objective problem (a type of formulation popular in this field but not explored in this report) and designing a genetic algorithm called "Ferret" to solve it. Genetic algorithms are well suited for parallel computing so solutions can be found quickly using this method.

Another solution is to produce a set of promising angle groups to use in the traditional process. So rather than the clinician having to come up with new groups of angles, the clinician is provided a collection

of approximately Pareto efficient groups of angles. In the multi-objective setting, efficient means that no objective can be better without making another worse. Cabrera-Guerrero et al. (2016) give a two-phase approach to perform this task.

7.3 Direct Aperture Optimization (DAO)

The output of the FMO phase is a collection of fluence maps and there is now the task of finding how this should be delivered. Remember that in Preciado-Walters et al. (2004) the FMO problem was solved by introducing apertures slowly using column generation. Why not add the collimator constraints into this construction process, so the apertures are deliverable, and retain the information of what apertures are chosen? This may not be an optimal decomposition, but at least provides us with a feasible solution.

This approach is seen Salari and Unkelbach (2013) and directly optimises shape and intensity of deliverable apertures. This combination of FMO and LSO is called Direct Aperture Optimization or DAO. Salari and Unkelbach (2013) also use column generation (which we looked at in FMO and LSO separately) to solve a multi-objective program describing the DAO problem.

Zhang et al. (2019) build on this work by using a technique called "region growth" to improve the efficiency of the column generation. This technique merges consecutive bixels that are in the same row and have negative beamlet potentials into a single larger block. The potential of the block is the sum of the beamlet potentials. The column generation proceeds as usual except now there are less beamlet potentials to consider.

7.4 Combining all three processes

It is shown in literature demonstrating combined approaches that this combination approach yields higher quality plans and in some cases plans can be produced faster. These combinations give rise to two possible two-step approaches. One approach is using FAO then LSO, and the other is using BAO then DAO. The natural question to ask is can BAO, FMO, and LSO be combined into one single planning process?

On the surface it certainly seems possible given that BAO can be combined with FMO and FMO can be combined with LSO. The main drawback of doing this is that the final single-step problem will inherit all of the challenges from each phase so that even if a mathematical description of the problem is produced, it will be hard to find a solution in any reasonable time.

Despite this, producing a mathematical formulation (potentially in the form of a large-scale MIP, for example) is the starting point. Then the structure of the problem can be analysed and potentially exploited so that finding a solution in good time is possible.

References

Ahmad, A., Ormiston-Smith, N., and Sasieni, P. (2015). Trends in the lifetime risk of developing cancer in Great Britain: comparison of risk for those born from 1930 to 1960. *British Journal of Cancer*, 112.

Ahuja, R. and Hamacher, H. (2005). A network flow algorithm to minimize beam-on time for unconstrained multileaf collimator problems in cancer radiation therapy. *Networks*, 45:36–41.

Aleman, D. M., Glaser, D., Romeijn, H. E., and Dempsey, J. F. (2010). Interior point algorithms: guaranteed optimality for fluence map optimization in IMRT. *Physics in Medicine and Biology*, 55(18):5467–5482.

Bertsimas, D., Cacchiani, V., Craft, D., and Nohadani, O. (2013). A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. *Computers & Operations Research*, 40:2187–2197.

Cabrera-Guerrero, G., Ehrgott, M., Mason, A., and Raith, A. (2016). A matheuristic approach to solve the multiobjective beam angle optimization problem in intensity-modulated radiation therapy. *International Transactions in Operational Research*.

Cancer Research UK (2017). Cancer incidence statistics. `https://www.cancerresearchuk.org/health-professional/cancer-statistics/incidence`. Accessed: 2021-04-18.

Craft, D. (2007). Local beam angle optimization with linear programming and gradient search. *Physics in Medicine and Biology*, 52(7):N127–N135.

Fiege, J., McCurdy, B., Potrebko, P., Champion, H., and Cull, A. (2011). Pareto: A novel evolutionary optimization approach to multiobjective IMRT planning. *Medical physics*, 38:5217–29.

Gören, M. and Taşkin, Z. (2015). A column generation approach for evaluating delivery efficiencies of collimator technologies in imrt treatment planning. *Physics in medicine and biology*, 60:1989–2004.

Hogstrom, K., Boyd, R., Antolak, J., Svatos, M., Faddegon, B., and Rosenman, J. (2004). Dosimetry of a prototype retractable eMLC for fixed-beam electron therapy. *Medical physics*, 31:443–62.

Kalinowski, T. (2010). A minimum cost flow formulation for approximated MLC segmentation. *Networks*, 57:135–140.

Langer, M., Thai, V., and Papiez, L. (2002). Improved leaf sequencing reduces segments or monitor units needed to deliver IMRT using multileaf collimators. *Medical physics*, 28:2450–8.

Preciado-Walters, F., Rardin, R., Langer, M., and Thai, V. (2004). A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming*, 101(2):319–338.

Romeijn, H. E., Ahuja, R., Dempsey, J. F., and Kumar, A. (2006). A new linear programming approach to radiation therapy treatment planning problems. *Operations Research*, 54:201–216.

Romeijn, H. E., Dempsey, J. F., and Li, J. G. (2004). A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology*, 49(10):1991–2013.

Salari, E. and Unkelbach, J. (2013). A column-generation-based method for multi-criteria direct aperture optimization. *Physics in Medicine and Biology*, 58(3):621–639.

Skinner, L., Fahimian, B., and Yu, A. (2019). Tungsten filled 3D printed field shaping devices for electron beam radiation therapy. *PLOS ONE*, 14:e0217757.

Tuncel, A. T., Preciado, F., Rardin, R. L., Langer, M., and Richard, J. P. (2012). Strong valid inequalities for fluence map optimization problem under dose-volume restrictions. *Annals of Operations Research*, 196(1):819–840.

Wang, X., Zhang, X., Dong, L., Liu, H., Wu, Q., and Mohan, R. (2004). Development of methods for beam angle optimization for IMRT using an accelerated exhaustive search strategy. *International journal of radiation oncology, biology, physics*, 60:1325–37.

Zhang, L., Gui, Z., Yang, J., and Zhang, P. (2019). A column generation approach based on region growth. *IEEE Access*, 7:31123–31139.