Direct Aperture Optimization for Proton Therapy Using a Multi Leaf Collimator

J. Unkelbach

Francis H Burr Proton Therapy Center
Report Number 2005-01

October 2005

Direct aperture optimization for proton therapy using a Multi Leaf Collimator

1. Introduction

The work considers the Direct Aperture Optimization (DAO) problem for proton therapy. In conventional photon IMRT using a Multi Leaf Collimator and the step and shoot method, the DAO problem refers to the direct optimization of MLC segment shapes. This is in contrast to the widely used two step approach of beamlet based optimization and subsequent sequencing.
Intensity modulated proton therapy (IMPT) is currently realized using a scanning beam. However, we assume that also for proton therapy a Multi Leaf Collimator can be applied in order to deliver intensity modulated proton beams.
In photon therapy, the DAO problem has two groups of variables:

1) The weights of the apertures
2) The leaf positions that determine the shape of the apertures

In proton therapy there will in general be two more variables per aperture:

1) The range of the proton beam
2) The modulation, i.e. the width of the SOBP

In order to deal with additional degrees of freedom, different approaches can be considered:

1) *Using a SOBP that covers the entire tumor*
   For example, the range and modulation of each beam direction is predetermined in such a way that the SOBP covers the entire tumor. This strategy has the advantadge that no patient specific hardware would be required. In addition, the optimization problem would formally be identical to the DAO problem in photon therapy because both additional parameters are fixed. The only difference would be a change in the dose contribution matrix of beamlets, i.e. the depth dose profile would be flat in the vicinity of the tumor instead of exponentially decaying. The downside would be that one major advantage of proton beams is unused, i.e. the potential of the proton beam to stop at the distal edge of the tumor.
2) *Using a SOBP plus a range compensator*
   The disadvantage of the previous method could be reduced by applying a compensator for each beam which adjusts the end of the SOBP to the distal edge of the tumor. This would reduce the dose burden of healthy tissue behind the tumor. However, the SOBP would produce high dose regions in front of the CTV.

In the case that multiple beams are used, these high dose regions would partially be avoided by the modulation of beam intensities like this is the case for photon IMRT. However, the approach requires a patient specific compensator for each beam direction.

3) *Using pristine peaks*
Probably, the full potential of proton therapy can only be exploit by allowing for different ranges and modulations for different apertures. Here, the range is considered as an additional parameter whereas the modulation as fixed. This means, pristine peaks or SOBPs with small modulation are considered. In principle an aperture corresponding to a proton beam with large modulation could be represented in terms of multiple apertures corresponding to pristine peaks and different ranges. This approach does not necessarily require patient specific devises. However, the range is an additional degree of freedom which complicates the DAO problem.

Treatment planning for method 1 and 2 could also be performed as beamlet based optimization and does not necessarily require DAO. In principle also method 3 could be subject to beamlet based optimization. For each beam direction there would be one beamlet per lateral position and per allowed range. However, sequencing of the resulting fluence maps would result in a large number of segments (at least if sequencing is performed for each range seperately). Therefore the need for DAO optimization is more indicated for this method.

2. The DAO problem

All further chapters deal with method 3 introduced above. It is assumed that pristine peaks or SOBPs with small modulation are used. Different apertures may have different range. The optimization problem is subject to three different types of variables which are treated in different wayes.

1) *The weights of the apertures*
The aperture weights are continuous variables. An objcetive function is chosen which is convex with respect the the aperture weights. Hence, for a given configuration of aperture shapes and ranges, one can optimize the aperture weights using a gradient method. Given a set of aperture shapes and ranges only the corresponding optimal weights are considered.

2) *The leaf positions*
Since a dose contribution matrix concept is applied for dose calculation, the leaf positions are varied in discrete steps corresponding to one beamlet. The optimization of the leaf positions is the major focus of this work and described in subsection 2.1.

3) *The range*
The range of the proton beam can also vary in discrete steps only due to the

underlying dose contribution matrix approach. A change of the range of an aperture will typically lead to a large change of the dose distribution. It is assumed that it is therefore difficult to treat the aperture range as a variable in the optimization process. Two approaches to deal with the range parameter are investigated:
First, each apperture is assigned a fixed range and the number of apertures is kept fixed. The disadvantage of the approach is, that for certain depths one may not need apertures to achive a high quality dose distribution, whereas for another depth it might be favorable to have a larger number of apertures. However, all chapters in the result part except section 7 deal with this simplified approach.
In the second approach range for each aperture is again fixed, however, the number of apertures is reduced during the optimization process. The optimization starts with several apertures for each range but during the optmization, apertures are deleted so that only a smaller number of apertures remains. In principle, the optimization algorithm has the freedom to choose these aperture ranges which are most useful. Section 7 sketches this approach, however, promising results were not obtained so far.

2.1. Optimization of the leaf positions

The optimization of the leaf positions is the major focus of the work. Essentially, a simulated annealing algorithm is used to assess the optimization problem.

2.1.1. Number of leaf configurations The aperture shapes can be parameterized by the leaf positions of all apertures. The total number of all possible leaf configurations represents the search space of the aperture optimization problem. The number of possible leaf configurations is huge. Let A be the number of apertures, let M_μ be the number of leaf pairs in aperture μ and let $N_{\mu i}$ be the number of allowed beamlets for leaf pair i in aperture μ. Without constraints on interdigitation or overtravel, the total number S of possible points in search space is given by

$$S = \prod_{\mu=1}^{A} \left[\prod_{i=1}^{M_\mu} \frac{(N_{\mu i}+1)(N_{\mu i}+2)}{2} \right] \tag{1}$$

For the 2D example considered below, the parameters $A = 36$, $M_\mu = 1$ and $N_{\mu i} = 9$ yield $S \approx 4.5 \cdot 10^{62}$. Hence, even for a very idealized example the number of possible states is enormous which makes exhaustive search impossible.

2.1.2. Physical equivalence of different states Many of these states are physically equivalent in the sence that they lead to the same effective fluence map of the beamlets.

1) All states that correspond to closed leaf pairs are physically equivalent since they deliver no dose, no matter where the leafs are positioned.

2) If multiple apertures exist for a given beam and range, permutations of these apertures are physically equivalent but correspond to different positions in search space.
3) Finally, solutions are physically equivalent when two states differ in the shape of one apperture, but the optimal weight of this aperture is zero for both states.

2.1.3. General optimization steps It is the aim to find the leaf configuration which minimizes the objective function, or at least some state with a near optimal objective value that corresponds to an acceptable dose distribution. The search strategy investigated here consists of three steps:

1) Choose an initial position in search space. This initial position is randomly generated, i.e. all leaf pairs of all apertures are considered independently and for each leaf pair one of the possible combinations of left and right leaf position is chosen with equal probability.
2) A simulated annealing like algorithm is applied to transform the initial position into a starting position for a local optimization algorithm. The simulated annealing method generates a random walk though search space where steps that improve the objective value are more likely than steps that lead to a worse state. During the simulated annealing process, the best solution obtained so far is saved.
3) Finally, the best state that was visited during the simulated annealing process is transformed to a final state in search space using an algorithm that only allows steps that improve the objective value. This method is refered to the "local optimization method" since it will typically applies only minor changes to the leaf positions instead of searching a significant portion of the search space. The final state is refered to as a solution.

2.1.4. The local optimization method When the temperature approaches zero, the simulated annealing method accepts only leaf moves that improve the objective value. Leaf moves are proposed on a random basis. Realistically, the simulated annealing method, therefore, becomes very inefficient for low temperatures since most of the proposed leaf moves are rejected. In practice, the simulated annealing step will be terminated at finite temperature after a specified number of iterations. Therefore, it is usually possible to improve the objective value by moveing individual leafs, i.e. the intermediate solution can be further improved locally. The local optimization method applied here proceeds according to the following steps:

1) Determine the change in the objective value for all possible moves of a single leaf by one bixel.
2) Determine the leaf moves that improve the objective value by more than a predefined threshold. The standard parameter value for this threshold is 0.001.

3) If such a move does not exist, the optimization is terminated. Otherwise, choose randomly one of these leaf moves, execute this move, adapt the aperture weights and go back to step 1).

Various modifications of the algorithm would be possible. Some of these include

1) One could always execute the leaf move that yields the largest improvement.
2) Apertures with zero weight could be treated differently. The simulated annealing optimization typically ends up with certain aperture weights set to zero. Therefore, it is usually possible to move a leaf without changing the objective function. Eventually, after executing several moves without change in the objective value, an improvement can be achieved. In the current implementation leafs corresponding to zero weight apertures are not moved due to the threshold on the desired improvement of the objective value.
3) In a second step one could consider the simultanious movement of two leafs in order to check for a possible improvement. In this case, the search could be restricted to those pairs of beamlets which "interact", i.e. there exists a voxel which receives dose from both beamlets under consideration.

2.2. The simulated annealing method

Simulated annealing is the major step in the leaf position optimization process. The simulated annealing process corresponds to a random walk through search space that will hopefully end up with an acceptable solution (or a good starting point for the local optimization method). The general proceeding is as follows where the seperate steps are described in more detail below.

1) Determine the initial temperature T_0.
2) Perform a random walk through search space consisting of a predefined number of steps.
3) Lower the temperature T according to a cooling schedule and decrease the step width $\sigma(T)$ for the random walk steps.
4) If the final temperature is reached, stop. Otherwise go back to step 2.

The random walk itself consists of the following steps:

1) Suggest a set of leaf moves that does not violate any constraints of the leaf positions.
2) Calculate the change in the objective value.
3) Accept or reject the leaf configuration depending on the change in cost and the current temperature.
4) Go back to 1 or stop if the required number of steps was performed.

The simulated annealing method involves tens of parameters that have or may have impact on the optimization performance. These parameters have to be determined by heuristic considerations.

2.2.1. Suggesting leaf moves To propose a change in leaf configuration, all leafs of all apertures are considered simultaniously. For each leaf a random number is drawn from a Gaussian distributionof width σ. The width $\sigma(T)$ depends on the Temperature T but is the same for all leafs. Leafs are considered independently. The continuous random number is than rounded to an integer. After a move is proposed for each leaf, those moves which violate any leaf constraints are canceled, i.e. the position of these leafs remains unchanged. In general, all leafs can be moved within one iteration. Practically, the proposed step width will be zero for some or most of the leafs when the distribution width $\sigma(T)$ becomes small.

2.2.2. Accept and reject leaf moves After a leaf move is proposed, the aperture weights for the new leaf configuration are optimized. Subsequently, the difference $\Delta F = F_{new} - F_{old}$ between new and old objective value is calculated. If $\Delta F < 0$ the objective value has improved and the change of leaf configuration is accepted. If $\Delta F > 0$ the change is accepted with the probability

$$P(T) = \exp\left(-\frac{\Delta F}{T}\right) \tag{2}$$

and rejected with the probability $1 - P$. In the latter case, the old leaf configuration is restored.

2.2.3. The cooling schedule The simulated annealing process is devided into N annealing steps. Each annealing step n consists of a specified number of accepted leaf configuration changes. After each annealing step the temperature is lowered according to a cooling schedule. Three different cooling schedules were considered. The temperature in annealing step n is determined according to

1) *exponential cooling:*

$$T_n = \exp\left(-\alpha n\right) T_0 \tag{3}$$

2) *fast annealing*

$$T_n = \frac{1}{\alpha n + 1} T_0 \tag{4}$$

3) *logarithmic cooling*

$$T_n = \frac{1}{\ln(\alpha n + 1) + 1} T_0 \tag{5}$$

where T_0 is the initial temperature and α is a "time scale" parameter. Practically, the number of annealing steps N, the initial temperature T_0 and the end temperature T_N are free parameters and are used to determine the time scale parameter α.
Under these conditions, the logarithmic cooling law corresponds to a very rapid cooling down to temperatures near the end temperature. Almost all the iterations are performed for low temperatures. Figure 1 shows the dependence of the temperature on the annealing step number for $N = 10$, $T_0 = 1$ and $T_N = 0.1$.

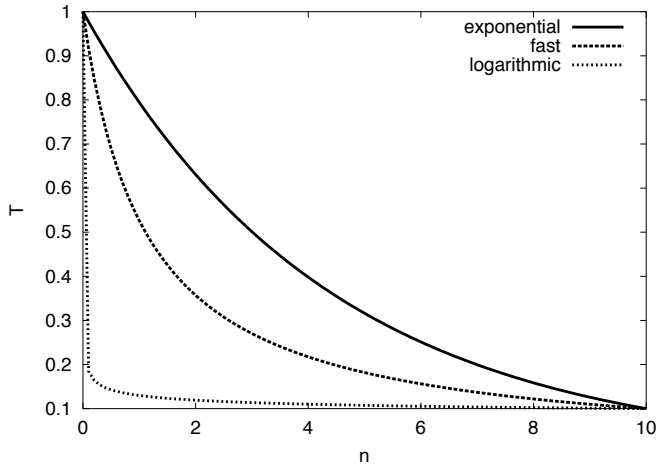

Figure 1. Dependence of Temperature on the annealing step number for different cooling schedules

2.2.4. The temperature dependence of the distribution width The width of the Gaussian distribution $\sigma(T)$ that generates the proposed leaf moves was assumed to be a linear function of the temperature, i.e. $\sigma(T)$ was chosen according to

$$\sigma(T) = \sigma_{T\to 0} + \frac{\sigma_0 - \sigma_{T\to 0}}{T_0} T \tag{6}$$

where σ_0 is the width for the initial temperature T_0 and $\sigma_{T\to 0}$ is the width when T approaches zero. For the initial width the choice $\sigma_0 = 1.0$ is made. $\sigma_{T\to 0}$ is chosen in such a way that the expected number of leaf moves per leaf configuration change approximately approaches some specified number m, i.e $\sigma_{T\to 0}$ is determined so that

$$1 - \int_{-0.5}^{0.5} \frac{1}{\sqrt{2\pi}\sigma_{T\to 0}} \exp\left(\frac{x^2}{2\sigma_{T\to 0}^2}\right) dx = \frac{m}{2\sum_{\mu=1}^{A} M_\mu} \tag{7}$$

holds. In practice, the number of leaf moves per iteration will only approximately reach m since some moves that violate constraints are omitted and the temperature does not go down to zero. The value of m was chosen to be $m = 4$.

2.2.5. Choosing the initial temperature To choose a realistic value for the initial temperature, the order of magnitude of typical changes in the objective value was estimated and the initial temperature was determined so that a predefined fraction of worsenings of this magnitude would be accepted:

1) Choose a random leaf configuration.
2) Propose a set of leaf moves according to the initial distribution width σ_0.
3) Calculate the change ΔF in the objective value

This process was repeated many times in order to calculate the average $\langle|\Delta F|\rangle$ of the absolute value of the change in the objective value. The initial temperature was then

set to

$$T_0 = -\frac{\langle|\Delta F|\rangle}{\ln(P)} \tag{8}$$

where the probability P was set to $P = 0.25$.

2.3. Speculations about finding a good solution

Practically, the number of iterations during the optimization process will be tiny compared to the number of possible leaf configurations S. Only a very small fraction of all aperture combinations will be explored during the optimization. The question is under which condition one can expect to find a good solution or even the optimal solution. Some speculations can be made, although very handwavy and not mathematically sound:

1) Looking at very idealized geometries (section 4) suggests that the DAO problem might have many states at distant locations in search space which have the same or approximately the same cost. In the field of statistical physics this feature corresponds to a degenerate energy ground state of the physical system and is also refered to "frustration". An example for systems known to exhibit degenerate ground states are "spin glasses".
2) If the DAO problem was actually frustrated, there might be some hope that a good solution can be found by searching only a small area of the search space with a relatively small number of iterations.
3) However, even if there are multiple near optimal solutions, the optimization may still get stucked in rather bad solutions surrounded by states with higher cost, so that it is unlikely to escape from these states.
4) It does not seem to be the case that there exists one optimal solution that is much better than all other solutions. However, if this was the case, it would probably be almost impossible to find this solution using simulated annealing and a rather small number of iterations.

Observations for this study indicate that in particular 3) seems to apply.

2.4. Questions addressed in this work

The following list summarizes the questions that are addressed in the subsequent result part of the report.

1) For very idealized geometries, DAO problems can be constructed that have perfect solutions in the sence that the desired dose distribution can be realized exactly. It is investigated whether the simulated annealing algorithm can reproduce such optimal solutions (section 4).
2) It is further investigated whether reasonable solutions can even be achieved by applying the local optimization method alone without a simulated annealing step before.

3) The performance of a single optimization run is characterized by different means. An attempt is made to characterize the walk in search space during the optimization (section 5).
4) For (slightly) more complex geometries, the quality of discovered solutions is discussed (section 6).
5) DAO optimization runs are performed for multiple initial states. It is investigated whether different runs lead to similar solutions or whether they end up with completely different aperture sets. An attept is made to find a quantitative measure that describes the diversity of different solutions (section 6.1).

3. Model of idealized geometry

In order to investigate the performance of simulated annealing for the DAO problem, a two-dimensional model of idealized geometry is considered. The model constist of a two-dimensional grid of voxels representing one slice of the patient. The grid is irradiated from up to four directions. The orientation of the beam directions is parallel/vertical to the grid orientation. The resolution of beamlets is identical to the resolution of voxels and lateral scattering is neglected. Hence, a single beamlet contributes dose only to a single column/row of the voxel grid. The dose distribution within one column/row was determined from a proton depth dose curve by taking the dose values at equidistant depths. The bragg peak of two neighboring beamlets (i.e. they correspond to the same beam direction and the same range), which irradiate two columns, are located in the same row and vice versa. The orientation of the MLC leafs is parallel to the slice of the patient for all beam directions.

3.1. The objective function

A standard quadratic objective function is used for dose optimization:

$$F = \sum_{i \in CTV} \alpha_{CTV}(D_i - D_{CTV})^2 + \sum_{i \in OAR} \alpha_{OAR} D_i^2 + \sum_{i \in UT} \alpha_{UT} D_i^2 \tag{9}$$

where α_{CTV}, α_{OAR} and α_{UT} are penalty factors for the CTV, the OAR and the Unclassified Tissue, respectively. D_{CTV} is the prescribed dose to the tumor. Standard parameters are $\alpha_{CTV} = 1.0$, $\alpha_{OAR} = 10.0$, $\alpha_{UT} = 0.1$ and $D_{CTV} = 1.0$.
The dose is evaluated using an underlying dose contribution matrix concept. The dose in voxel i is given by

$$D_i = \sum_{\mu=1}^{A} d_\mu^i w_\mu \tag{10}$$

where w_μ is the weight of aperture μ and d_μ^i is the dose contribution of aperture μ to voxel i. The dose contributions d_μ^i of the apertures are calculated from the dose

contribution of beamlets and are updated after each leaf configuration change:

$$d_\mu^i = \sum_{j=L+1}^{R-1} d_{\mu j}^i \tag{11}$$

where L is the last bixel covered by the left leaf, R is the first bixel covered by the right leaf and $d_{\mu j}^i$ is the dose contribution of bixel j in aperture μ to voxel i.

3.2. Interpreting graphics output

In the following sections, color plots are used to display dose distributions and apertures. The center of these plots show dose distributions where brown color corresponds 1.0, i.e. the prescribed dose and light blue is 0.0. The outer parts display the apertures. The fluence is color coded, again with brown corresponding to 1.0 and light blue to 0.0. The color bars beside the dose array and the apertures code the position of the bragg peak of the respective aperture within the dose array.

4. Reproducing an ideal solution

In this section, the idealized patient anatomy in figure 5a is considered. In addition, the following settings apply:

1) The dose contribution to voxels behind the bragg peak of a beamlet is neglected.
2) For each beam direction one aperture per range is allowed.
3) The unclassified tissue surrounding tumor and OAR is not penalized in the objective function, i.e. $\alpha_{UT} = 0.0$.

Under these conditions, multiple solutions exist that deliver the prescribed dose to all tumor voxels and no dose to the OAR. The ideal objective value of zero can thus be achieved. It is investigated whether such ideal solutions (which can easily be constructed geometrcally) can be reproduced by the simulated annealing algorithm or even by the local optimization alone.

4.1. Applying local optimization only

In a first attempt it is investigated whether reasonable solutions can be found by only applying the local optimization method described in section 2.1.4. No simulated annealing is involved. 100 optimization runs where performed, starting from different random initial leaf configurations. Figure 2a shows a histogram of the objective values obtained. The obtained solutions cover a wide range of dose distributions of different quality. The best solution found is suprisingly good and is shown in figure 3. A complete sparing of the OAR is enforced by the high penalty factor for the OAR. The dose distribution within the CTV is quite homogeneous. However, the actual set of apertures is not intuitive. Figure 4 shows an example for an average quality solution. The dose within the

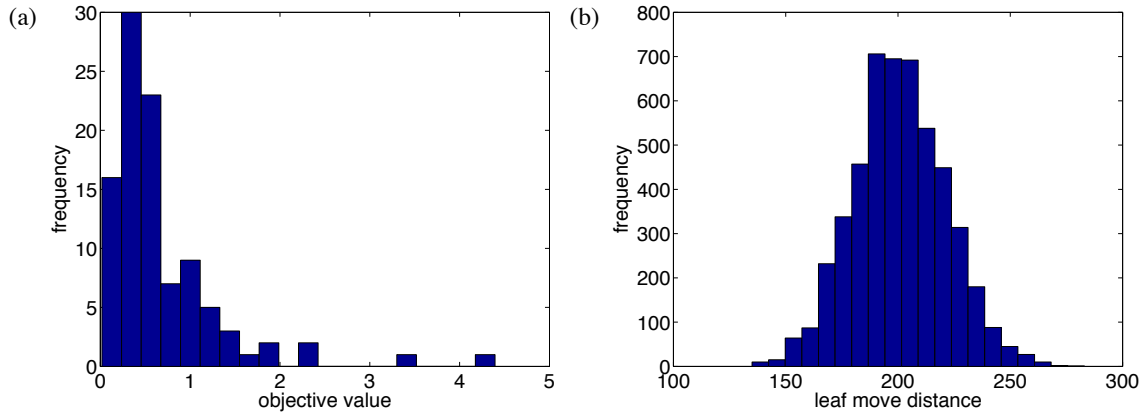

Figure 2. Histogram of the objective values obtained for 100 optimization runs for local optimization only

CTV is quite inhomogeneous and would not be acceptable.

Figures 2b,5b and 5c address the diversity of different solutions. Figure 2(b) shows a histogram of the pair wise distance of the 100 solutions in search space. As a measure of distance the number of required leaf moves to transfer one aperture into another is used. This means, two apertures sets have distance one if one leaf has to be moved by one beamlet in order to make both solutions identical. The average distance is in the order of 200 single leaf moves. This number can be compared to the average distance of completely random aperture sets. For random apertures, the average is empirically found to be approximately the same, i.e. only slightly larger. This means, that random initial states do not evolve towards a certain region in search space on a global scale. Under the local optimization method solutions evolve towards completely distinct states (or remain distinct).
Figure 5b shows the average dose distribution resulting from the superposition of the 100 solutions. The average dose distribution is quite homogeneous within the CTV, although individual solutions produce strongly inhomogeneous dose distributions. However, deviations from the prescribed dose are obviously randomly distributed around the prescribed dose when a large set of solutions is considered. Figure 5c shows the spatial distribution of the standard deviation of the dose. The largest values occur in the unclassified tissue surrounding the tumor, indicating that the dose delivered to voxels may come from different beamlets for different solutions. The diversity of solutions will be discussed in more detail in section 6.1.

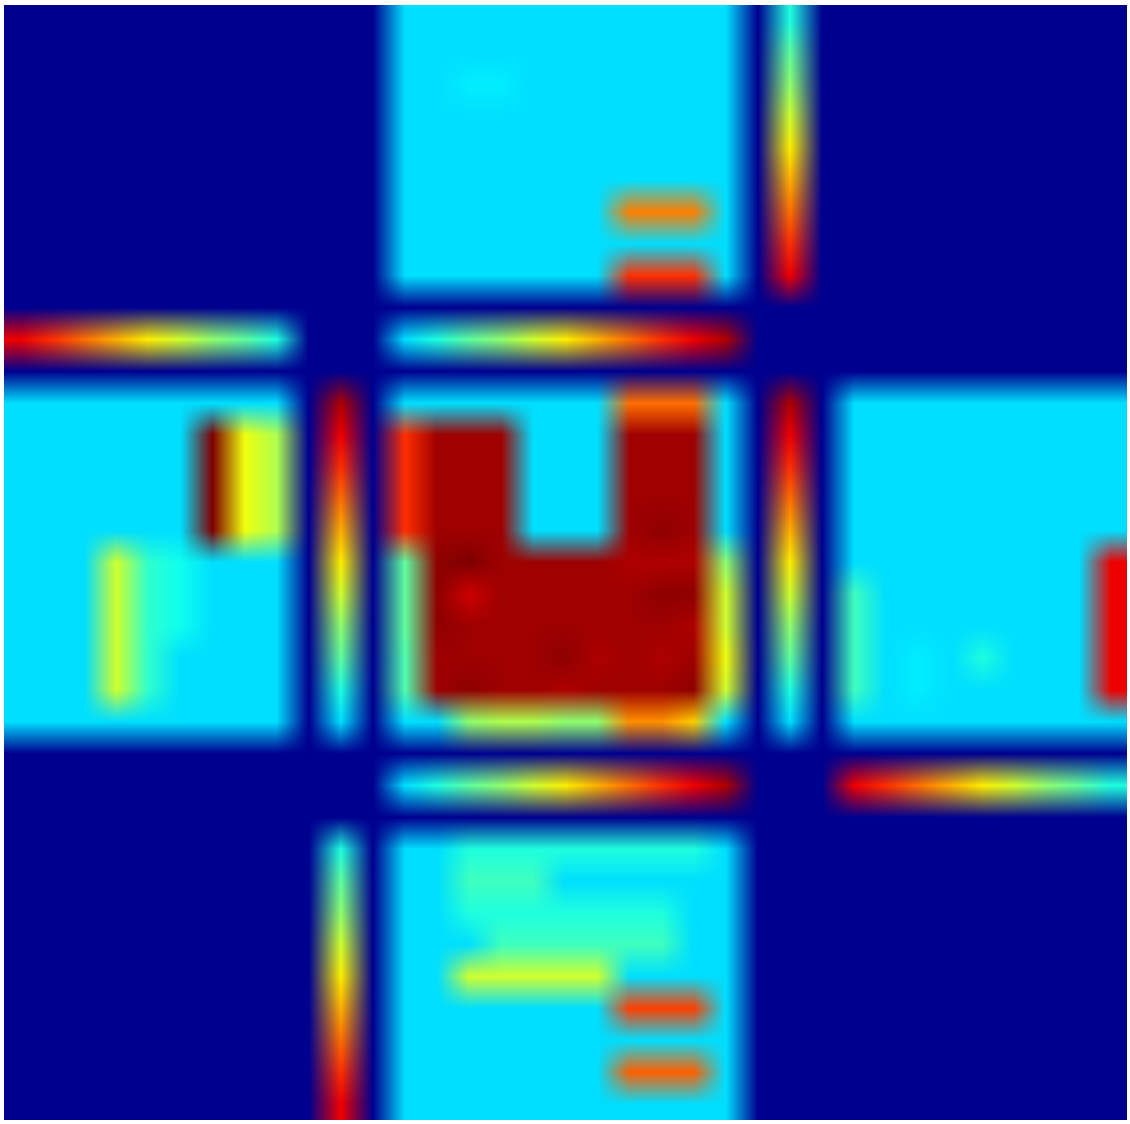

Figure 3. Best solution found in 100 local optimizations, corresponding to an objective value of 0.018

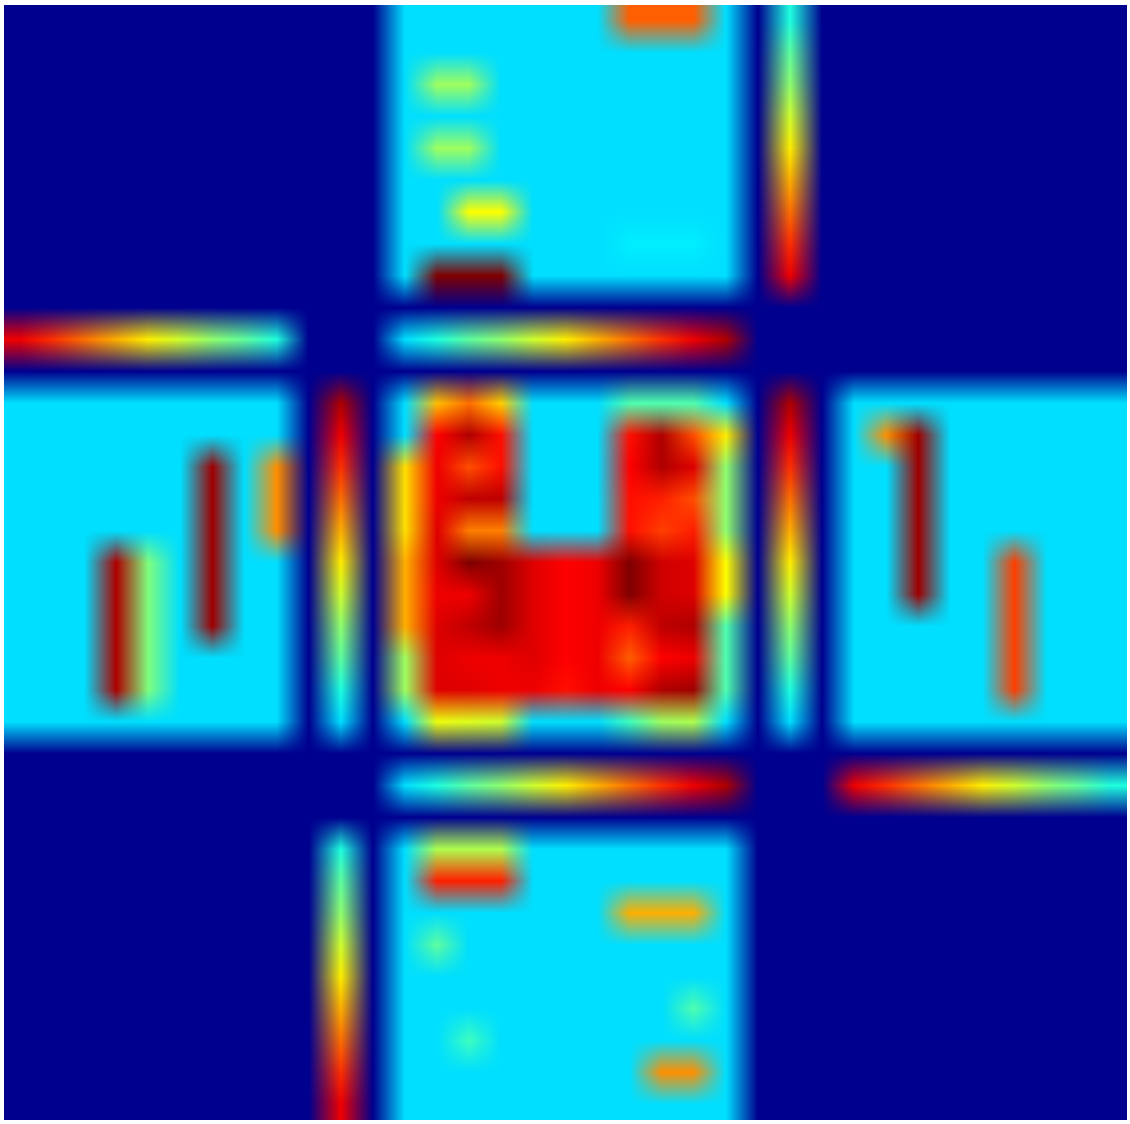

Figure 4. Average quality solution corresponding to an objective value of 0.639

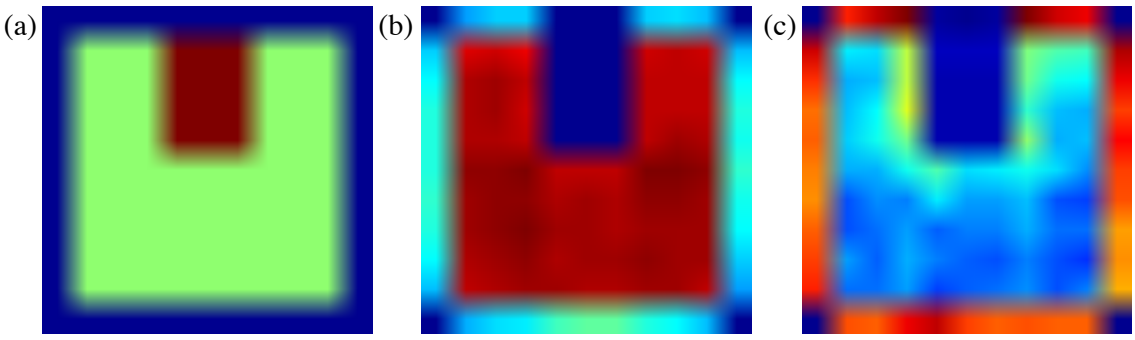

Figure 5. (a) Geometry of the tumor (green), the OAR (brown) and the unclassified tissue (blue) (b) Average dose (c) Standard deviation (both averaged over 100 optimization runs)

4.2. Simulated annealing using a logarithmic cooling schedule

In the next step, a simulated annealing step is performed before the local optimization method is applied. The logarithmic cooling schedule is applied. The end temperature was 2% of the initial temperature, 50 annealing steps were performed and one annealing step consisted of 200 accepted leaf configuration changes. Again, 100 optimization runs were performed. Figure 6 shows the histogram of the obtained objective values. When compared to the pure local optimization, the average quality of the solution has drastically improved. However, only once a perfect solution was found, i.e. the prescribed dose in every CTV voxel is realized and the OAR receives no dose. Figure 7 shows the solution which is a bit unexpected. The figure indicates that there will be quite a large number of perfect solutions, some of which one may not think of initially. Figure 8 shows a solution that corresponds to an objective value of 0.0207 which approximately represents to the average objective value. For some parts of the CTV, an intuitive amd easy way of delivering the dose was found, however, not for the whole tumor.

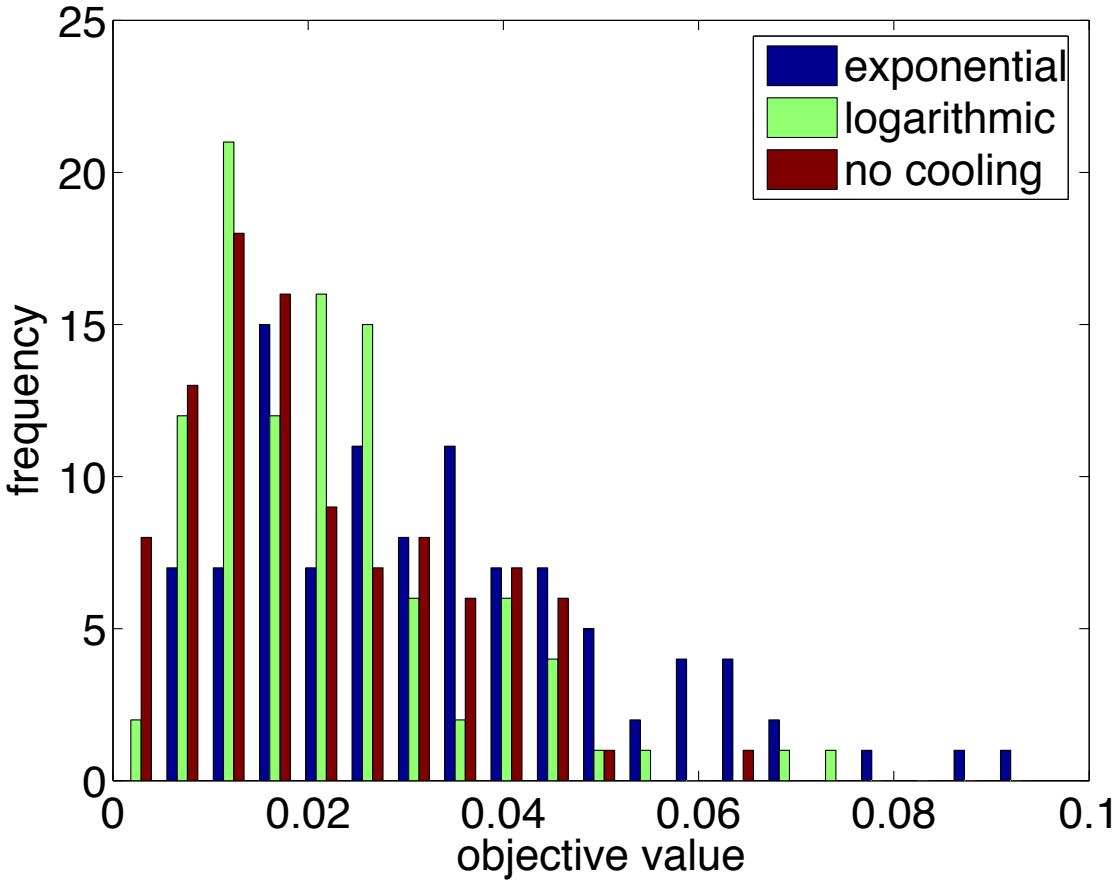

Figure 6. Histogram of the objective values obtained for 100 optimization runs for simulated annealing using different cooling schedules

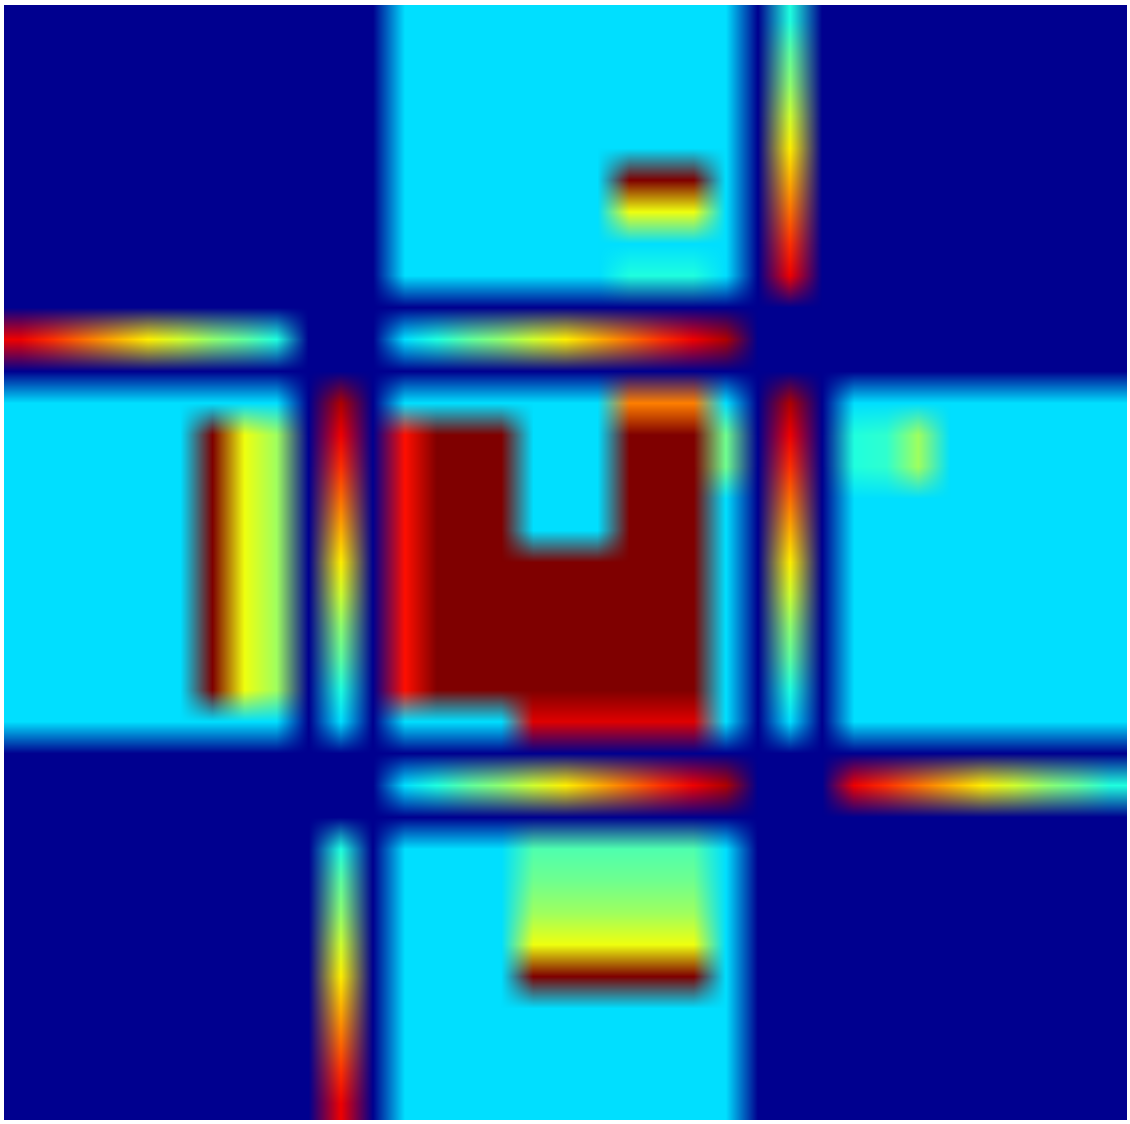

Figure 7. Perfect solution found using logarithmic cooling

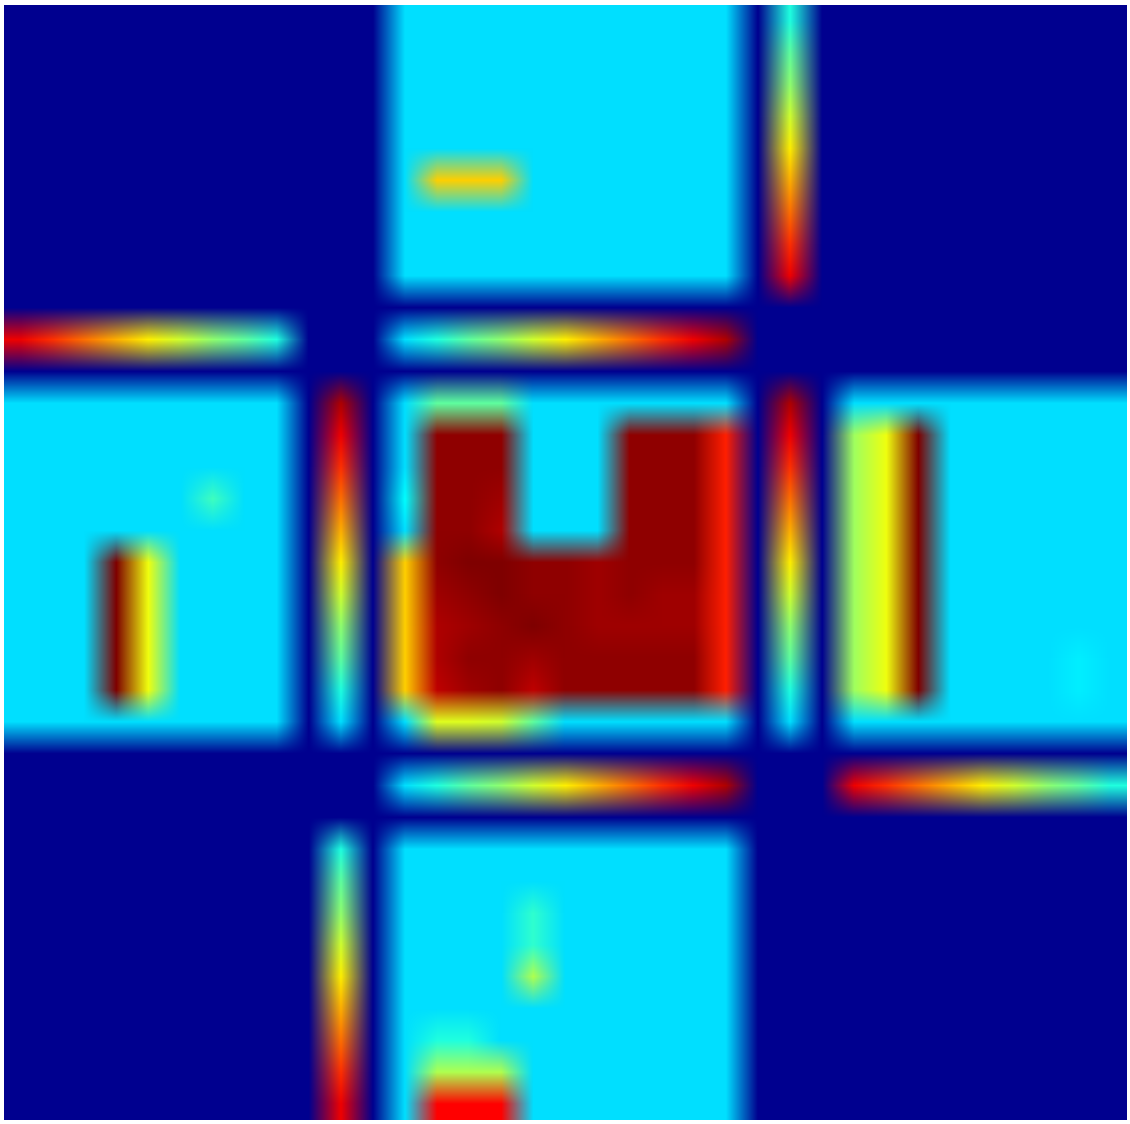

Figure 8. Average quality solution found using logarithmic cooling corresponding to an objective value of 0.0207

4.3. Using other cooling schedules

The performance of different cooling schedules was compared. Figure 6 compares the histograms of the objective values from 100 optimization runs for three different cooling schedules: First the logarithmic schedule, second, an exponential cooling law and third, no cooling with a constant temperature corresponding to the end temperature of both other cooling schedules. The exponential cooling schedule permormed worse compared to both other methods. No perfect solution was found and the histogram of the objective values is shifted towards higher values. For both, the logarithmic and the exponential cooling, the end temperature was 2% of the initial temperature and 50 annealing steps were performed. For the logarithmic cooling each annealing step consisted of 200 accepted leaf configuration changes, whereas for the exponential cooling 1000 accepted leaf moves were required. The computation time is approximately proportional to the total number of proposed leaf configuration changes including the rejected ones. Hence, the computation time was larger for exponential cooling, although by less than a factor of 5.0 since the ratio of accepted and rejected leaf moves is larger for higher temperatures. Seemingly, the random walk process at high temperatures does not yield much benefit, at least not at the expense of reducing the number of iterations at low temperatures. Similar results like for the logarithmic cooling schedule were obtained when no cooling was performed at all and the temperature was set to 2% of the initial temperature from the first iteration (figure 6). In this case even four perfect solutions were found. (However, the computation time for 200 accepted leaf configuration changes was slightly larger for this schedule since the number of rejected moves is increased.)

In comparison to the pure local optimization, introducing a simulated annealing procedure drastically improves the average quality of the solutions found and in particular avoids extremly bad solutions. However, this comes with a large increase in computation time. Hence, a large number of local optimizations runs (exploring many initial states) could be performed within the same amount of computation time. And eventually, some of these runs lead to a reasonably good solution. For example, the best solution found for pure local optimization competes with average quality solutions of the simulated annealing approach. Hence, for practical purpose, such a brute force method of exploring a huge number of initial conditions must not be fully discarded.

5. Characterizing a single optimization run

Figure 9 shows the objective value as a function of the iteration number (were one iteration is one accepted leaf configuration change) for an exponential cooling schedule. For comparison, figure 10 shows the objective value as a function of the iteration number for the logarithmic cooling schedule. For the exponential schedule there is a rather smooth drop of the average objective value. For the logarithmic schedule there is a steep drop after the first annealing step. After that, the average objective value stays more or less the same.

Figures 9b and 10b show the objective values for temperatures near the final temperature. The aperture sets explored correspond to objective values in the range from 0.2 to 0.6, approximately. Practically, these objective values are quite large in the sence that they correspond to rather bad dose distributions. From this perspective, it seems reasonable to further lower the temperature. On the other hand, leafs are moved in steps of one bixel. Hence, changes to the dose distribution can not always be small, especially when an aperture with large weight is changed. Hence, lowering the temperature to much will cause certain apertures to freeze out and remain unchanged because every attempt to change it will be rejected (see section 5.1).

Figure 11a shows the number of rejected leaf configuration changes as a function of temperature for the exponential cooling schedule. For low temperatures a relatively large fraction of all suggested leaf moves is rejected, hence making the algorithm quite inefficient. Figure 11b shows the average number of leaf moves per leaf configuration change as a function of temperature for both the accepted leaf moves and the rejected ones. Generally, the number of moved leafs per iteration decreases when the cooling proceeds since the width of the Gaussian distribution becomes smaller (section 2.2.4). Especially for low temperatures, leaf configuration changes that propose a smaller number of leaf moves are more likely to be accepted as configuration changes that want to move a larger number of leafs.

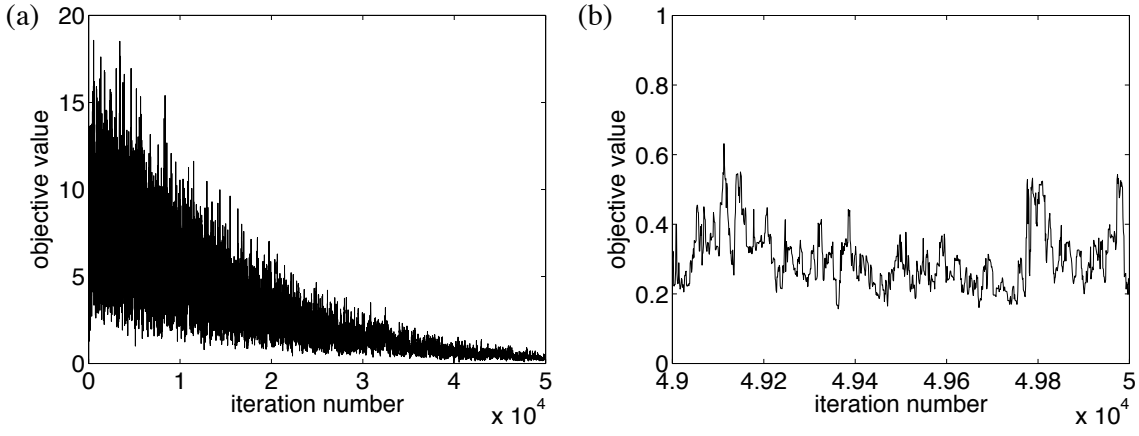

Figure 9. Objective value as a function of the iteration number for exponential cooling

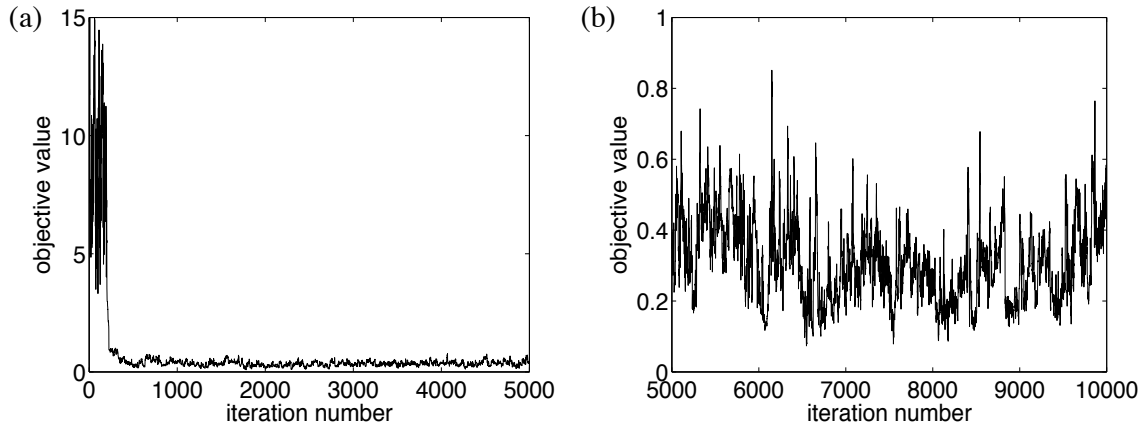

Figure 10. Objective value as a function of the iteration number for the logarithmic cooling schedule

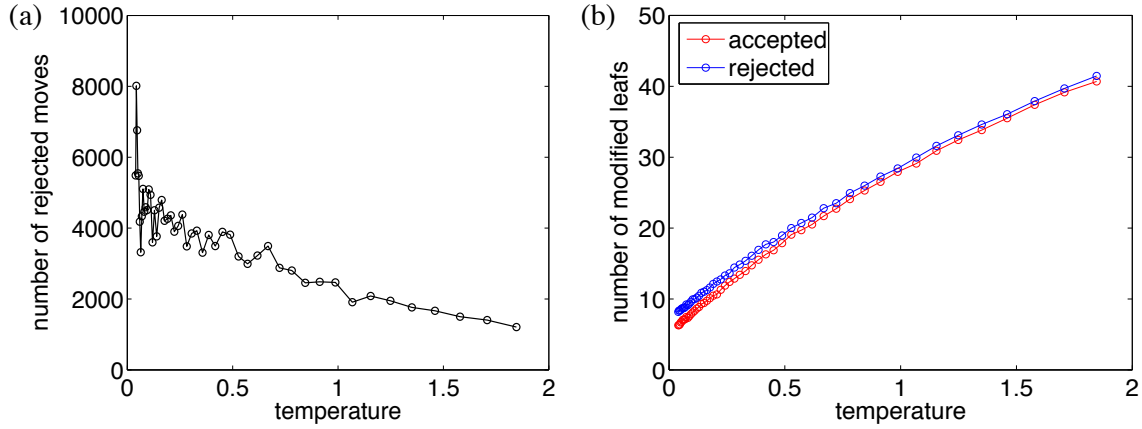

Figure 11. (a) Number of rejected leaf configuration changes as a function of temperature for exponential cooling (b) Average number of moved leafs per iteration as a function of temperature

5.1. Characterizing the walk in search space

One of the major questions that arise, concerns the set of states which are finally explored by the simulated annealing method. Does the random walk in search space remain restricted to a very limited area so that all states explored are rather simular? Or does the method consider significantly different states on the way? Visualizing the walk in search space is not easy, however, an attempt can be made.

The logarithmic cooling schedule is considered first. Figure 12 shows for each beamlet the frequency that this beamlet is changed from open to closed or vice versa. The main observation is as follows. There are rows of beamlets, corresponding to a certain aperture that are rarely changed. In contrast, there are beamlets which are subject to much more frequent changes. Figure 13 shows the average fluence of the beamlets during the same optimization run. The frequency of beamlet changes from open to

close is obviously correlated with the average fluence of these bixels (or the weight of the corresponding aperture). The rows of beamlets in figure 12 which are rarely changed correspond to apertures that deliver a high dose to the tumor, whereas the apertures that deliver a low dose are modified more frequently. Hence, simulated annealing at low temperatures as this is the case for the logarithmic cooling schedule seems to confine itself to certain regions of the search space were some key apertures are almost fixed and only the residual degrees of freedom are further varied. The center of figure 12 shows the averaged absolute change of the dose between to iterations for each voxel. It seems consistent with the frequency of beamlet changes, however, does not contain new information. The center of figure 13 shows the averaged dose distribution over the optimization process.

For the exponential cooling schedule, iterations occur at higher temperatures making it easier to accept worsenings of the objective value. As a consequence, the distribution of beamlet changes from open to close is much more washed-out (figure 14).

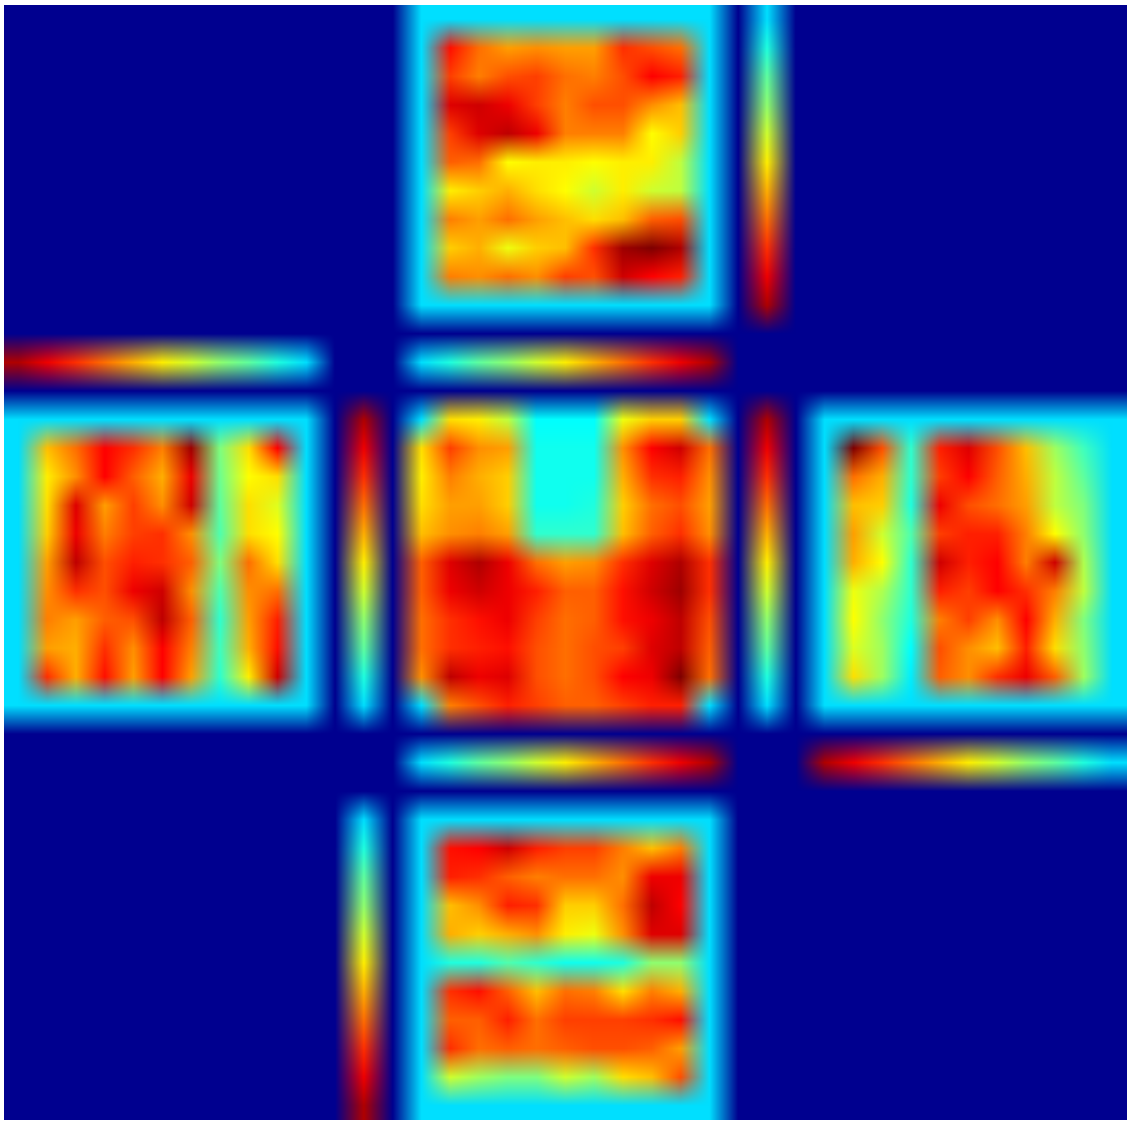

Figure 12. bixel display: Frequency of changing a beamlet from open to close and vice versa, dose display: Averaged absolute change of dose values between to iterations (for the logarithmic cooling schedule)

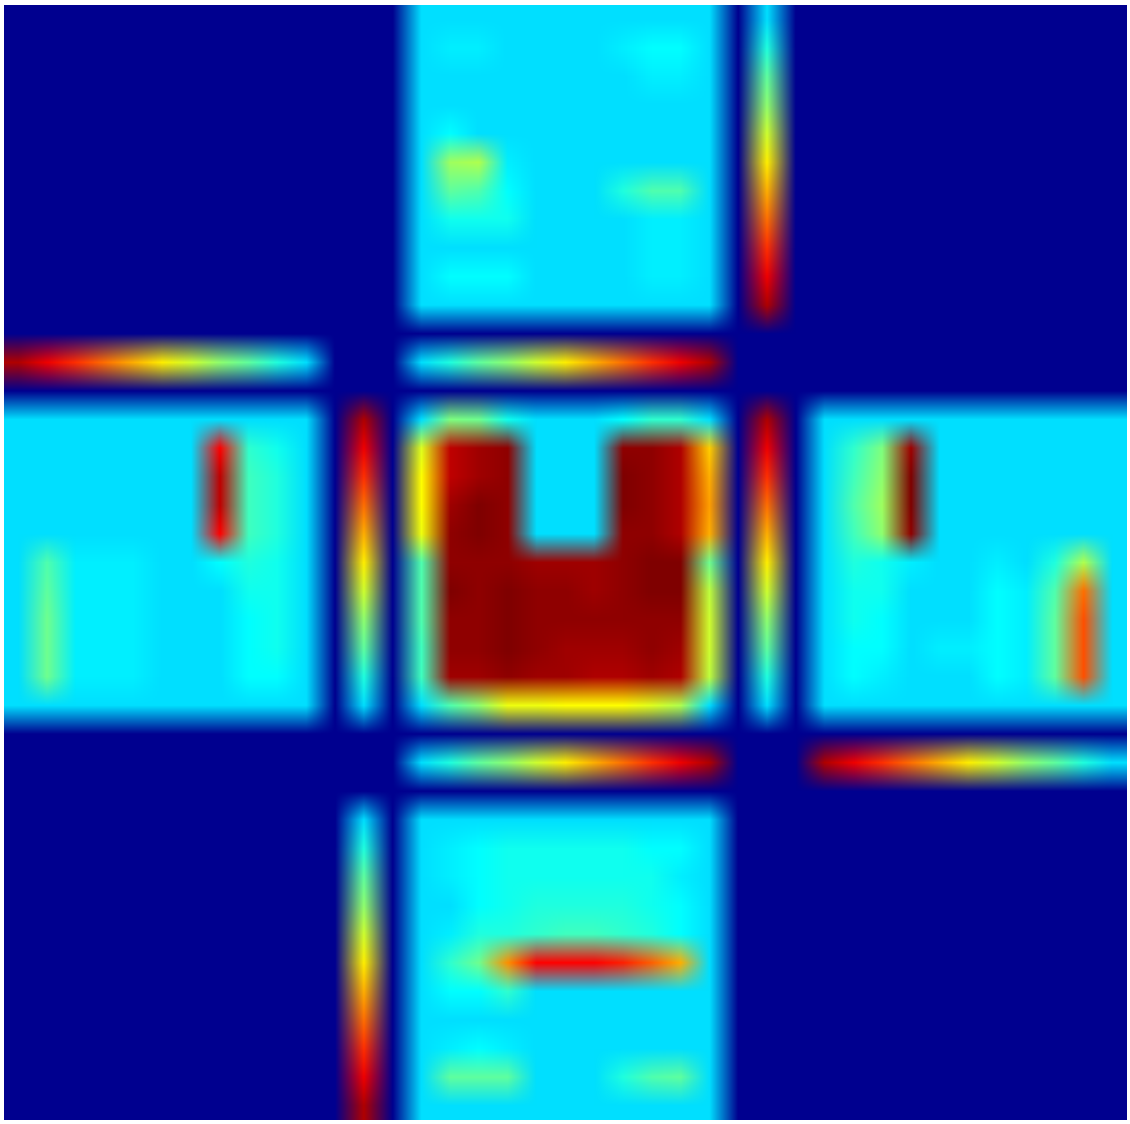

Figure 13. bixel display: Average fluence of beamlets during the optimization, dose display: Averaged dose of voxels during the optimization

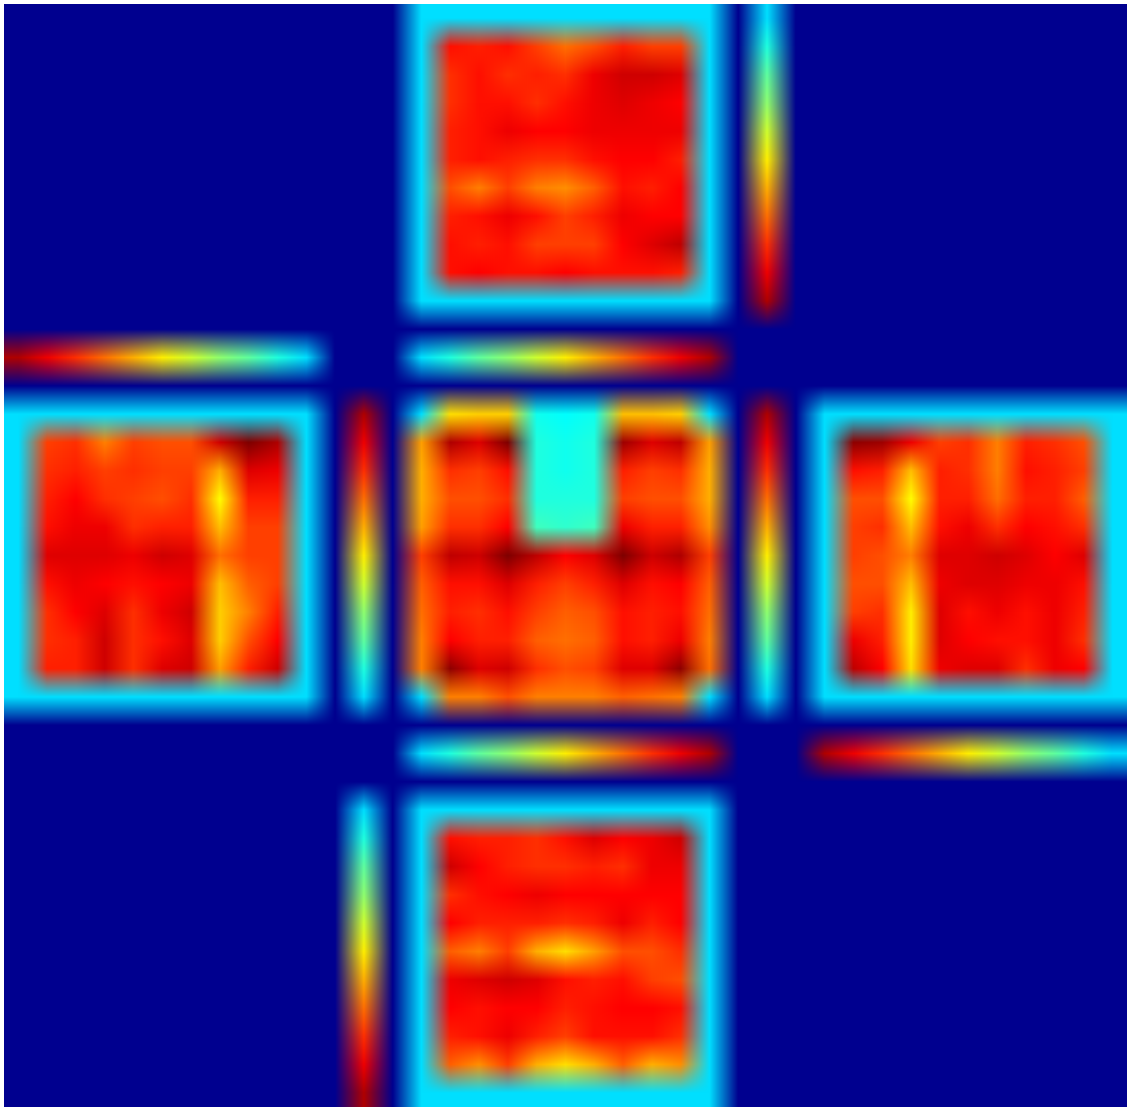

Figure 14. bixel display: Frequency of changing a beamlet from open to close and vice versa, dose display: Averaged absolute change of dose values between to iterations (for the exponential cooling schedule)

6. A more complex geometry

The tumor geometry in figure 5a is particularly simple because the tumor can be divided into three or four rectangles. Realistically, the tumor boundary will have a more complex shape. It is investigated whether reasonable solutions can be found for the slightly more complex geometry in figure 20a. In order to enforce a conformal dose distribution, dose to the unclassified normal was penalized in the objective function. The results shown were obtained using a logarithmic cooling schedule, 50 annealing steps consisting of 200 accepted leaf configuration changes each, the final temperature was 2% of the initial temperature, one aperture per range and beam direction is used and 100 optimization runs were performed.

Figure 17 shows the best solution found. For this solution, the maximum tumor dose is 1.07, the minimum tumor dose is 0.89 and the standard deviation of the spatial dose distribution in the tumor is 0.036. Figure 18 shows the worst solution found. For this solutions, the maximum tumor dose is 1.11, the minimum tumor dose is 0.84 and the standard deviation of the spatial dose distribution in the tumor is 0.050. Generally, it becomes much more difficult to deliver a homogeneous dose distribution to the tumor if the boundary has an irregular shape. Furthermore, the sparing of healthy tissue at the edges of the tumor is often not achieved.

Figure 16a shows the dependence of the obtaind minimum objective value on the number of apertures used. For each optimization run, the total number of available apertures is fixed and given by 36. However, after the simulated annealing process, some apertures are in a non-usable shape and the weight of these apertures is set to zero. Figure 16a shows that the solution is usually better when more apertures are finally used.
Figure 16b shows the correlation of the best objective value obtained after the simulated annealing process and the final objective value after the subsequent local optimization step. A good best solution after simulated annealing usually leads to a better final solution.

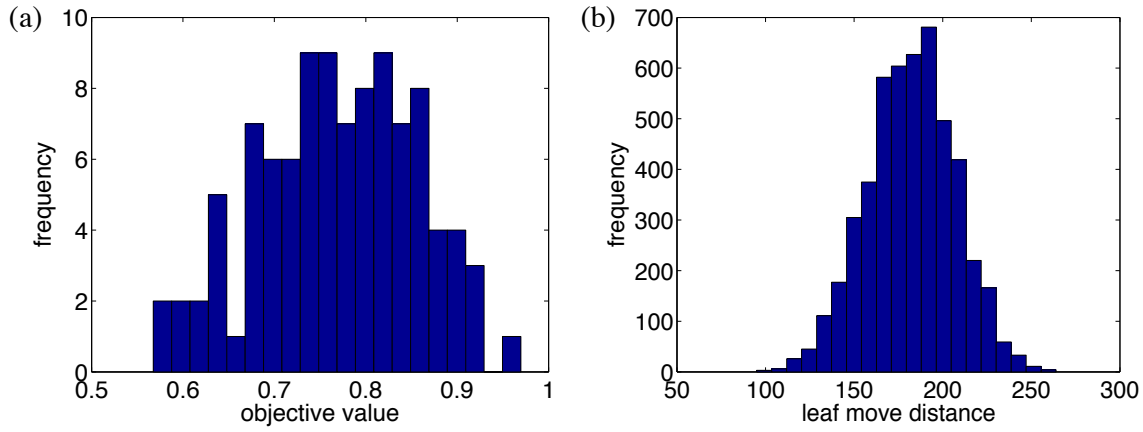

Figure 15. (a) Histogram of the objective values obtained for 100 optimization runs for logarithmic cooling (b) Histogram of the pair wise leaf move distance of the solutions

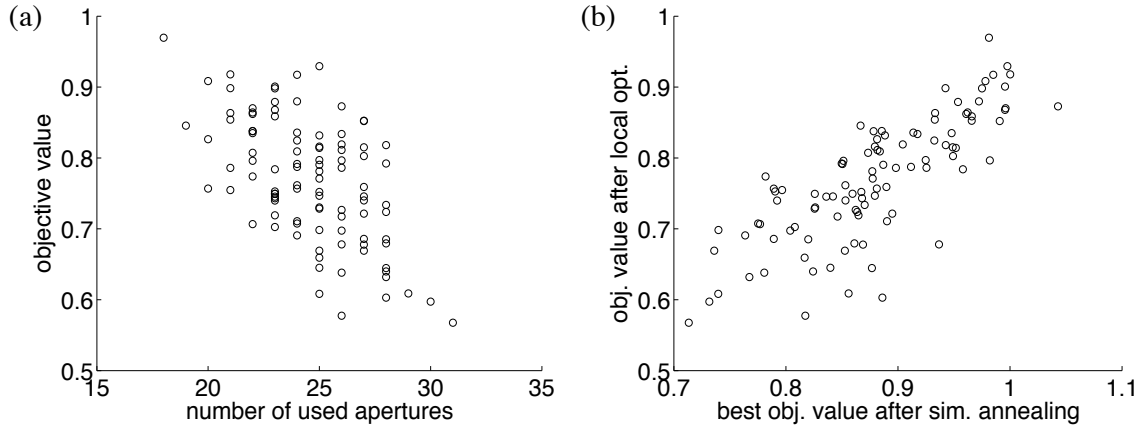

Figure 16. (a) Correlation of the objective value after optimization and the number of used apertures (b) Correlation of the best objective value after simulated annealing and after subsequend local optimization

6.1. Diversity of solutions

This subsection discusses the diversity of solutions obtained in different optimization runs. Figure 15b shows the histogram of the pair wise leaf move distance for the 100 solutions. Hence, on the average approximately 180 times a single leaf has to be moved by one beamlet in order to transform one solution into another. The average distance of random leaf configurations is only slightly larger. Note that, the maximum distance of two states is $36 \cdot 10 \cdot 2 = 720$ which corresponds to moving both, left and right leaf from the extrem left to the extrem right position or vice versa. This indicates that solutions from different optimizations are at distant positions in search space. Although this measure describes the seperation of states in search space reasonably well, it does not necessarily mean that the obtained solutions are physically different. For example, two sets of apertures may differ in leaf positions that belong to apertures that deliver almost no dose.
As a better indicator for the physical distinctness of solutions, the frequency that a beamlet is used can be considered. Figure 19 shows for each bixel the frequency that it was part of an aperture with non-zero weight at the end of the optimization. Obviously, there are bixels which are used in almost every solution. This holds, for example, for the beamlets irradiating from left and right direction and stopping in front of the OAR. For some bixels the symmetry of the geometry and the fact that only one aperture per range was allowed has to be considered when interpreting figure 19. The beamlets irradiating from below and placing the bragg peak at the distal edge of the tumor are used in approximately 50% of the cases. Hence in almost 100% of the solutions either the bixels left of the OAR or the bixels right of the OAR are used. Some beamlets that would deliver dose to the OAR are obviously never used. However, there remains a large number of bixels which are occasionally used but not always. Hence, this class of bixels seem to cause the distinctness of solutions. In summary, the different solutions have some common features since some beamlets or used in almost every solution. However, there remains a large set of degrees of freedom that is used in different ways.

It is observed that individual solutions often do not yield a homogeneous dose distribution within the tumor. The question arises whether cold spots and hot spots do always occur at the same positions for different solutions or whether they are rather randomly distributed in space. To assess this question, figure 20b shows the average dose distribution that results from suming up the dose from different solutions. Figure 20c shows the corresponding standard deviation of the dose in each voxel. On average, the inhomogeneities of individual solutions level out and the sum of all 100 solutions yields a pretty homogeneous dose distribution on the tumor. The deviations from the prescribed dose in individual solutions seem to be randomly distributed around the prescribed dose. Figure 20c in addition shows that the standard deviation within the tumor is quite constant. Hence, hot and cold spots may occur in all parts of the tumor almost equally distributed. The largest differences in dose distributions occur at the

edges of the tumor. Some solutions find a way to spare the unclassified tissue at these voxels, others do not.

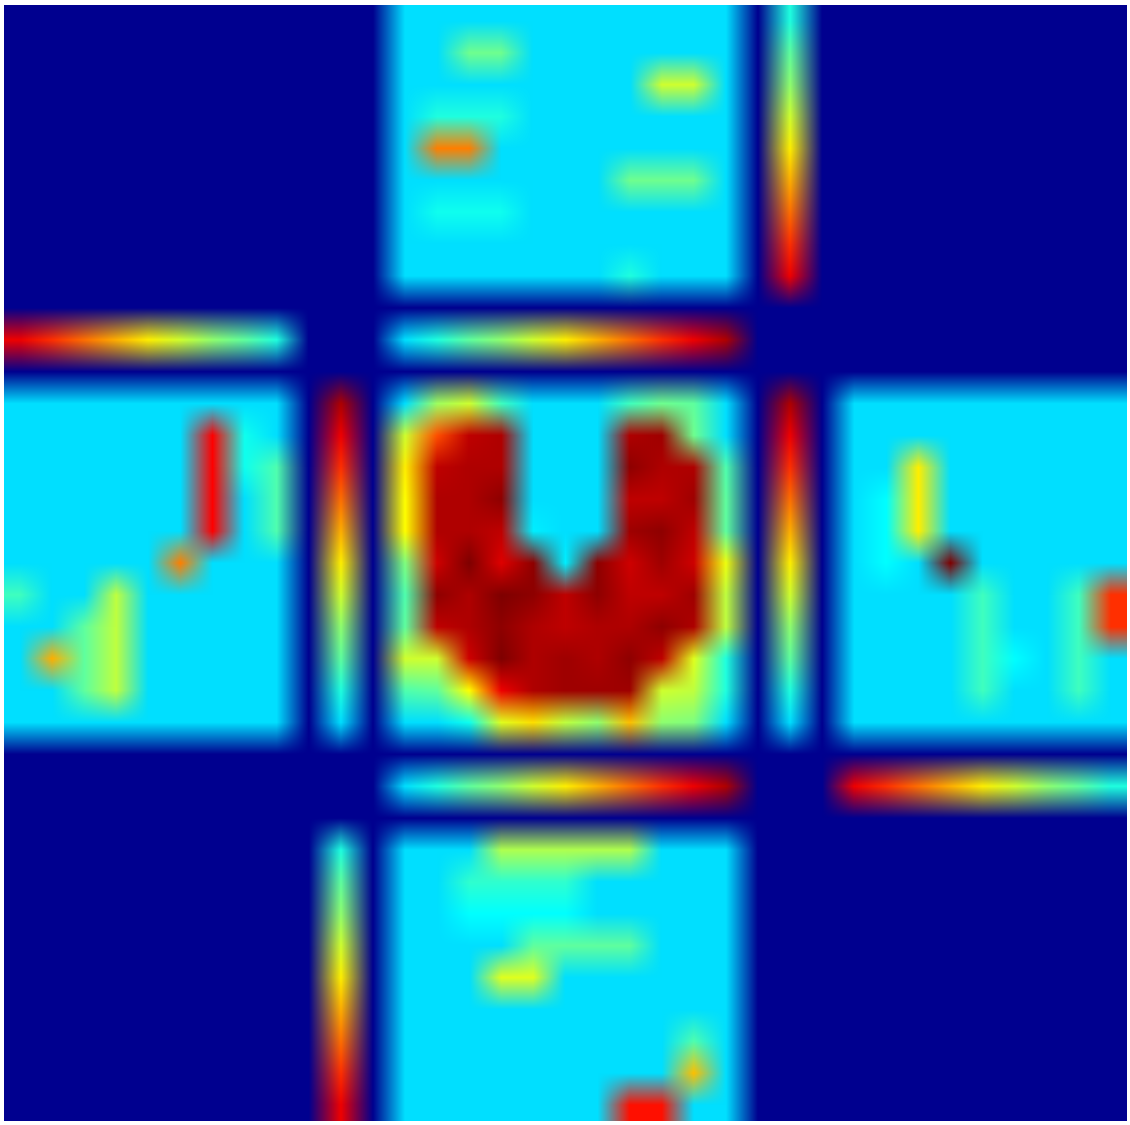

Figure 17. Best solution found in 100 local optimizations, corresponding to an objective value of 0.568

6.2. Making use of the diversity of solutions?

The inability to find solutions that deliver a homogeneous dose to the tumor may be a problem. For the example presented here, cold spots up to 85% of the prescribed dose occured in some solutions. Concerning this problem, the diversity of different solutions and the observation that cold spots are located in different positions may yield advantages. If different solutions are applied in different fractions, the total dose distribution delivered can be quite homogeneous even though the dose in single fractions

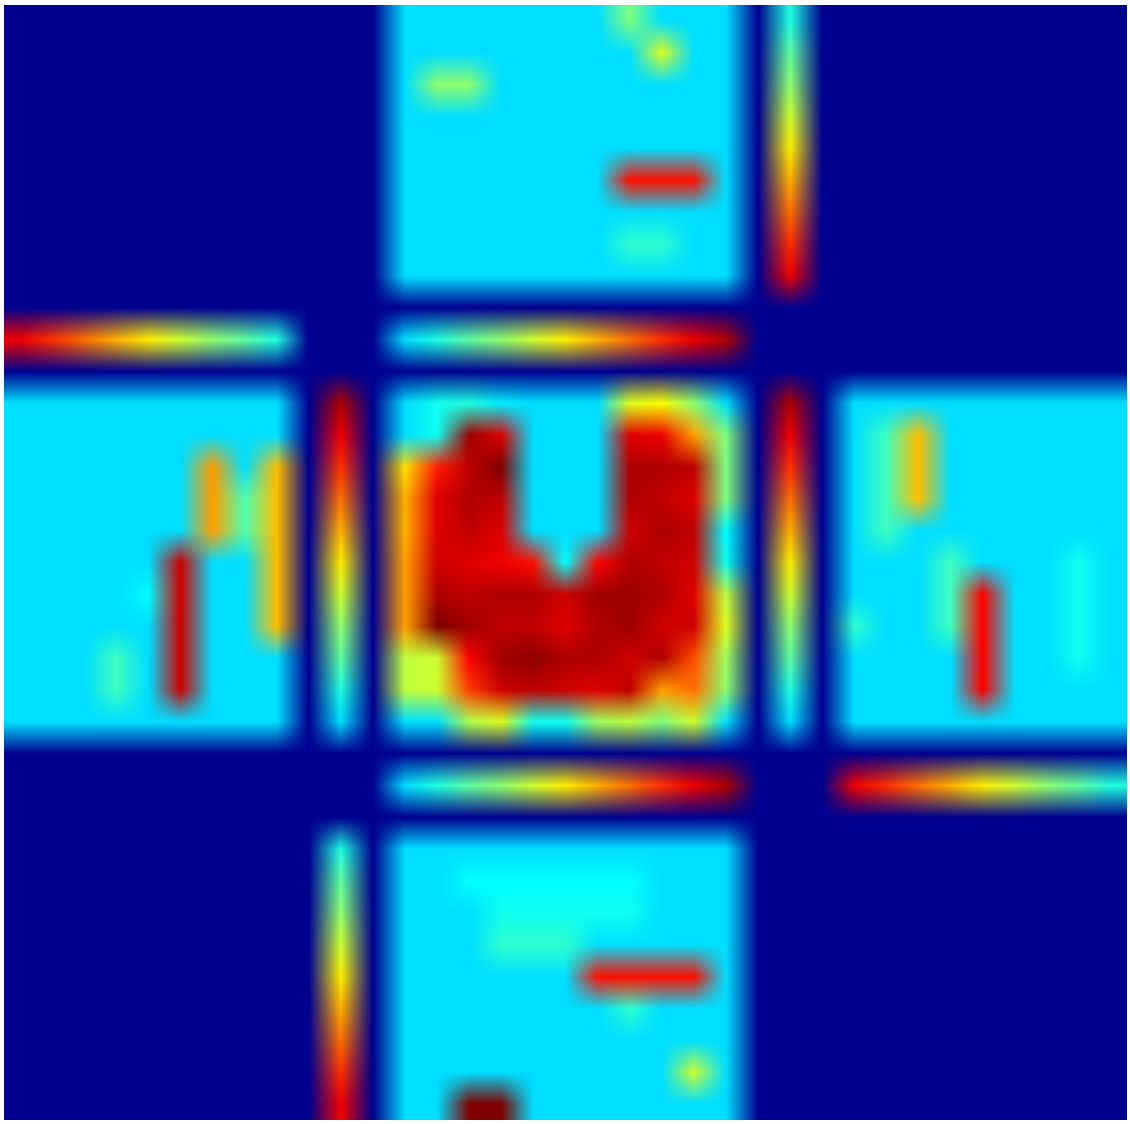

Figure 18. Worst solution found in 100 local optimizations, corresponding to an objective value of 0.917

may show hot and cold spots. ‡ Assuming that a pool of K solutions was generated, mixed integer programming could be used to determine which solutions are used for treatment and for how many fractions. Let N_k be integer variables that describe the number of fractions that are irradiated using solution k with dose distribution D_k. Let Q be the total number of fractions. A mixed integer program could optimize some objective function that depends linearly on the total dose distribution $D = \sum_{k=1}^{K} N_k D_k$, subject to the constraint $\sum_{k=1}^{K} N_k = Q$. It would still be the aim of the simulated annealing method to produce homogeneous dose distributions for each fraction. The secondary

‡ Irradiating different plans over the coarse of treatment may have further advantages in the presence of range uncertainties. The impact of uncertainties of the proton range may be reduced if dose gradients are not always located at the same position.

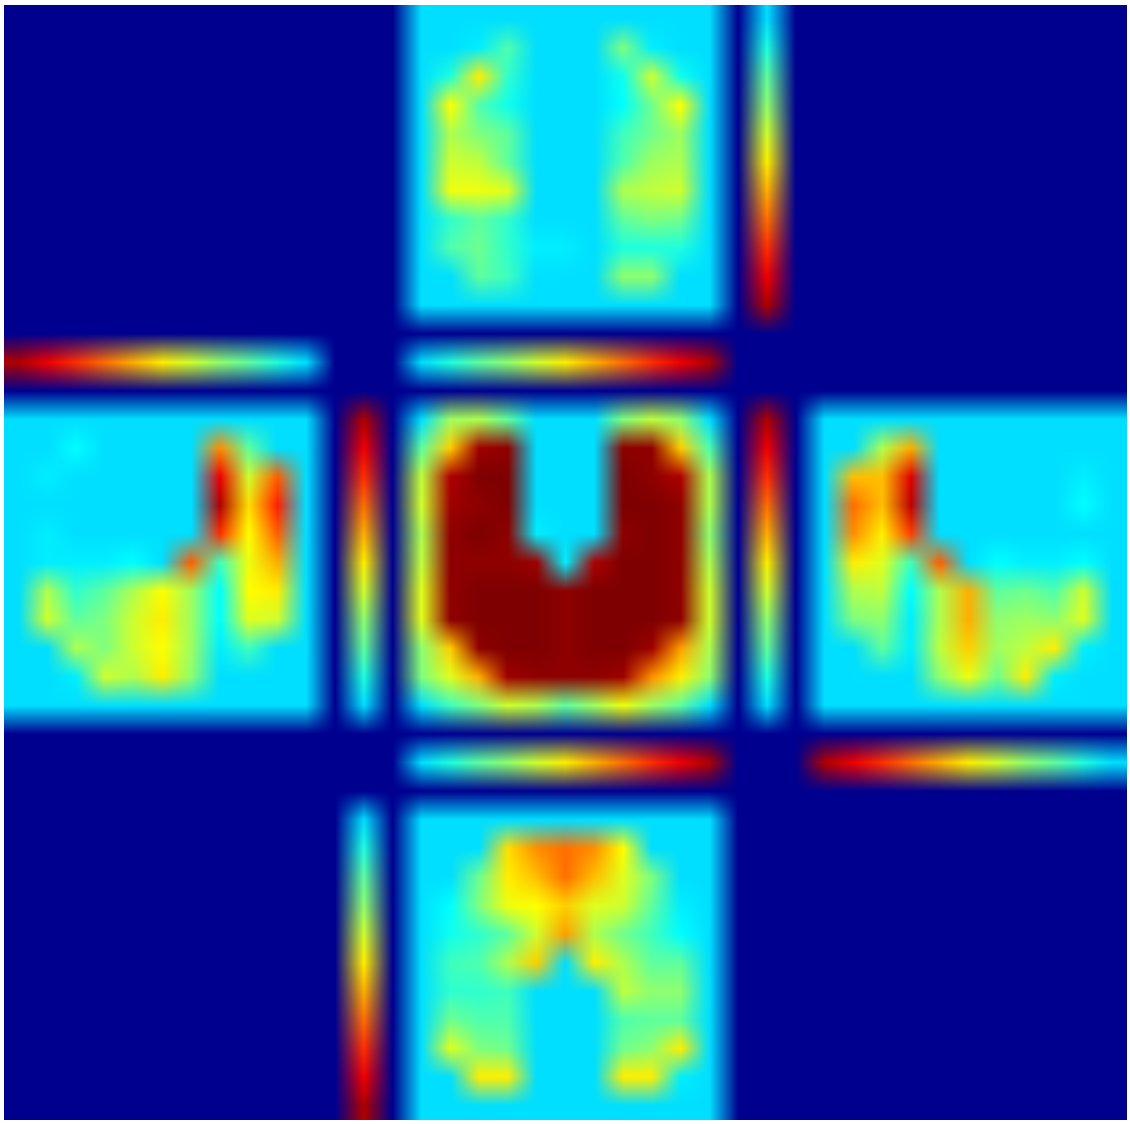

Figure 19. Relative frequency for beamlets to be used. The center shows the average dose distribution over 100 optimization runs

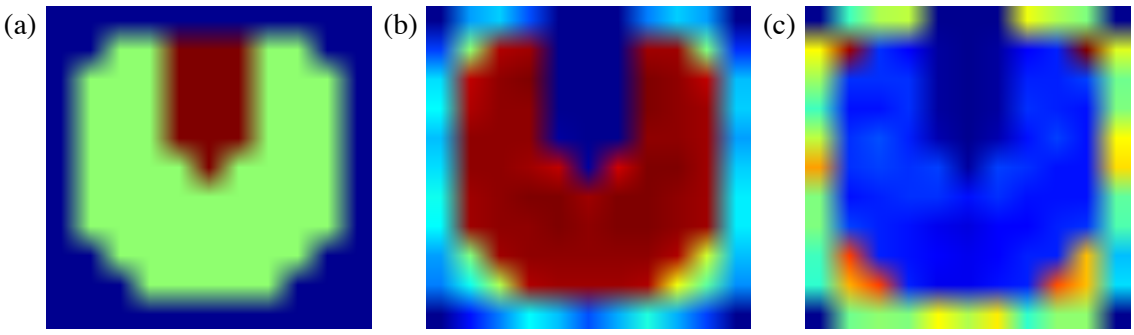

Figure 20. (a) Geometry of the tumor (green), the OAR (brown) and the unclassified tissue (blue) (b) Average dose (c) Standard deviation (both averaged over 100 optimization runs)

optimization would only compensate for the inability of the simulated annealing method to do so, either because there is no satisfying solution or because it is not found.

7. Deleting Apertures during the optimization

As mentioned in section 2, another approach was investigated in order to deal with the range as another degree of freedom for each aperture. In order to allow for a variable number of apertures per range, the optimization starts with several apertures per range. During the optimization some of the apertures are deleted until a specified number of apertures remains. The problem with this approach is to define a suitable criterion to decide which aperture should be deleted.
The criterion investigated here is based on the total dose contribution of an aperture to the CTV. Apertures that contibute a low dose to the CTV are considered less important and the aperture with the lowest contribution is deleted. Two modifications of this approach are investigated. The first approach considered the best aperture set after each annealing step, the second considered the average CTV dose contributions during the annealing step. The optimization runs are performed as follows:

1) The logarithmic cooling law was applied.
2) 50 cooling steps were performed, each consisting of 200 accepted leaf moves and the final temperature was 2% of the initial temperature.
3) The optimization started with 72 apertures in total, i.e. 2 per beam direction and range.
4) After each of the first 36 annealing steps, one aperture was deleted.
   method a) The best aperture set explored in this annealing step is considered. The aperture with the lowest CTV dose contribution is deleted and the weights are reoptimized. The reduced set of apertures is used as a starting position for the next annealing step.
   methob b) During each annealing step the average CTV dose contributions are calculated. After one annealing step, the aperture with the lowest average CTV dose contribution is deleted and the remaining apertures are used as starting position for the next annealing step.
5) The remaining 14 annealing steps were performed with a constant number of 36 final apertures.

Figure 21 shows the histogram of the obtained objective values in 100 optimization runs for both strategies. For comparison, the histogram in figure 15 is shown, where a fixed set of 36 apertures was used, i.e. one per range. Both strategies that delete apertures yield better objective values on average compared to the fixed aperture set. Method a) which considers the CTV dose contributions in the best aperture set instead of the average during the annealing step performed best. Figure 22 shows the best solution found. To some extent, the solution makes use of the additional degrees of freedom in an intuitive way. For example, in the beam coming from below, there is an aperture on

either side of the OAR that positions the bragg peak at the distal edge of the tumor.

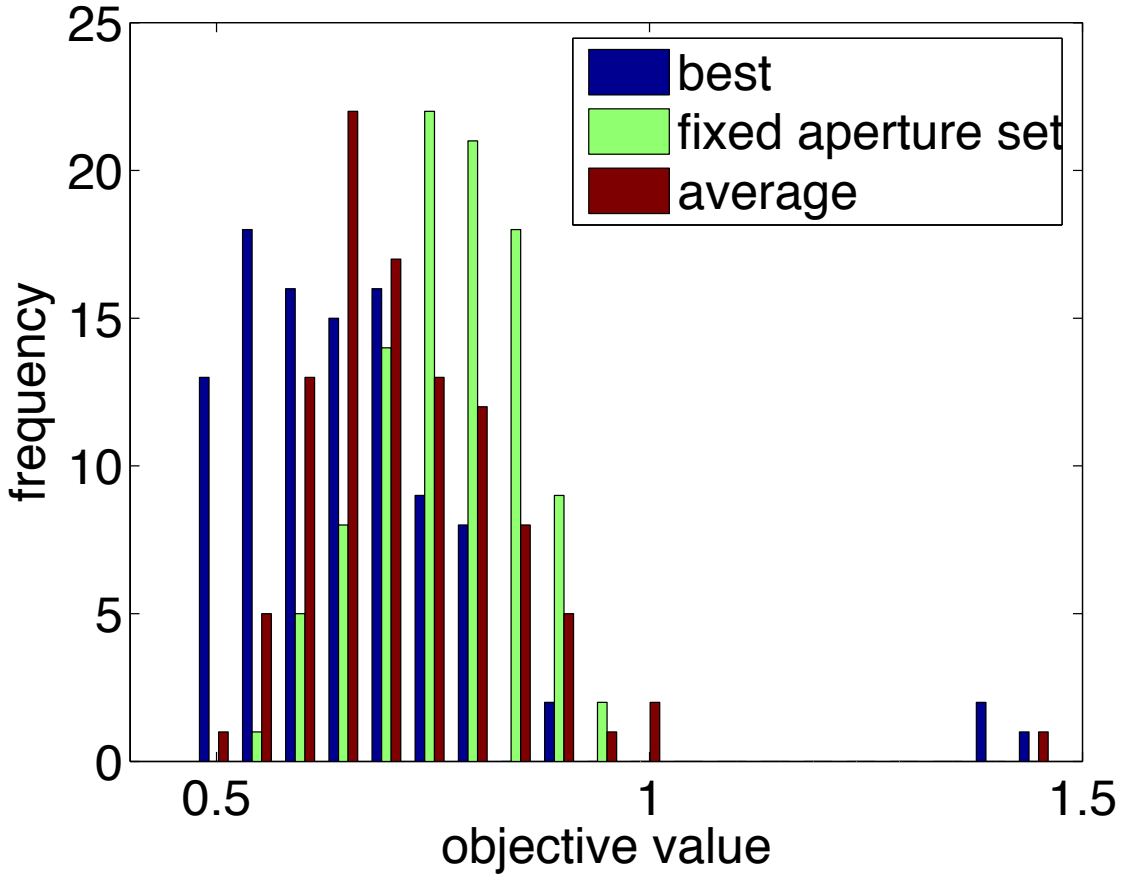

Figure 21. Histogram of objective function values in 100 optimization runs for different strategies: (best) deleting the aperture with lowest CTV dose in the best aperture set, (average) deleting the aperture with lowest average CTV dose during the annealing step

The histogram in figure 21 also shows that in some cases, the strategy of deleting apertures leads to worse results. Figure 23 shows the solution that corresponds to an objective value of 1.39. The reason for the high objective value is the underdosing of a CTV voxel adjacent to the OAR. In the initial set of 72 apertures, there are only 4 apertures that can irradiate this voxel without delivering a high dose to at least one OAR voxel. These 4 apertures are all deleted, thus making a good dose distribution impossible. If an aperture contributes a rather low dose to the CTV, it may still be important when its dose contribution can not be replaced by other apertures. On the other hand, an apertures that contribute a higher dose to the CTV may not be required urgently since eventually it can be replaced by other apertures.
Practically, the approach of deleting apertures during the optimization process may require a good set of initial apertures, i.e. the apertures should be reasonably distributed over the allowed ranges. Eventually, a good guess on how many apertures are needed for

a certain range may be obtained be considering the fluence maps of the beamlet-based optimized solution.

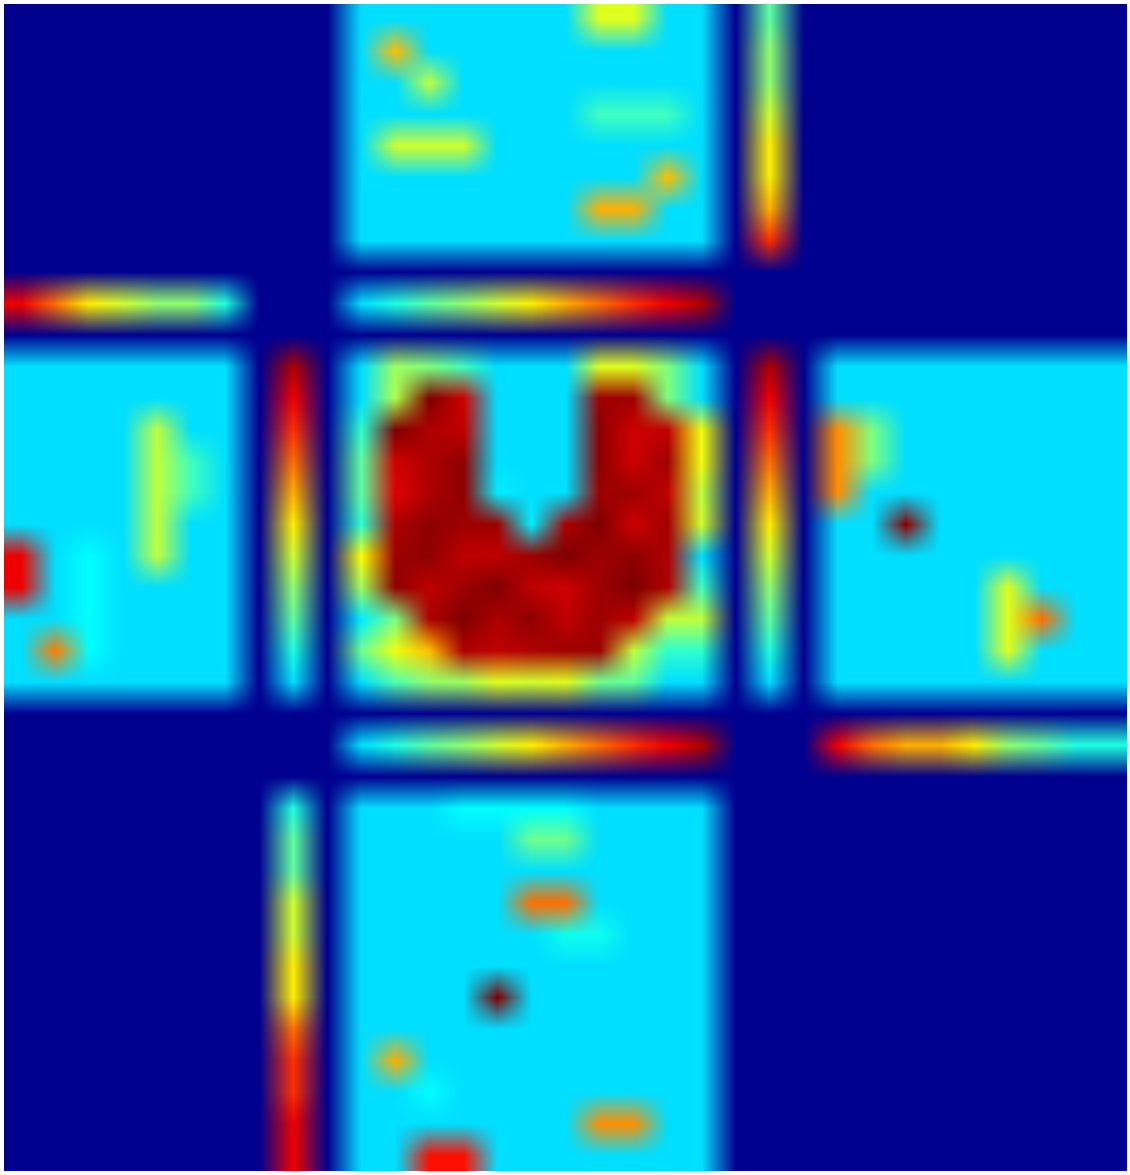

Figure 22. Best solution found after deleting 36 of 72 initial apertures during the simulated annealing process

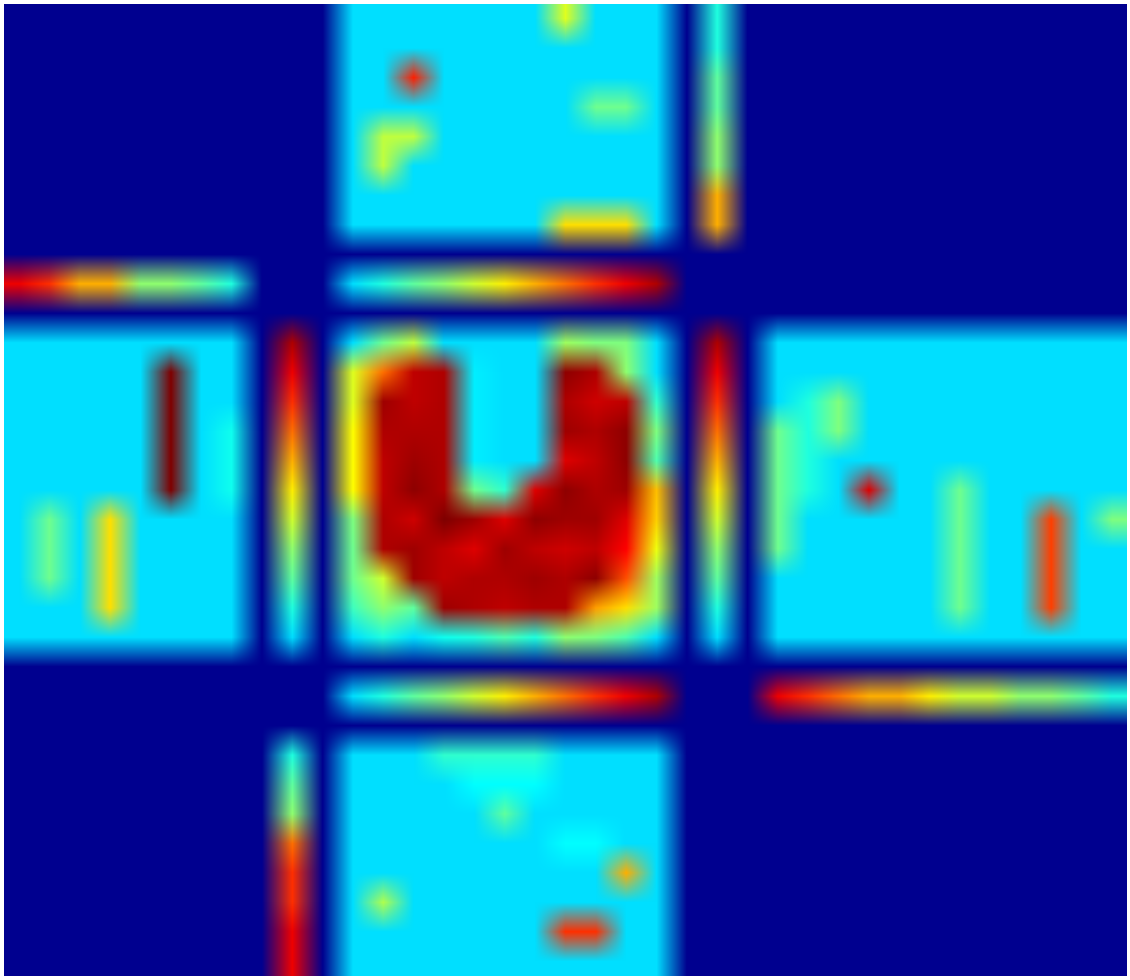

Figure 23. Bad solution resulting after deleting the wrong apertures during the simulated annealing process