

Treatment planning of intensity modulated proton therapy by using sparsity inducement for spot selection

Joris Mühlsteff

Master Thesis
Applied Mathematics



Treatment planning of intensity modulated proton therapy by using sparsity inducement for spot selection

by

Joris Mühlsteff

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Wednesday June 30, 2021 at 11:00 AM.

Student number: 4319966
Project duration: July 2, 2020 – June 30, 2021 (including 6 week break)
Thesis committee: Dr. ir. M. Keijzer, TU Delft, supervisor
Dr. ir. M. B. van Gijzen, TU Delft
Dr. ir. S. Breedveld, Erasmus University Medical Center, Rotterdam
Ir. M. Oud, Erasmus University Medical Center, Rotterdam

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

In intensity modulated proton therapy (IMPT), patients are irradiated with small spots, that deliver a local dose to the tumor. The number of possible spots to choose from is virtually infinite, but practically limited, which requires a spot selection. This spot selection should result in an optimal treatment plan, i.e., to deliver a sufficient dose to the tumor, while sparing the healthy surrounding tissue. These trade-offs make treatment planning in radiotherapy a multi-criteria optimization problem. The current approach for this spot selection by the Erasmus Medical Center (MC) is an iterative resampling method which uses a trial and error principal. A random sampled spot selection is made, bad spots are removed, and new spots are randomly added.

The research goal of this project is to improve the current spot selection method, by inducing sparsity on spot selections with the ℓ_1 -norm, without decreasing the plan quality of the current solution. This idea was already explored in Janssen [1], and considered viable. Sparse solutions are beneficial for optimization problems since they reduce the problem size and have higher probability of producing qualitative solutions. To achieve these goals, the *Sparsity-Induced-Spot-Selection* (SISS) method was developed. Contrary to the iterative resampling approach, the SISS method uses a top-down approach. Starting with a large spot coverage, it selects as little relevant spots as possible through the use of the ℓ_1 -norm to induce sparsity, until an acceptable treatment plan using as little spots as possible is made.

The developed method was validated on a test set consisting of 10 head and neck patients. Using the SISS method, an average spot selection of 1159 spots was produced, compared to a solution of 1074 spots for the resampling method. For the average patient, 6 out of 10 *Organs-at-risk* (OAR) received a lower dose with the SISS method than with the resampling method. The remaining OARs all received a marginal dose surplus of 0.6 Gy, with a maximum of 2.6 Gy. The target volumes in the tumor also received a similar dose to the resampling method, with the near-minimum dose of the tumor receiving a dose shortage of 0.2 Gy, and the near-maximum dose of the tumor receiving a dose surplus of 0.5 Gy, both being considered as marginal differences.

The SISS method produces a comparable spot selection and shows no decline in plan quality of the dose distributions. Using the ℓ_1 -norm to induce sparsity on spot selections in treatment planning is feasible. The computation time of the SISS method could be reduced using a voxel reduction, although this reduction in computation time is not guaranteed in practice due to robust optimization being applied.

Acknowledgements

This project was carried out under the supervision of dr. ir. Sebastiaan Breedveld, dr. ir. Marleen Keijzer and ir. Michelle Oud. I would first of all like to thank Sebastiaan Breedveld for offering me the opportunity to carry out this thesis project at the Erasmus MC, and for guiding me through the research. I would also like to thank Marleen Keijzer as a supervisor from the TU Delft, for supporting me throughout the project and providing useful insight during the weekly meetings. Without her connections at the Erasmus MC I would not have been able to do this research. I would also like to thank Michelle Oud for offering her help and sharing her knowledge during the research, which was very much appreciated. Finally, I would like to thank my friends and family for supporting me throughout this project. It has been a great experience at which I will gladly look back.

Contents

1	Introduction	1
1.1	Proton therapy	1
1.2	Treatment planning	2
1.3	Proton beam scanning	3
1.4	Problem description and proposal for a solution	3
1.4.1	Current method: IMPT through iterative resampling	3
1.4.2	Sparsity inducement using norms	4
1.5	Design of this study	4
1.6	Thesis outline	5
2	Sparsity inducement	7
2.1	ℓ_1 -norm	7
2.2	Geometric example	8
2.3	ℓ_p -norm for $0 \leq p < 1$	8
3	Multi-criteria optimization	11
3.1	Pareto optimality	11
3.2	Approaches for multi-criteria optimization	12
3.2.1	Weighted-sum method	12
3.2.2	ϵ -constraint method and $2pec$ method	13
3.2.3	Relation between $2pec$ and weighted-sum	13
4	Methodology	15
4.1	Overview of the SISS method	15
4.2	Iterative resampling method	16
4.2.1	Resampling solution	16
4.2.2	Conversion from $2pec$ to weighted-sum	16
4.3	ℓ_1 -norm addition	17
4.3.1	Norm weight	17
4.3.2	Threshold for removal of redundant spots	17
4.4	Deliverability settings	18
4.5	Alternative weight distributions	18
5	Results	21
5.1	ℓ_1 -norm addition	21
5.1.1	Norm weight	21
5.1.2	Threshold	24
5.1.3	Spot selection and dose distribution	25
5.2	Deliverability settings	26
5.3	Alternative weight distributions	28
5.4	Voxel reduction	30
6	Discussion	35
6.1	Norm weight	35
6.2	Threshold	35
6.3	Spot selection and dose distribution	35
6.4	Voxel reduction	36
6.5	Weight distribution	36
7	Conclusion	37

A Dose-Volume Histograms	39
A.1 Dose-Volume Histograms after $2p\epsilon$ optimization	39
A.2 Dose-Volume Histograms after Pareto projection	42
References	47

1

Introduction

In 2020, cancer was diagnosed in around 19.3 million patients worldwide, and radiotherapy was used in about 50% of the cases, sometimes in combination with chemotherapy or surgery [2]. In a radiotherapy treatment, the patient is irradiated with ionizing radiation to damage the malignant tumor cells and eradicate the tumor. Unfortunately, an inevitable side effect is the damaging of healthy cells around the tumor. This may lead to certain complications which can negatively impact the patients quality of life. On the other hand, there is the potential for cure, i.e., to irradiate the tumor with a sufficient dose and destroy it completely. This balancing between sparing healthy surrounding tissue and irradiating malignant tumor cells makes radiotherapy treatment planning a multi-criteria optimization problem, with the aim of a treatment that results in the highest quality of life for the patient. Since each patient is anatomically unique, this requires individually balanced plans.

1.1. Proton therapy

Radiation can be performed using different sources, including photon therapy, proton therapy and brachytherapy. The conventional method is photon therapy, which uses X-rays for irradiation. In this project however, we will focus on irradiation by proton therapy, which is widely considered to be the most technologically advanced method for irradiation currently available [3]. When a patient is irradiated with protons, beams of high energy protons are delivered into the tumor.

The main advantage of proton therapy over photon therapy is that the dose deposition is better suited for irradiating the tumor and sparing the healthy tissue of the patient. In conventional photon therapy, X-rays start with their peak-dose right at the beginning of their path inside the patient's body, while the deposited dose deeper in the patient decreases gradually. At the depth of the tumor, the intensity has decreased significantly, while the healthy tissue on the path to the tumor has received a high dose, thus damaging the tissue, as shown in Figure 1.1. Even when passing the tumor, the photons still deposit a dose (exit dose).

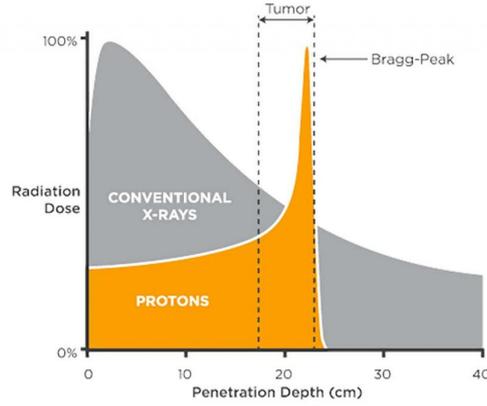


Figure 1.1: Bragg curve of protons and X-rays, showing the difference in location of the Bragg peak. Image taken from [4].

In proton therapy, the dose starts out low when entering the body and reaches its highest point, the Bragg peak, preferably at the location of the tumor. Behind the tumor, the dose quickly decreases to zero, thus having no exit dose. Varying the initial energy of the protons results in a different depth of the Bragg peak. Proton therapy thus allows a local dose deposition. One of the challenges of proton therapy addressed in this research is to determine where the volumes of high dose delivered by the protons should be placed in the patient to ensure sufficient dose to the entire tumor.

The local dose deposition causes another challenge in proton therapy: small anatomical variations or patient-setup variations on the treatment couch will cause the Bragg-peak to reach its maximum at a different location, resulting in underdosages in the tumor, and higher doses in healthy tissue. A method to mitigate this problem is robust treatment planning, where a desired dose is ensured under different predefined uncertainty conditions [5]. Due to the explorative nature of this project, plan robustness has not been included in this work. However, since robust treatment planning is captured in a canonical optimization problem formulation, it is expected that the findings of this work can be applied to robust treatment planning for further exploration.

1.2. Treatment planning

The treatment plan describes all the machine parameters required to deliver the dose distributions in an optimal way, i.e., to deliver a sufficient dose to the tumor while sparing the surrounding healthy tissue as much as possible. At the Erasmus Medical Center (MC), treatment planning is guided by a wishlist. The wishlist reflects the clinically desired doses in the volumes and describes the constraints and prioritized objectives, such as a minimum radiation doses for the *Clinical-Target-Volume* (CTV), and maximum radiation doses for *Organs-At-Risk* (OARs), as shown in Table 1.1. The resulting dose distribution must satisfy all constraints, while balancing the objectives in a clinically favorable way.

Table 1.1: Example of a wishlist with objectives and constraints, for a simplified head and neck case. The number of the objective indicates its priority in the optimization, with a lower number having the higher priority. For objective 1, the mean dose in the right parotid should be minimized with a limit of 26 Gy. Furthermore, the first constraint guarantees that the PTV (Planning-Target-Volume) receives no more than 49.22 Gy.

Objectives	Volume	Type	Limit (Gy)
1	Parotid (R)	minimize mean	26
2	Parotid (L)	minimize mean	26
Constraints	Volume	Type	Limit (Gy)
	PTV	maximum	49.22
	PTV	minimum	43.7

At this moment in clinical practice, the radiation therapist tweaks the parameter of planning objectives it-

eratively to obtain a treatment plan best suited for the unique anatomy of the patient. With this procedure, the final dose distribution depends generally on human factors. This does not guarantee optimality of the trade-offs made between the different objectives.

1.3. Proton beam scanning

A beam delivery technique in proton therapy is called *intensity modulated proton therapy* (IMPT), which uses protons for irradiation, with a non-uniform dose that can be adjusted to the size and location of the tumor. Treatment plans for IMPT are generated using Erasmus-iCycle, an in-house developed treatment planning system. A patient that undergoes IMPT is irradiated by multiple proton beams from different angles. The proton beam originates from a source and goes through a magnetic scanner, where the beam is diverted into several smaller beamlets, called *spots*, to a certain off-axis position. The depth depends on the energy of the beam, see Figure 1.2.

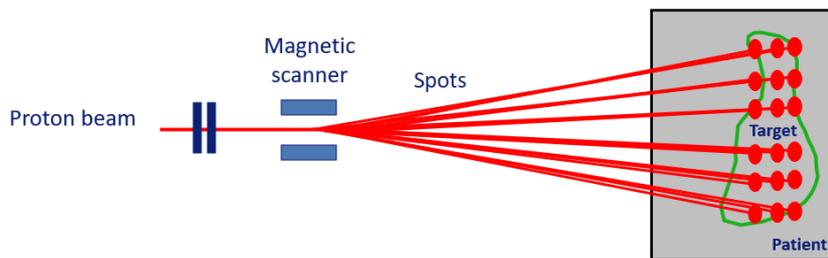


Figure 1.2: Schematic representation of proton beam scanning for IMPT. The patient's volume is discretized into 3D voxels for the optimization. Image taken and modified from Tony Lomax.

The intensity of each spot is expressed in monitor units (MU). The relation between the monitor unit of a spot and its intensity is linear, so a higher intensity results in a higher delivered dose.

1.4. Problem description and proposal for a solution

The number of possible spot positions is infinite, as spots may be placed virtually anywhere. Unfortunately, using a high spot density would make the optimization problem too large to solve. Furthermore, not all spots can be used due to practical limitations in the treatment device. Therefore, a spot selection is required to determine which spots to use to deliver a sufficient dose while sparing the healthy tissue as much as possible.

1.4.1. Current method: IMPT through iterative resampling

The current spot selection method uses a resampling approach [6]. It starts by randomly selecting 3000 spots out of millions of candidate spots. Then a treatment plan is created by converting the objectives and constraints of the patient to a weighted-sum problem, which is solved for the optimal MU per spot. Next, the spots below a certain threshold are removed and up to 3000 new spots are randomly added to the remaining ones, after which the process repeats until the last iteration, which is determined in advance to fit the patient. In the last iteration, spots under the minimum MU are removed and the solution is projected onto a Pareto front (see Section 3.1). This process is known as a Pareto projection.

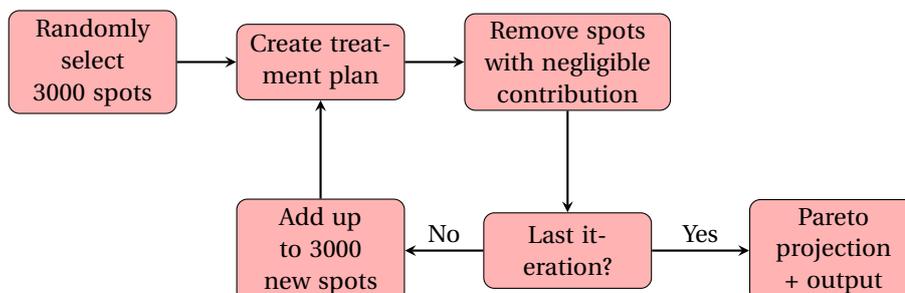


Figure 1.3: Schematic representation of the iterative resampling process of the current method used by the Erasmus MC.

The minimum MU is required by Holland Proton Therapy Center (HollandPTC), which is an outpatient center for proton therapy in Delft, to guarantee a stable beam during irradiation. The entire process is visualized in Figure 1.3.

By using this resampling approach, 24000 spots are tried, of which approximately 1000 spots end up in the final plan. In the experience of the Erasmus MC, approximately 1000 spots are enough to create a treatment plan with acceptable plan quality. There are however three main points of improvement to be made with this resampling approach:

1. The method for the selection of spots should be a more mathematically substantiated method that ideally results in a better spot selection. Random sampling and trial and error does not guarantee qualitative solutions.
2. The transparency of the current approach can be improved. There are a number of complex steps involved, which we want to avoid by using a more straightforward, intuitive approach to this problem.
3. The computational time takes too long. This iterative approach generates a treatment plan for every iteration, which is inefficient and could be time consuming when the optimization problems is large.

1.4.2. Sparsity inducement using norms

In this research project we aim to improve the spot selection approach by incorporating sparsity inducement using norms. The idea of using norms to induce sparsity was already explored in Janssen [1]. Janssen showed positive results for spot reduction by implementation of the ℓ_1 -norm, essentially a distance measuring function, on a small dataset and encouraged testing on a larger dataset.

A vector is considered sparse if it contains mostly zeros and a few non-zeros. When a solution has many zeros, there are fewer active spots (spots taken into the optimization). This would result in fewer variables in the optimization, which decreases the problem size, and also in a potential reduction in treatment time. Given these benefits, we wish to develop a new approach for the spot selection, which uses the ℓ_1 -norm to produce sparse solutions. In Chapter 2 it is further explained why the ℓ_1 -norm induces sparsity on solutions. The incorporation of the ℓ_1 -norm in the spot selection led to the development of the *Sparsity-Induced-Spot-Selection* (SISS) method. A concise overview of the SISS method is shown in Figure 1.4.

1. Take the spots that were randomly sampled in the resampling method.
2. Select a relevant subset of approximately 1000 spots.
3. Create a treatment plan based on this subset of spots.

Figure 1.4: Concise overview of the proposal for the SISS method.

Step 1 differs from the original resampling approach. Instead of iteratively adding up to 3000 spots and removing bad spots, the SISS method uses a top-down strategy. It starts with a large spot coverage, i.e., approximately 24000 spots in our dataset, and selects as few spots as possible to find a good subset of spots. The selection of this relevant subset in step 2 is the kernel of the SISS method, and brings the concept of sparsity inducement to the problem. After this selection, step 3 converts the solution to an executable plan for the radiation therapist. The overview of the SISS method in Fig. 1.4 is expanded and further detailed in Chapter 4.

1.5. Design of this study

The general approach of this research is to implement the SISS method and compare the results of the spot selection to that of the original resampling solution, using a dataset of 10 patients, each head and neck cases. The implementation of the SISS method is done in Matlab, while the solver for the optimizations runs in C++. To evaluate whether the spot selection produces a clinically acceptable treatment plan, dose distributions of both the SISS method and the resampling method are made and compared with each other.

The dose distributions are represented in a cumulative 2D *Dose-Volume Histogram* (DVH), a commonly used tool to visualize the 3D dose distributions. An example is shown in Fig. 1.5

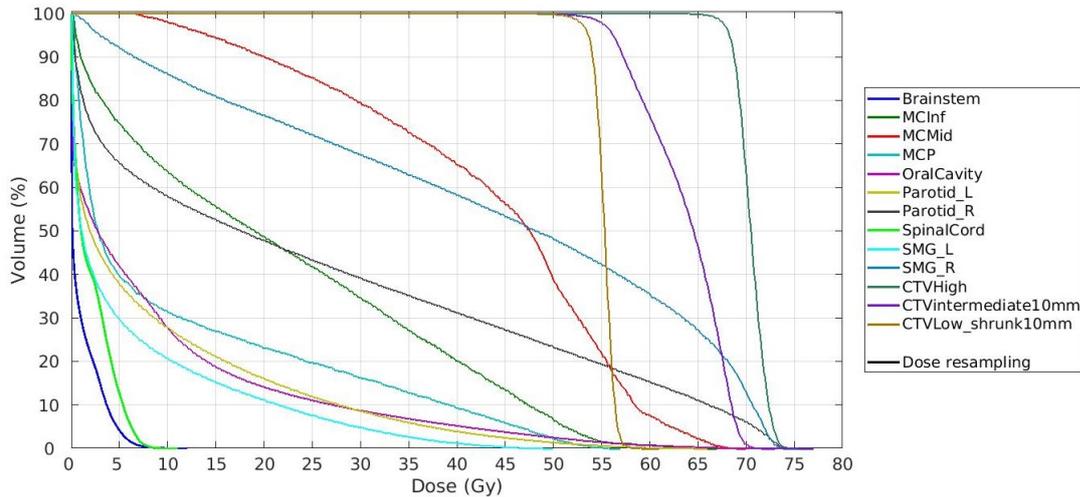


Figure 1.5: Example of a Dose-Volume Histogram (DVH) of a head and neck case.

The horizontal axis represents the dose delivered in Gray (Gy), and the vertical axis represents the relative volume of a structure that receives at least a certain dose level. For example, 80% of the volume of the MCMid gets a dose of 30 Gy, while only 40% of the volume receives 50 Gy. In general the CTV volumes receive a higher dose, as these volumes comprise the tumor, and the other volumes (OARs) receive a lower dose. Note that the average patient has 10 to 20 volumes in and around the tumor, so showing the dose distributions of all volumes would decrease the readability of the DVH. Therefore a selection is made of relevant volumes, which are related to the benchmark parameters in Table 1.2.

To understand these benchmark parameters, we must first define serial and parallel OARs. An organ may be divided into sub-components. A *serial* OAR is an organ-at-risk that fails if one of its components is disabled. Examples include the spinal cord and the brain stem. A *parallel* OAR is an organ-at-risk where all the sub-components must be disabled for the organ to fail. Examples include the lungs, liver and kidneys. Because of this high risk of failure, the threshold for a surplus dose in a serial OAR is low.

Table 1.2: Overview of four different benchmark parameters used to evaluate the quality of dose distributions.

Dose	Explanation	Applicable to	Preferable dose
D0.03cc	Dose that 0.03cc of the volume receives.	Serial OARs	Low
Dmean	Mean dose over the entire structure.	Parallel OARs	Low
D2%	Dose that 2% of the volume of the structure receives. This is the near-maximum dose.	CTV	High
D98%	Dose that 98% of the volume of the structure receives. This is the near-minimum dose.	CTV	High

The DVHs and benchmark parameters are used when evaluating the quality of dose distributions. The DVHs give a general overview of the similarity of the dose distribution to the resampling solution, while the benchmark parameters highlight important features that are not immediately evident from the DVH, such as the D2% dose.

1.6. Thesis outline

In Chapter 2, it is explained why and how the ℓ_1 -norm induces sparsity. A geometrical example is provided as well as a demonstration of the effect of other ℓ_p -norms on sparsity of solutions. In Chapter 3, the different multi-criteria optimization techniques used in this research are explained. In Chapter 4 the SISS method is explained in more detail, by introducing different input parameters and additional deliverability settings. Chapter 5 contains the results and analysis of the previously described methods, as well as a potential method

for reducing computation time, by reducing the number of voxels in the patients volume. In Chapter 6 we provide a discussion on the results, and the thesis is finalized with a conclusion in Chapter 7.

2

Sparsity inducement

In this chapter, the functioning of the ℓ_1 -norm and its involvement in sparsity inducement are discussed in more detail. A geometric example is provided as well as a look at other norms for sparsity inducement.

2.1. ℓ_1 -norm

The resampling solution produces a vector with the MU per spot. A way to reduce the delivery time of this vector is by making it sparser, thus having it contain more zeros. To do so, we will use the ℓ_p -norm. The ℓ_p -norm is defined as follows:

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{1/p}. \quad (2.1.1)$$

This project will focus on the ℓ_1 -norm, which is defined as

$$\|\mathbf{x}\|_1 = \sum_{i=1}^n |x_i|. \quad (2.1.2)$$

To explain why the ℓ_1 -norm induces sparsity on the spot vector, we show the difference between an ℓ_1 -norm penalty and an ℓ_2 -norm penalty, and explain why the latter norm does not induce sparsity. We take a simple vector $\mathbf{x} = (x_1, x_2) = (1, \epsilon) \in \mathbb{R}^2$, where $\epsilon > 0$ is small. The ℓ_1 -norm and ℓ_2 -norm on this vector are respectively:

$$\|\mathbf{x}\|_1 = 1 + \epsilon, \quad (2.1.3)$$

$$\|\mathbf{x}\|_2 = \sqrt{1 + \epsilon^2}. \quad (2.1.4)$$

Now we penalize both terms of \mathbf{x} with $\delta \leq \epsilon$ and compare the effects for the ℓ_1 -norm and the ℓ_2 -norm:

Table 2.1: Example of the effect of the ℓ_1 -norm and the ℓ_2 -norm on magnitude reduction.

	ℓ_1 -norm	ℓ_2 -norm
penalize x_1	$1 - \delta + \epsilon$	$\sqrt{1 - 2\delta + \delta^2 + \epsilon^2}$
penalize x_2	$1 + \epsilon - \delta$	$\sqrt{1 + \epsilon^2 - 2\delta\epsilon + \delta^2}$

Note that with the ℓ_2 -norm, penalizing the smaller term x_2 results in a smaller reduction in norm than doing so to the larger term x_1 . Thus when penalizing a vector with the ℓ_2 -norm, smaller elements tend not to go

to 0, as the reduction from ϵ to 0 becomes smaller as ϵ gets smaller. For the ℓ_1 -norm however, the reduction is always δ , regardless of the magnitude of the element. For this reason we will look into the ℓ_1 -norm for sparsity inducement.

2.2. Geometric example

A simple geometric example is presented to show the effect of the ℓ_1 -norm, when compared to the ℓ_2 -norm. Take a simple linear program in 2D, as seen in Equation (2.2.1), where the objective is to minimize some vector $\mathbf{f} = (x_1, x_2)$:

$$\begin{aligned} \text{minimize } f &= (x_1, x_2) \\ x_1, x_2 &\in \mathbb{R} \end{aligned} \tag{2.2.1}$$

We start by adding the familiar ℓ_2 -norm to the objective function. Now we have to find the point $\mathbf{x}^* = (x_1^*, x_2^*) \in \mathbb{R}$ for which $\|\mathbf{x}^*\| < \|\mathbf{x}\| \quad \forall \mathbf{x} \in \mathbb{R}^2$. This can be visualized as increasing the value of the ℓ_2 -norm, which is essentially increasing a circle around the origin, until it touches the vector \mathbf{f} , as seen in Figure 2.1a. Applying the same procedure to the ℓ_1 -norm results in increasing a diamond around the origin, as seen in Figure 2.1b.

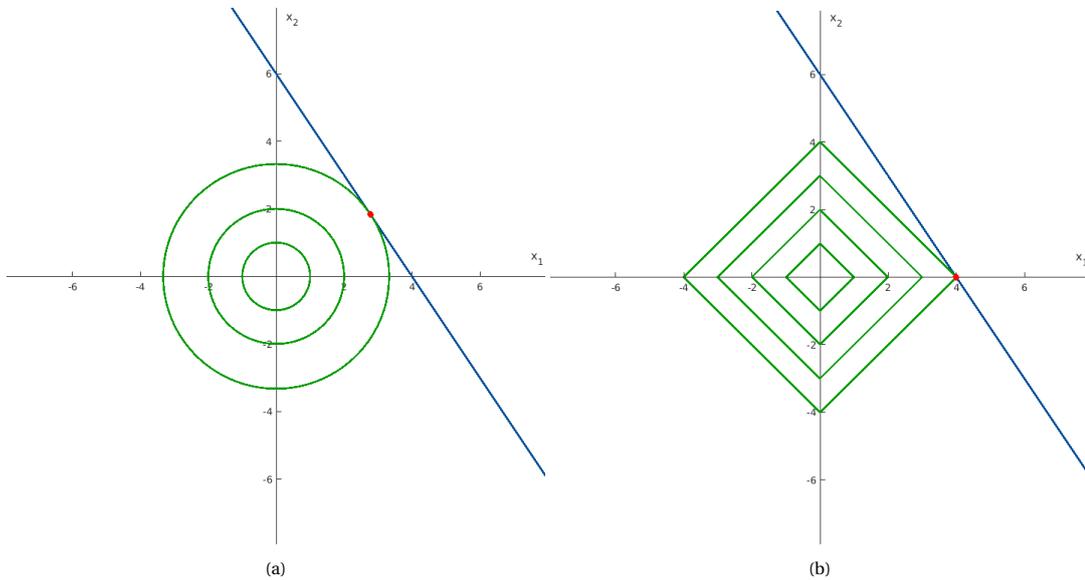


Figure 2.1: Visualization of adding increasing ℓ_p -norms to the objective function for $p = 2$ (a) and $p = 1$ (b). The increasing ℓ_2 -norm takes the shape of increasing circles, while the increasing ℓ_1 -norm takes the shape of increasing diamond. In both figures, the blue line is the objective function, given by $f(x_1)$, and the red dot is the intersection of the objective function with the minimum ℓ_p -norm.

The value of \mathbf{f} which has minimum ℓ_p -norm is indicated in both cases by the red dot. This shows the difference between the two norms, as the solution in Figure 2.1a is not a sparse solution, since it does not lie on one of the axes, but the solution in Fig. 2.1b does, since it lies on the x -axis. It should be noted that the ℓ_1 -norm does not always guarantee sparse solutions, as the slope of the vector \mathbf{f} could be similar to the slope of one of the edges of the ℓ_1 -norm shape, but it does *increase* the chance of a sparse solution.

2.3. ℓ_p -norm for $0 \leq p < 1$

While Figure 2.1 shows that the shape of the ℓ_1 -norm induces more sparsity than the ℓ_2 -norm, the shape of the ℓ_p -norm for $0 \leq p < 1$ seems to induce even more sparsity, as seen in Figure 2.2. Ideally we would want to use the ℓ_0 -norm, as all of its points lie on the axes, thus making it a highly sparse solution. However, for $0 \leq p < 1$, the ℓ_p -norm does not meet the requirements of the mathematical definition of a norm. For a function $\|\cdot\|$ to be considered a norm, it must satisfy three properties:

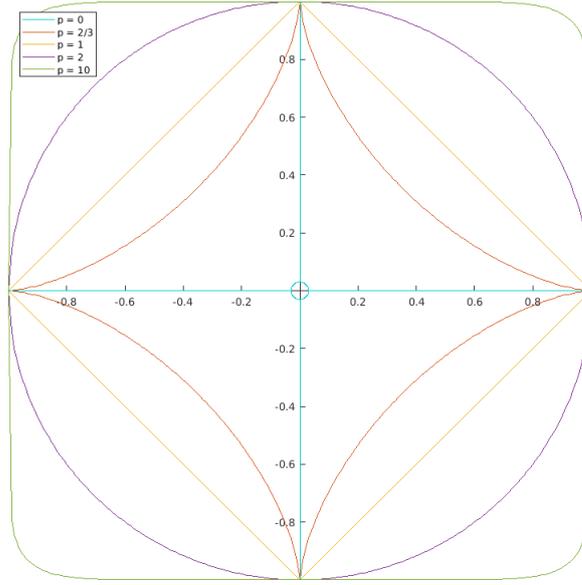


Figure 2.2: Geometric representation the ℓ_p -norm for $p = 0, 2/3, 1, 2$ and 10 . The shape of the norm expands from a non convex plus sign to a convex square with rounded corners. The ℓ_0 -norm has a small circle around the origin, since $(0,0)$ is the only point on the axes not included in the norm.

1. (Positive definite): If $\|x\| = 0$, then $x = \mathbf{0}$.
2. (Homogeneity): $\forall x \in \mathbb{R}^n, k \in \mathbb{R} : \|kx\| = |k|\|x\|$.
3. (Triangle inequality): $\forall x, y \in \mathbb{R}^n : \|x + y\| \leq \|x\| + \|y\|$.

If $p = 0$, then the ℓ_0 -norm seems undefined in Equation (2.1.1), but it becomes the number of non-zero elements in a vector. Furthermore, the ℓ_0 -norm does not meet the homogeneity property. Take $x_1 = 1, x_i = 0$ for $i \neq 1$ and $k = 2$. Then

$$\|kx\|_0 = \|2 * 1\|_0 = 1 \neq 2 = |2|\|1\|_0 = |k|\|x\| \quad (2.3.1)$$

If $0 < p < 1$, then the ℓ_p -norm is well defined by Equation (2.1.1), but the triangle inequality does not hold anymore. Take for example $x_1 = y_1 = 2$ and $x_i = y_i = 0$ for $i \neq 1$. Then

$$\|x\|_p + \|y\|_p = 2 + 2 = 4 \quad (2.3.2)$$

But

$$\|x + y\|_p = (2^p + 2^p)^{1/p} = (2^{p+1})^{1/p} = 2 \cdot 2^{1/p} > 4 \quad (2.3.3)$$

Besides the fact that the ℓ_0 -norm does not meet the mathematical definition, there are also practical reasons for favoring the ℓ_1 -norm over the ℓ_0 -norm or any ℓ_p -norm with $0 < p < 1$. Note that if $0 < p < 1$, the ℓ_p -norm becomes a convex constraint, which is known to cause computational difficulties [7]. Furthermore, the optimization problem with ℓ_0 -norm is known to be NP-hard [8]. Finally, it has been proven that the ℓ_1 -norm is the best convex approximation to the ℓ_0 -norm [7]. For these reasons, we use the ℓ_1 -norm, as it is more likely to generate sparse solutions than the ℓ_2 -norm, and does not have the computational difficulties of the ℓ_0 -norm or any ℓ_p -norm with $0 < p < 1$.

3

Multi-criteria optimization

In treatment planning, there are usually multiple objectives to be achieved: delivering a sufficient dose to certain volumes, and sparing healthy surrounding tissue. This makes radiotherapy treatment planning a multi-criteria problem. In this chapter, multi-criteria optimization is introduced by explaining the optimization methods used in the current resampling method, as well as an important feature in radiotherapy treatment planning.

3.1. Pareto optimality

When dealing with optimizing multiple objectives, the ideal situation occurs when one objective cannot be improved without worsening another objective. Consider an optimization problem with objectives f_1, \dots, f_n . Let $f^*(x) = f_1^*, \dots, f_n^*$ be a solution of the problem. Then we say that $f^*(x)$ is *Pareto-optimal* if there is no other feasible solution that improves the value of one objective f_i^* without worsening another objective f_j^* , where $i \neq j$. The set of all Pareto optimal solutions is referred to as the *Pareto-front*. To illustrate the Pareto front, we consider an example of two objective functions, $f_1(x) = 1 + x^2$ and $f_2(x) = 3 + (x - 1)^2$ that need to be minimized:

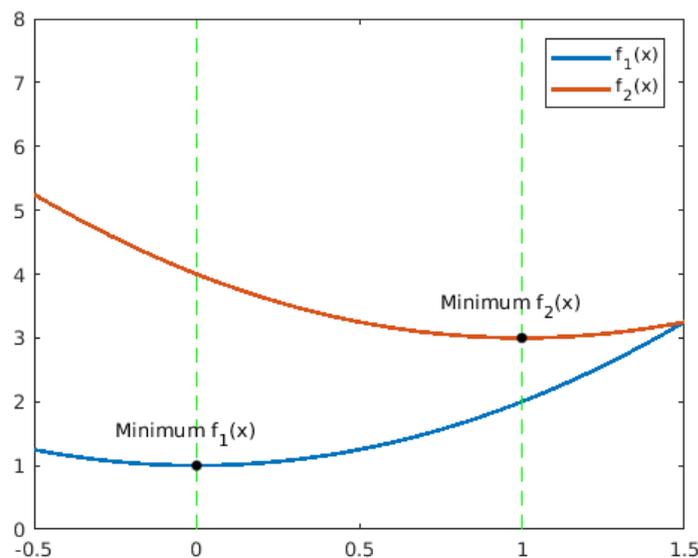


Figure 3.1: Plot of functions $f_1(x) = 1 + x^2$ and $f_2(x) = 3 + (x - 1)^2$, with their global minima. The region between the green lines is called the *trade-off region*, where the optimal value of the minimization problem is to be found by making trade-offs between both functions.

Both functions decrease for $x \leq 0$ and increase for $x \geq 1$, but for $0 < x < 1$, $f_1(x)$ increases while $f_2(x)$ decreases. This is the region where trade-offs between $f_1(x)$ and $f_2(x)$ can be made. Now a Pareto front can be generated using the weighted-sum method, with weights $w_1 = t$ and $w_2 = 1 - t$ for $t \in [0, 1]$, as shown in Fig. 3.2.

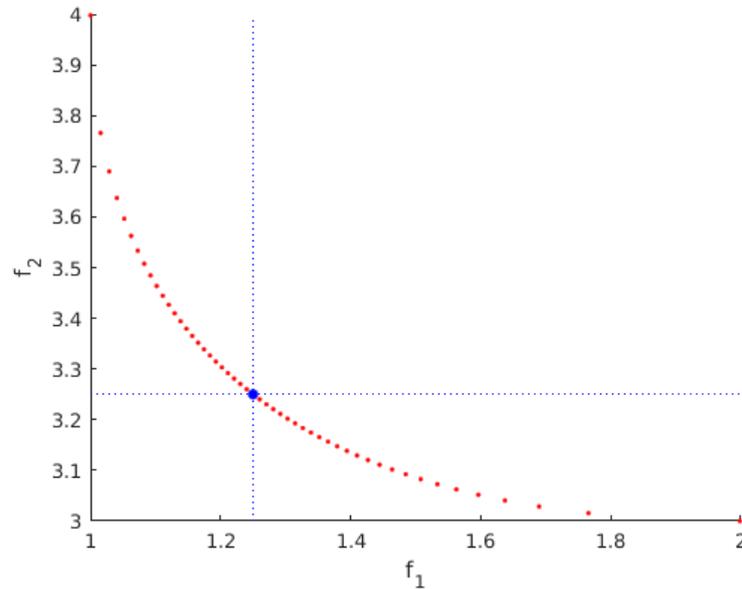


Figure 3.2: Pareto front for the minimization of $f_1(x)$ and $f_2(x)$. Every red dot represents a solution for a different $t \in [0, 1]$. The blue point at the intersection represents the optimal solution $(f_1, f_2) = (1\frac{1}{4}, 3\frac{1}{4})$ for $t = 0.5$.

The generated Pareto front in Figure 3.2 shows the optimal solution to be $(f_1, f_2) = (1\frac{1}{4}, 3\frac{1}{4})$ for $t = 0.5$.

In radiotherapy, Pareto-optimality translates to not being able to deliver more dose to a tumor without damaging more healthy tissue, or reducing the dose in one organ without increasing the dose to others. The goal in radiotherapy is to find a solution that is Pareto optimal. The current planning system optimization with Erasmus-iCycle is based on a prioritized wishlist, which contains all the targets. Erasmus-iCycle optimizes objectives from a wish-list by priority. If the highest priority objective meets its hard constraint, the next objective is optimized, until none of the objectives can be optimized any further. In this way, the solution generated by Erasmus-iCycle is always Pareto optimal.

3.2. Approaches for multi-criteria optimization

The current method for IMPT at the Erasmus MC uses multiple optimization methods for multi-criteria problems. In this section we introduce the weighted-sum method, the ϵ -constraint method, and its extension into the *2pec* method.

3.2.1. Weighted-sum method

One of the most straightforward methods for solving multi criteria problems is the weighted-sum method. It operates by assigning a weight w_i to the objective functions $f_i(x)$ for $i = 1 \dots n$, and summing them together so that the optimization problem becomes:

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n w_i f_i(x) \\ &\text{subject to} && \mathbf{g}(x) \leq \mathbf{0} \end{aligned} \tag{3.2.1}$$

where $\mathbf{g}(x)$ are hard constraints which are to be met at all times.

3.2.2. ϵ -constraint method and $2pec$ method

The ϵ -constraint method works with a priority-based optimization, where one objective at the time is optimized, while keeping the other objectives constrained. This way each objective is optimized only once. The ϵ -constraint method can be extended into the 2 -phase- ϵ -constraint ($2pec$) method.

The first phase of the $2pec$ method consists of assigning a goal b_i to each objective $f_i(x)$ and minimizing each objective separately to its goal, while keeping higher prioritized objectives constrained. If it is possible to minimize an objective $f_i(x)$ below its goal b_i , it is more desirable to minimize another (lower prioritized) objective $f_j(x)$ to its goal b_j than to minimize f_j to its fullest extend, because setting the constraint for $f_i(x)$ higher than its goal b_i increases the chance of sparing the lower prioritized objective f_j [9]. The i 'th step of the first phase of the optimization problem then becomes:

$$\begin{aligned} & \text{minimize} && f_i(x) \\ & \text{subject to} && \mathbf{g}(x) \leq \mathbf{0} \\ & && f_k(x) \leq \epsilon_k \quad \text{for } k \in \{1, \dots, i-1\} \end{aligned} \quad (3.2.2)$$

where ϵ_k is a bound that is chosen according to the following rule:

$$\epsilon_k = \begin{cases} b_k & f_k(x^*)\delta < b_k \\ f_k(x^*)\delta & f_k(x^*)\delta \geq b_k \end{cases} \quad (3.2.3)$$

where δ is a slight relaxation to prevent numerical problems and x^* is the result of the previous iteration. Note that in the first iteration of Equation (3.2.2), there is no other $f_k(x)$ that is bounded as a constraint.

In the second phase of the $2pec$ method, all objectives which met their goals are minimized to their fullest, while keeping all other objectives constrained:

$$\begin{aligned} & \text{minimize} && f_i(x) \\ & \text{subject to} && \mathbf{g}(x) \leq \mathbf{0} \\ & && f_k(x) \leq \epsilon_k \quad \text{for } k \in \{1, \dots, n\} \setminus i \end{aligned} \quad (3.2.4)$$

Note that the second phase of the $2pec$ method resembles the ϵ -constraint method, and therefore their solutions have the same properties [9].

3.2.3. Relation between $2pec$ and weighted-sum

It is possible to switch between the $2pec$ method and the weighted-sum method. The conversion from the weighted-sum method to the $2pec$ method is not relevant for this research, therefore only the conversion from the $2pec$ method to the weighted-sum method is considered. For this conversion, it is required to construct weights for the weighted-sum optimization. To obtain these weights, the problem in Equation (3.2.2) is rewritten as an unconstrained problem, called the *Lagrangian*:

$$L(x, \nu, \lambda) = f_n(x) + \sum_{i=1}^{n-1} \nu_i (f_i(x) - \epsilon_i) + \sum_{j=1}^m \lambda_j g_j(x), \quad (3.2.5)$$

which is to be minimized with respect to x, ν and λ , where ν and λ are nonnegative vectors of *Lagrange multipliers*. Now we can take these Lagrange multipliers for the constrained objectives from the last iteration of the $2pec$ optimization, and use them as the weights for the weighted-sum problem. This approach results in an identical optimal solution [9]. To get to the optimal triplet (x^*, ν^*, λ^*) , a primal-dual interior-point method is used, which is further explained in Chapter 4.

4

Methodology

In this section our proposed SISS method is further explained, as well as additional research in support of the SISS method. The research in this project is essentially a continuation of Janssen [1]. Janssen showed positive results for spot reduction by implementation of the ℓ_1 -norm on a small dataset and encouraged testing on a larger dataset.

The initial research consisted of reproducing the results from Janssen [1], to get a better understanding of the problem. The dataset used by Janssen consisted of a single prostate cancer patient with a resampling solution of 1705 spots. When the reproduction of the results turned out succesfull, the same approach was applied to a larger dataset, which consisted of 10 patients (each a head and neck case). This dataset contains actual, anonymized patients that were treated at the Erasmus MC, and is used throughout the project to validate the SISS method.

4.1. Overview of the SISS method

As explained in Section 1.4.2, the SISS method starts with a large spot coverage, randomly spread over the target, from which a relevant subset is selected, upon which a treatment plan is based. This concise description of the method is further detailed in Figure 4.1.

1. Take a large, randomly sampled spot coverage (approximately 24000 spots).
2. Convert the resampling problem to a weighted-sum problem.
3. Add the ℓ_1 -norm to weighted-sum problem.
4. Solve the weighted-sum problem for the optimal MU per spot.
5. Remove spots below a certain threshold.
6. Re-optimize spots with a $2p\epsilon c$ optimization.
7. Create a treatment plan with the resulting spot selection.

Figure 4.1: Overview of the SISS method.

Steps 1 and 2 detail the setup of the method and conversion of the problem to a practical form. Steps 3 and 4 introduce the main idea of sparsity inducement to this form and result in the first solution. Step 5 serves as an improvement of the solution found in Step 4, while steps 6 and 7 re-optimize the solution and convert it to a clinical irradiation plan.

4.2. Iterative resampling method

4.2.1. Resampling solution

As mentioned in Section 1.4.1, the current method for a spot selection uses an iterative resampling approach, which adds spots the solution on a trial and error base and removes bad spots from the solution. Using this approach, a spot coverage has been generated for the 10 test patients, as seen in Table 4.1.

Table 4.1: Number of spots evaluated and number of selected spots in the final solution of the resampling approach for 10 test patients.

Patient	Spot coverage	Resampling solution
1	23894	1289
2	23881	1064
3	23909	1225
4	23879	990
5	23871	961
6	23901	1185
7	23917	1306
8	23888	1060
9	23857	873
10	23832	793
Average	23882	1074

The spot coverage contains close to 24000 spots on average per patient. From this spot coverage, only the spots with a minimum MU of 1.33 have been selected in the solution, to guarantee a stable beam during irradiation at HollandPTC. This selection results in a resampling solution of 1074 spots on average per patient. Results of previous research [6, 10, 11] indicate that a spot selection of approximately 1000 spots is sufficient for a proper dose distribution. With the use of the SISS method, the goal is to achieve a spot selection with a similar or lower number of spots on average, while maintaining or improving the plan quality of the dose distributions.

4.2.2. Conversion from $2pec$ to weighted-sum

The results from the resampling method serves as the basis for the SISS method. We take the uniformly sampled spot coverage of the tumor, from Table 4.1. This spot coverage, which includes the spots from the resampling solution, is the framework from which we will construct our spot selection. To be able to perform an optimization with the ℓ_1 -norm, we first have to convert the resampling problem, which is a $2pec$ optimization problem, to a weighted-sum optimization problem, since performing a $2pec$ optimization with 24000 spots would take approximately 10 hours to compute, thus making it highly impractical. The weighted-sum optimization problem is of the form

$$\begin{aligned}
 & \text{minimize} && \sum_{i=1}^n w_i f_i(x) \\
 & \text{subject to} && \mathbf{g}(x) \leq \mathbf{0}
 \end{aligned} \tag{4.2.1}$$

where the $f_i(x)$'s are the objectives, such as minimizing the dose in the left parotid, and w_i are the accompanying weights. The objectives are taken from a wishlist, as seen in Table 1.1, and the weights are obtained from solving the Lagrangian in Equation (3.2.5) for the triplet (x^*, ν^*, λ^*) by using a *primal-dual interior-point method*.

An interior point method is a type of algorithm that is used to solve linear and nonlinear convex optimization problems. It takes a standard optimization problem of the form:

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{subject to} && \mathbf{g}(x) \leq \mathbf{0} \\ & && x \geq 0 \end{aligned} \tag{4.2.2}$$

and converts it to:

$$\begin{aligned} & \text{minimize} && f(x) - \mu \sum_{i=1}^n \ln(x_i) \\ & \text{subject to} && \mathbf{g}(x) \leq \mathbf{0} \end{aligned} \tag{4.2.3}$$

Equation (4.2.3) prevents violation of the constraints by adding a barrier term $\ln(x_i)$ to the objective function. The method then iterates through the search space by decreasing the value of the parameter μ , which pushes the solution to its barrier value. The primal-dual aspect of the name refers to the method also updating dual variables.

The primal-dual interior-point method guarantees optimal and feasible solutions if the problem is well-defined, i.e., bounded and feasible [12]. The solution of the primal-dual interior-point method represents the optimal MU per spot, as mentioned in step 4 of Figure 4.1.

4.3. ℓ_1 -norm addition

Now that we have set up the optimization problem in a proper form, we add the ℓ_1 -norm to the problem, as well as introducing two parameters to the problem: norm weight α and threshold θ .

4.3.1. Norm weight

After obtaining the weights from the Lagrangian, we add the ℓ_1 -norm to the objective function with weight α . This way the problem in Eq. (4.2.1) becomes

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^n w_i f_i(x) + \alpha \|\mathbf{x}\|_1 \\ & \text{subject to} && \mathbf{g}(\mathbf{x}) \leq \mathbf{0} \end{aligned} \tag{4.3.1}$$

where the parameter α represents the weight of the ℓ_1 -norm. A higher value of α will result in a heavier penalty by the ℓ_1 -norm, which results in a sparser solution, and thus in a bigger spot reduction than a lower value of α . The value of α is determined by comparing the effect of different norm weights on the plan quality of the dose distribution. Before we are able to compare these different norm weights, we first look at the second parameter of the problem, as it influences the choice for a norm weight as well.

4.3.2. Threshold for removal of redundant spots

The optimization problem in Equation (4.3.1) has the spot coverage from Table 4.1 as its input. The solution of Equation (4.3.1) is therefore a vector with the size of the spot coverage in Table 4.1, approximately 24000. As mentioned in Section 4.2.2, this number of spots is too high for a *2pec* optimization. Therefore a threshold θ is imposed on the solution, which effectively has two consequences:

1. *Removal of redundant spots* - The magnitude of the MU of the spots ranges from 10^{-8} to 10^1 . Since the magnitude of the MU determines the intensity of the pencil beam, the difference in magnitude implies that certain spots have an intensity that is approximately 10^9 times greater than other spots. Furthermore, the frequency of magnitudes is not uniformly distributed, and the majority of the spots

has a small MU. A threshold θ removes spots with an MU of negligible contribution, and only keeps spots that have a valuable contribution to the irradiation.

2. *Reduction of variables* - A reduction of variables decreases the problem size, which is potentially beneficial for the optimization time. The optimization is performed with fewer decision variables, thus reducing the chance of long optimization times.

The value of θ is determined by evaluating the effect of different values of θ on the spot selection and the dose distribution. The number of remaining spots after setting the threshold should be high enough to be able to produce a spot selection of approximately 1000 spots, which is the desired solution size.

After setting the threshold θ and removing spots with $MU < \theta$, a *2pec* optimization is performed on the solution of the weighted-sum optimization problem in Eq. (4.3.1) to take the clinically desired trade-offs into account.

4.4. Deliverability settings

Steps 1 to 6 of the SISS method, as described in Figure 4.1, result in a spot selection through sparsity induction. To convert this spot selection to a clinical treatment plan, two deliverability settings are required, which are summarized in Fig. 4.2.

1. Remove spots with $MU < 1.33$.
2. Re-optimization with Pareto projection.

Figure 4.2: Deliverability steps required to produce a clinical treatment plan out of the solution of the SISS method. The steps can essentially be placed between step 6 and 7 of Fig. 4.1.

The first step of the deliverability settings is the removal of spots with $MU < 1.33$. This seems to obviate the need of a threshold θ in Section 4.3.2. However, the purpose of the minimum MU of 1.33 is to guarantee a stable beam during irradiation at HollandPTC. This purpose is not associated to contributing to a better spot selection, which is the purpose of the threshold θ .

Since the removal of spots with $MU < 1.33$ will require an additional re-optimization of the spots in Step 3, it is also obligatory to set a constraint during the optimization, which requires the $MU \geq 1.33$ for all spots, to guarantee only spots with $MU \geq 1.33$ are selected in the final solution.

In initial stages of the research, the re-optimization was performed through another *2pec* optimization, but it was later decided to substitute the *2pec* optimization with a Pareto projection, as earlier mentioned in Section 1.4.1. The Pareto projection only performs a single optimization, whereas the *2pec* optimization optimizes each objective separately in its first phase, and possibly a second time in the second phase if its goal is not met in the first phase. This substitution of a *2pec* with a Pareto projection is therefore beneficial in time savings on the optimization. The Pareto projection includes a minimum MU of 1.33 in its implementation, so it also obviates the need for setting the $MU \geq 1.33$ constraint.

4.5. Alternative weight distributions

The weights that are used in the weighted-sum optimization problem in Section 4.2.2 are extracted from the Lagrangian in Equation (3.2.5) and are normalised values between 0 and 1, as seen in Table 4.2. The third column of Table 4.2 shows the priority of each objective, where 0 is the highest priority, such as in constraints and the high CTVs, and higher numbers indicate lower priorities, such as the brainstem and spinalcord.

Table 4.2: Overview of objectives, represented by their volume names, and their accompanying weights. This particular example contains the first 10 objectives and weights from patient 1. Note that the CTVHigh, CTVintermediate10mm and CTVLow_shrunk10mm appear twice. First as constraints, which must be satisfied all the time, thus maximum weight 1, and secondly as objectives, thus having a lower weight. Note that the three constraints all have priority 0, which is the highest priority, as the constraints must be satisfied at all times.

Volume name	Weight	Priority
CTVHigh	1	0
CTVintermediate10mm	1	0
CTVLow_shrunk10mm	1	0
CTVHigh	0.047482	1
CTVintermediate10mm	0.148177	1
CTVLow_shrunk10mm	0.050694	1
CTVcombined_ring0-10mm	0.001091	2
CTVcombined_ring10-15mm	2.04e-5	2
Brainstem	0.002652	4
SpinalCord	0.00245	4

The weights in the weighted-sum optimization play a key role in deciding which spots the optimization selects for the solution. Using the weights from the resampling solution is undesirable since the resulting solution from the SISS method will depend on the resampling solution. To evaluate the influence of weight distributions other than the one generated by the resampling solution, four alternative weight distributions are applied in the ℓ_1 -norm optimization.

1. *All weights equal to 1* - The first distribution sets all weights equal to 1. This practically removes the weights from the objective function, as all objectives in the objective function become equally important.
2. *Priority-based* - The second distribution determines a weight relative to the priority, by using $w_i = 1/p_i$, where p_i is the priority of objective f_i , as seen in Table 4.2. Note that this formula is only applied to the objectives, as the constraints must always have the highest priority, since they are always to be satisfied, and therefore must have the highest weight as well.
3. *Average of 10 patients* - The third distribution uses an approximation of the weight, by using the average weights of the volume for the nine other test patients in the dataset (who are all similar tumors).
4. *Constraints only* - The last distribution is similar to the first one, as it sets the weights of the constraints and the ℓ_1 -norm equal to 1, and the weights of the objectives equal to a proportionally low value, in this case $1/30$.

5

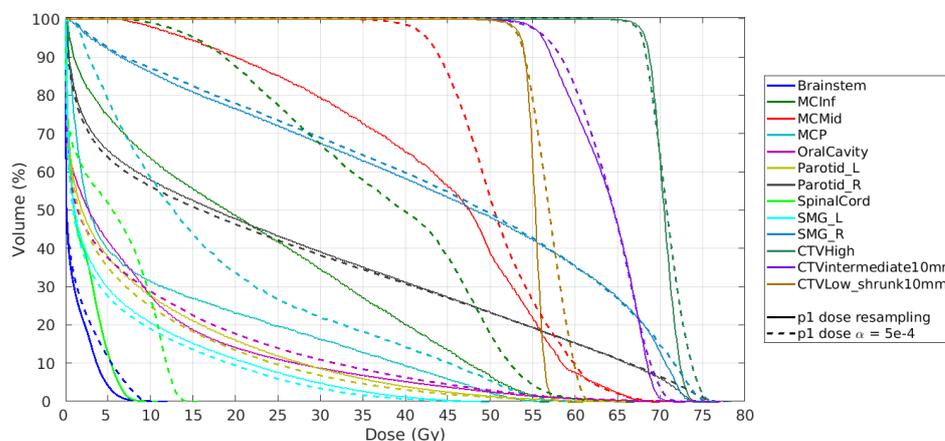
Results

5.1. ℓ_1 -norm addition

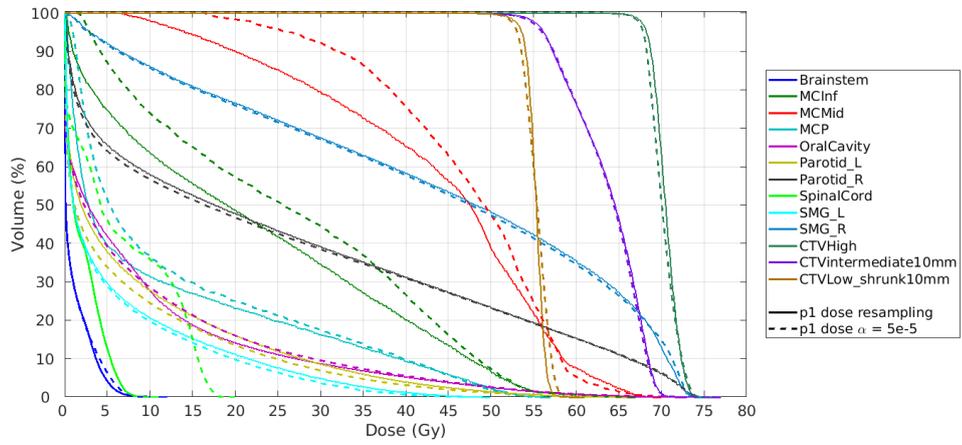
The results of the ℓ_1 -norm optimization include the tuning of parameters α and θ , as well as the resulting spot selection and dose distribution. The tuning of parameters was done without the deliverability settings, as the priority was the testing of the effect of the ℓ_1 -norm addition on the spot selection, so the DVHs in this section are computed without the steps in Figure 4.2, but this does not significantly impact the results.

5.1.1. Norm weight

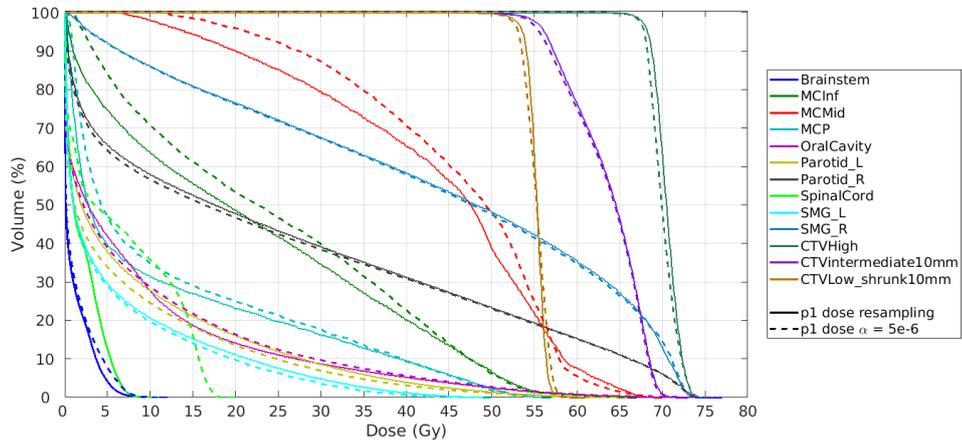
Based on research done by Janssen [1], the value of α will be somewhere in the interval of $[5e-8, 5e-4]$. To decide the value of the norm weight, five different values of α are selected, each in a different order of magnitude of the interval: $5e-4$, $5e-5$, $5e-6$, $5e-7$ and $5e-8$. For every value of α , the weighted-sum optimization problem in Equation (4.3.1) is solved and a dose distribution is made based on the resulting spot selection. Finally, a DVH is made to visualize the results.



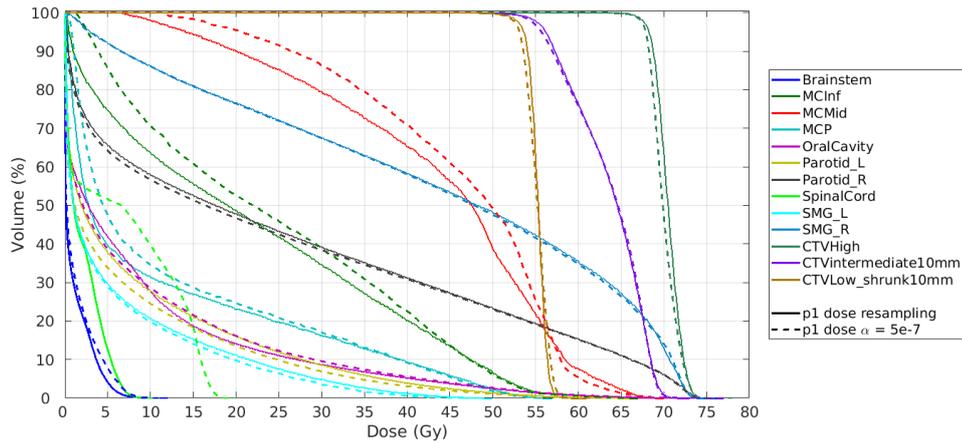
(a) Dose distribution with $\alpha = 5e-4$



(b) Dose distribution with $\alpha = 5e-5$



(c) Dose distribution with $\alpha = 5e-6$



(d) Dose distribution with $\alpha = 5e-7$

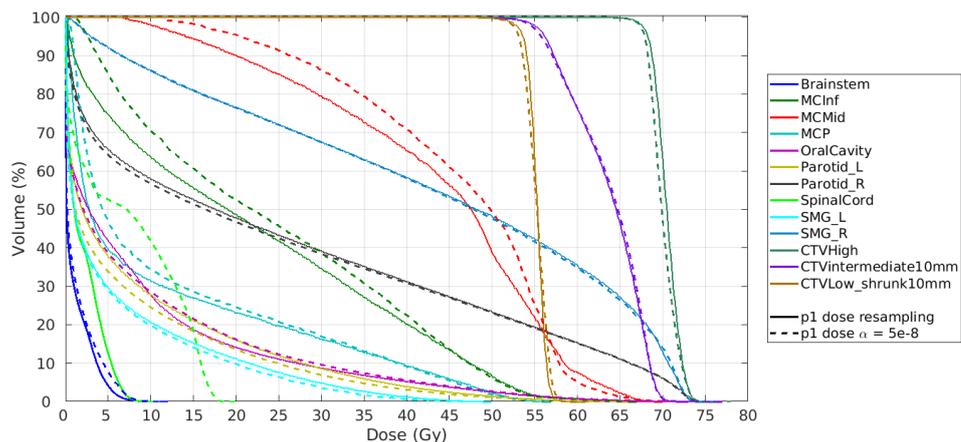
(e) Dose distribution with $\alpha = 5e-8$

Figure 5.1: Overview of 5 different DVHs for different norm weights. All the optimizations were performed on the data of patient 1.

Table 5.1: Overview of the benchmark parameters for the dose distribution based on the weighed-sum optimization with 3 different norm weights: $\alpha = 5e-6$, $\alpha = 5e-7$ and $\alpha = 5e-8$. Every triplet of volumes is assigned a color based on its size compared to the others. For the OARs, the lowest dose gets a green color, and the highest dose a red color, since a lower dose is preferable. For the CTVs, the color grading is reversed, since a higher dose is preferable. Some rows contain the same number for different α 's, but are still colored differently, this can be attributed to rounding of values.

		D0.03cc			
		Resampling	Dose 5e-6	Dose 5e-7	Dose 5e-8
Brainstem	10.1	8.9	8.9	8.9	
SpinalCord	9.8	18.0	18.2	18.1	

		Dmean			
		Resampling	Dose 5e-6	Dose 5e-7	Dose 5e-8
SMG_L	6.2	5.6	5.6	5.6	
SMG_R	43.3	43.0	43.1	43.1	
Parotid_L	8.5	7.3	7.2	7.2	
Parotid_R	26.0	25.3	25.3	25.3	
OralCavity	8.9	9.0	8.9	8.9	
MCMid	43.9	45.7	45.7	45.7	
MCSup	58.6	59.2	59.2	59.2	
MCInf	22.5	24.1	24.0	23.9	

		D2%			
		Resampling	Dose 5e-6	Dose 5e-7	Dose 5e-8
CTVHigh	73.5	73.3	73.3	73.3	
CTV54	73.1	72.9	72.9	72.9	

		D98%			
		Resampling	Dose 5e-6	Dose 5e-7	Dose 5e-8
CTVHigh	67.4	67.2	67.3	67.3	
CTV54	53.4	53.2	53.2	53.2	

Figure 5.1 shows that the dose distribution for $\alpha = 5e-4$ has significantly more deviation from the resampling solution than the other norm weights. The dose distribution for $\alpha = 5e-5$ is also slightly worse than the dose distributions for $\alpha = 5e-6$, $\alpha = 5e-7$ and $\alpha = 5e-8$, since the MCMid and MCInf receive a higher dose, while the other volumes receive a similar dose. The dose distributions for $\alpha = 5e-6$, $\alpha = 5e-7$ and $\alpha = 5e-8$ all have marginal differences, so we examine these dose distributions in more detail, by looking at the benchmark parameters: D0.03cc, Dmean, D2% and D98% doses, which are explained in Table 1.2.

As seen in Table 5.1, all three norm weights produce good quality dose distributions, as the dose in the OARs is almost everywhere lower or equal than the resampling dose, and the dose in the CTV is almost everywhere equal to the resampling dose. Only the SpinalCord seems to get a significantly higher dose than the resampling dose, which can be accounted to a trade-off made by the solver. To determine which norm weight has the best dose distribution, a score has been assigned of green = 1, yellow = 2 and red = 3 to each cell, where in the OARs the lowest dose receives a 1, and in the CTVs the highest dose receives a 1. Summing the scores for all rows and evaluating the minimum score gives a score of 29 for $\alpha = 5e-6$, 26 for $\alpha = 5e-7$ and 22 for $\alpha = 5e-8$. This indicates that $\alpha = 5e-8$ has the best dose distribution, although the differences are minimal, as indicated by the DVHs in Figure 5.1.

Since the aim is to balance the reduction of the number of spots with the plan quality of the dose distributions, the spot selection is compared as well to make a well considered decision. Since all previously shown dose distributions are made with the dataset of patient 1, all solutions have 23894 spots in their selection, as seen in Table 4.1. However, the magnitude of the MU of the spots differs for each α . This brings us to the next parameter in our solution: a threshold θ for removing spots from the solution.

5.1.2. Threshold

It is desirable to only remove spots from the solution with a negligible contribution. Removing too many spots will result in an insufficient dose on the tumor, while removing too few spots will not reduce the problem size. To find a suited value of θ , an evaluation of the frequency of orders of magnitude is required. Figure 5.2 shows the number of spots below various thresholds, computed for the three norm weights chosen in Section 5.1.1. It can be seen that for all three values of α the majority of the spots have an MU < 0.1. To test the contribution of these spots, a comparison is made between the spot selection without any threshold θ , and the spot selection on which the threshold $\theta = 0.1$ is imposed (so spots with MU < 0.1 are removed from the solution), as well as a comparison in dose distributions.

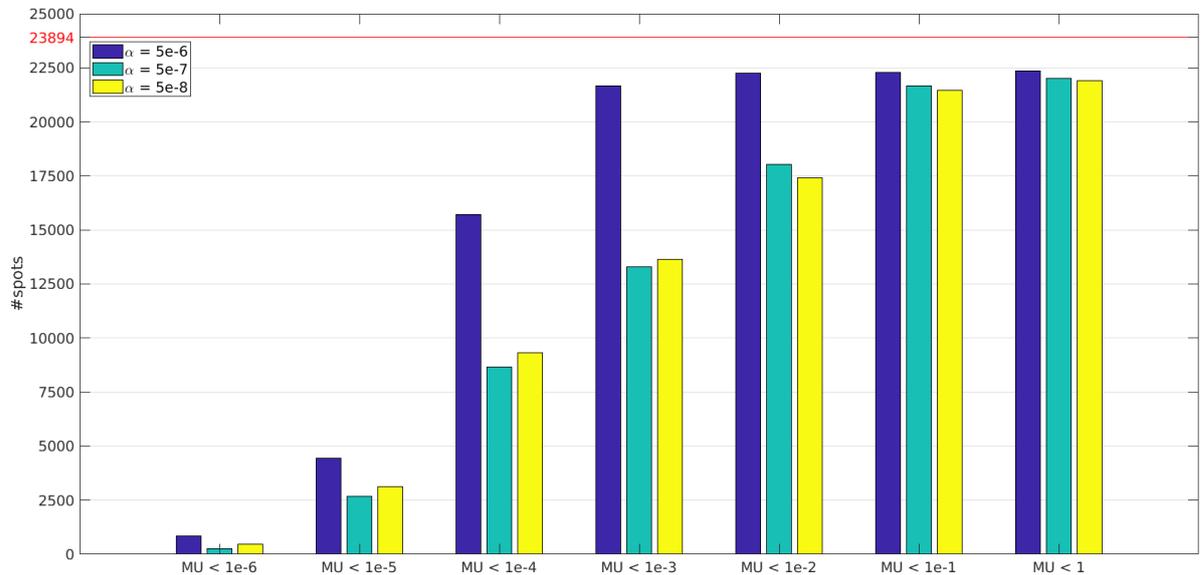


Figure 5.2: Frequency table for the number of spots in the solution of the ℓ_1 -norm optimization for seven different thresholds, computed for all three norm weights $\alpha = 5e-6$, $\alpha = 5e-7$ and $\alpha = 5e-8$ on patient 1. The red line is the total number of spots in the solution.

Figure 5.3 shows that removing the spots with MU < 0.1 has a negligible effect on the plan quality, since the lines for the dose of the resampling solution and the dose of the solution with removed spots have marginal

differences for all three norm weights. Therefore we can safely remove spots with $MU < 0.1$ from the solution of the ℓ_1 -norm optimization.

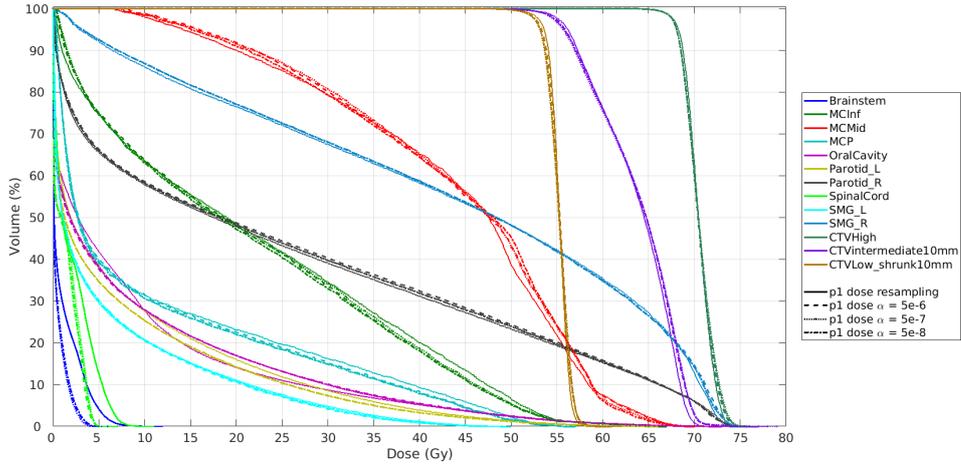


Figure 5.3: DVH of norm weights $\alpha = 5e-6$, $\alpha = 5e-7$ and $5e-8$ plotted all together against the DVH of the resampling dose. Every dose distribution was made after removing spots with $MU < 0.1$ and performing a $2pec$ optimization on the remaining spots.

Finally, the spot selections for the different norm weights are evaluated. As mentioned in Section 4.3.2, the spots that remain after removing spots with $MU < 0.1$, first must be optimized using a $2pec$ optimization, to take the clinical trade-offs into account. The $2pec$ optimization results in 1857 spots for $\alpha = 5e-6$, 2234 spots for $\alpha = 5e-7$ and 2427 spots for $\alpha = 5e-8$. As expected, this is inversely proportional to the plan quality of the dose for each norm weight, making all three values for α a valid choice for the norm weight. To properly balance the plan quality of the dose distribution with the number of spots in the solution, $\alpha = 5e-7$ is chosen as norm weight.

5.1.3. Spot selection and dose distribution

After choosing parameters $\alpha = 5e-7$ and $\theta = 0.1$, the first six steps of the SISS method, as seen in Figure 4.1, are executed, which results in the following spot selection for the 10 test patients.

Table 5.2: Overview of the number of spots selected for all 10 test patients in three phases: in the spot coverage, in the resampling solution, and in the SISS method (after the $2pec$ optimization).

Patient	Spot coverage	Resampling solution	SISS (2pec)
1	23894	1289	2234
2	23881	1064	1578
3	23909	1225	1974
4	23879	990	1747
5	23871	961	1701
6	23901	1185	1861
7	23917	1306	2034
8	23888	1060	1817
9	23857	873	1566
10	23832	973	1151
Average	23882	1074	1766

Table 5.2 shows that the average number of spots selected by the SISS method is 698 spots higher than the average number of spots in the resampling solution. However, the solution of the SISS method still contains

spots that need to be removed for deliverability of the treatment plan, which is further explained in the next section. First a comparison is made of the dose distribution of the solution of the SISS method and the resampling solution. This comparison is visualized in Figure 5.4, and the individual DVHs are seen in Appendix A.1.

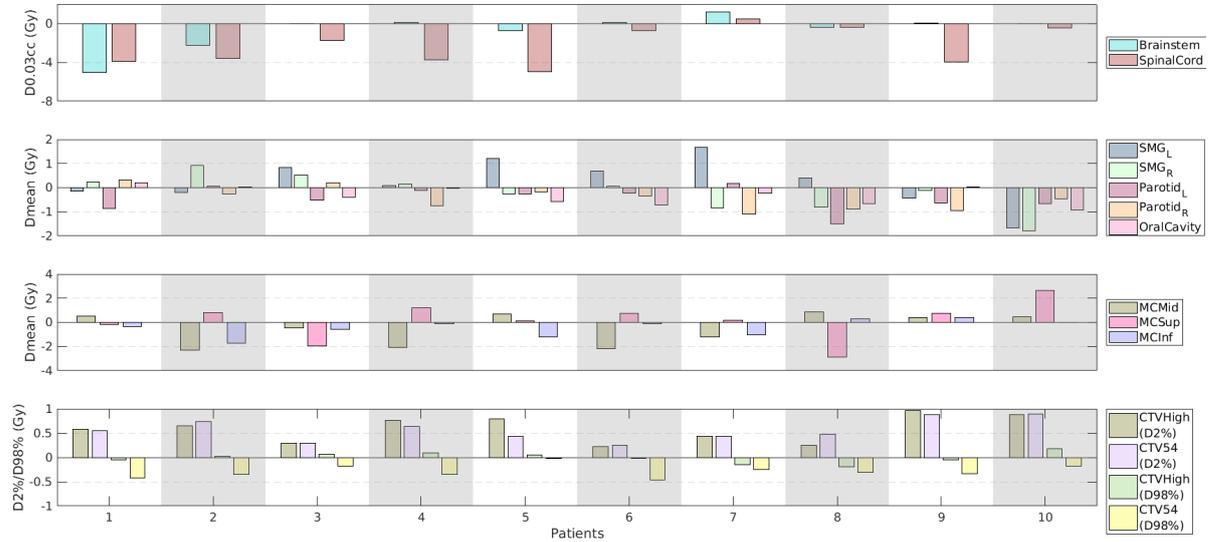


Figure 5.4: Bar charts for the differences between the dose of the resampling method and the dose after the $2pec$ optimization, as part of the SISS method. The bars represent the dose of the resampling method, subtracted from the dose of the SISS method (after the $2pec$ optimization). Therefore, for the D0.03cc and Dmean bar charts (OARs), if a bar is negative, it is in favor of the SISS method. For the D2%/D98% bar chart (CTV), the principle is reversed: a positive bar is in favor of the SISS method.

The bar charts in Figure 5.4 show that on an average of 10 patients, 6 out of 10 OARs receive a lower dose with the SISS method than with the resampling method. For the OARs that receive a higher dose with the SISS method, the average surplus dose is 0.6 Gy, with a maximum of 2.7 Gy, found at the MCSup for patient 10. The bar charts also show that the doses in the CTVs have marginal differences. While the near-minimum dose of the tumor (D98%) receives on average 0.1 Gy less with the SISS method than with the resampling dose, the near-maximum dose of the tumor (D2%) receives on average a dose surplus of 0.6 Gy.

5.2. Deliverability settings

As mentioned in Section 4.4, the spot selection from Table 5.2 must be converted to a clinical treatment plan that is executable by the radiation therapist. This conversion consists of removing spots with $MU < 1.33$, setting the constraint $MU \geq 1.33$ in the optimization, and performing a Pareto projection on the remaining spots. This results in the spot selections seen in Table 5.3.

Table 5.3 shows that setting the minimum MU at 1.33 results in a reduction of spots by 607 spots on average, in comparison to the $2pec$ optimization. This reduction of spots makes the spot selection of the SISS method deviate less than 8% from the desired solution size of approximately 1000 spots, as mentioned in Section 4.2.1. Based on this new spot selection, dose distributions have been made for the 10 test patients. The DVHs based on the dose distributions are shown in Appendix A.2. Using the benchmarks parameters from Table 1.2, an overview is presented in Table 5.4 of the delivered doses for the 10 test patients.

Table 5.3: Overview of the spots selected for the 10 test patients at four different stages. The first and second column contain respectively the spots from the spot coverage and from the resampling solution. The third and fourth column contain the spots generated by the SISS method after respectively the $2pcc$ optimization and the Pareto projection.

Patient	Spot coverage	Resampling solution	SISS (2pcc)	SISS (Pareto projection)
1	23894	1289	2234	1397
2	23881	1064	1578	1097
3	23909	1225	1974	1394
4	23879	990	1747	1150
5	23871	961	1701	1026
6	23901	1185	1861	1189
7	23917	1306	2034	1384
8	23888	1060	1817	1213
9	23857	873	1566	916
10	23832	793	1151	825
Average	23882	1074	1766	1159

Table 5.4: Overview of the benchmark parameters from Table 1.2, applied on the dose of the SISS method, after performing the final Pareto projection. Note that the values are rounded to one decimal for readability, which accounts for the 0.0 Gy in the Brainstem for patient 9. The high doses in the Brainstem and SpinalCord for patient 7 can be attributed to the tumor being positioned next to the brain stem.

		D0.03cc (Gy)									
		Patient 1	Patient 2	Patient 3	Patient 4	Patient 5	Patient 6	Patient 7	Patient 8	Patient 9	Patient 10
Brainstem		4.0	0.6	0.6	1.0	1.8	0.9	70.9	0.3	0.0	0.1
SpinalCord		4.8	4.2	1.0	5.2	4.8	1.2	18.9	1.5	1.7	6.3

		Dmean (Gy)									
		Patient 1	Patient 2	Patient 3	Patient 4	Patient 5	Patient 6	Patient 7	Patient 8	Patient 9	Patient 10
SMG_L		5.9	8.6	55.7	15.5	60.6	63.7	67.2	58.2	46.6	27.3
SMG_R		43.4	70.6	60.7	66.8	23.4	41.2	32.3	32.1	26.3	26.8
Parotid_L		7.5	7.0	16.9	10.9	17.2	28.1	35.6	13.7	60	9.5
Parotid_R		26.1	16.6	32.1	22.4	11.5	5.5	12.2	7.5	8.2	6.9
OralCavity		8.9	44.3	26.0	23.0	20.9	23.1	7.8	21.1	4.4	9.6
MCMid		43.5	47.4	56.8	50.6	42.0	31.4	38.9	64.9	64.8	66.8
MCSup		58.3	50.4	21.7	61.6	59.4	66.9	51.9	28.0	39.7	38.0
MCInf		21.3	8.6	11.1	23.1	14.5	13.3	18.3	40.9	67.0	46.1

		D2% (Gy)									
		Patient 1	Patient 2	Patient 3	Patient 4	Patient 5	Patient 6	Patient 7	Patient 8	Patient 9	Patient 10
CTVHigh		73.8	73.5	73.3	73.7	73.9	73.5	73.8	73.6	73.2	73.6
CTV54		73.4	73.0	72.8	73.0	72.9	73.0	73.4	72.6	72.2	72.3

		D98% (Gy)									
		Patient 1	Patient 2	Patient 3	Patient 4	Patient 5	Patient 6	Patient 7	Patient 8	Patient 9	Patient 10
CTVHigh		67.4	67.7	67.4	67.7	67.7	67.5	67.2	67.0	67.8	68.0
CTV54		53.2	52.5	52.9	52.5	52.5	52.8	52.9	52.8	52.6	52.7

To compare the results from Table 5.4 with the dose distributions from the resampling method, their differences are visualized in Figure 5.5. This comparison shows that on an average of 10 patients, 6 out of 10 OARs receive a lower dose with the SISS method after deliverability settings, than with the resampling method. For the OARs that received a higher dose with the SISS method after deliverability, the average surplus dose is 0.6 Gy, with a maximum of 2.6 Gy, found at the MCSup for patient 10. The differences in dose for the CTVs are also marginal. While the near-minimum dose of the tumor (D98%) receives on average 0.2 Gy less with the SISS method, the near-maximum dose of the tumor (D2%) receives an average dose surplus of 0.5 Gy. These results show that the reduction in spots do not decrease the quality of the dose distributions.

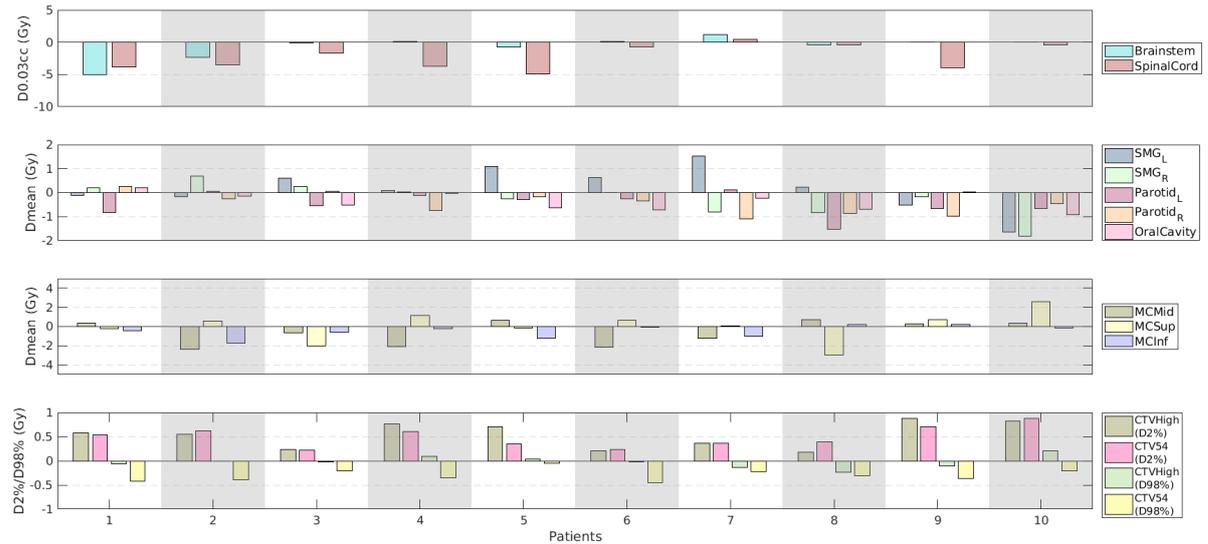
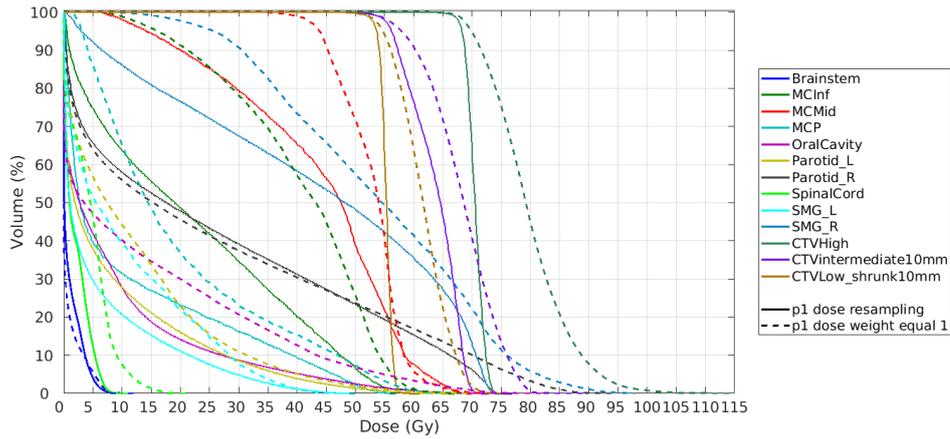


Figure 5.5: Bar charts for the differences between the dose of the resampling method and the dose after the $2pec$ optimization, as part of the SISS method. The bars represent the dose of the resampling method, subtracted from the dose of the SISS method (after the $2pec$ optimization). Therefore, for the D0.03cc and Dmean bar charts (OARs), if a bar is negative, it is in favor of the SISS method. For the D2%/D98% bar chart (CTV), the principle is reversed: a positive bar is in favor of the SISS method.

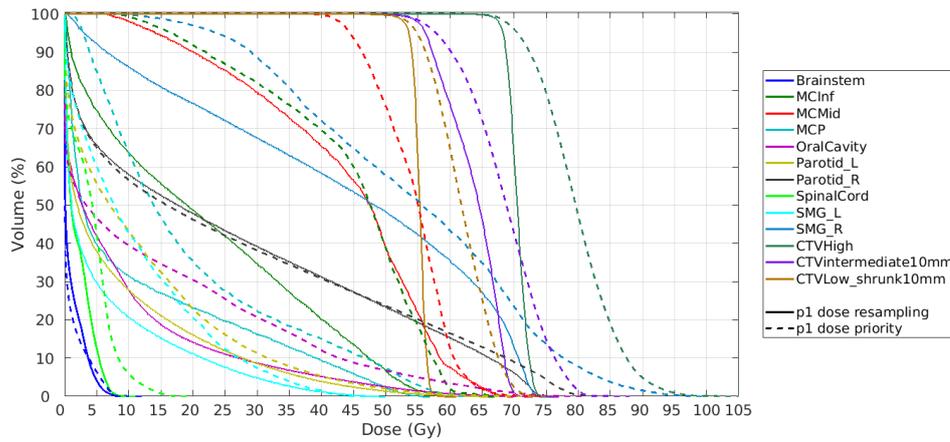
5.3. Alternative weight distributions

To evaluate the influence of weight distributions other than the one generated by the resampling solution, four alternative weight distributions are applied in the ℓ_1 -norm optimization. The four alternative weight distributions (equal to 1, priority-based, average, constraints) resulted in spot selections of respectively 417 spots, 404 spots, 364 spots and 449 spots. Based on these spot selections, the dose distributions were computed, which resulted in the DVHs in Figure 5.6.

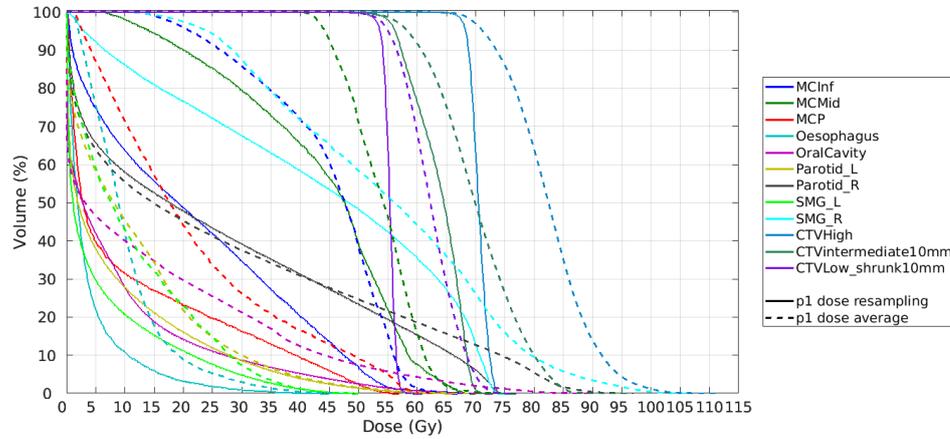
Figure 5.6 shows that for all four weight distributions, the resulting dose distributions deteriorated in comparison to the resampling dose. Weight distributions 1,2 and 4 show dose surpluses in almost all the OARs, that are of unacceptable magnitude, i.e., 90% of MCMid receiving 45 Gy instead of 20 Gy in Figure 5.6a. The CTVs of these weight distributions also exceed the resampling dose significantly, i.e., max dose of CTVHigh reaching 105 Gy, instead of 75 Gy in Figure 5.6b). Weight distribution 3 shows lower doses in the OARs compared to the other distributions, although still surpluses compared with the resampling solution, but the CTVs also exceed the resampling dose significantly.



(a) Weight distribution 1: all weights set equal to 1.



(b) Weight distribution 2: weights based on the priority of the volume.



(c) Weight distribution 3: average weight of the 9 other patients

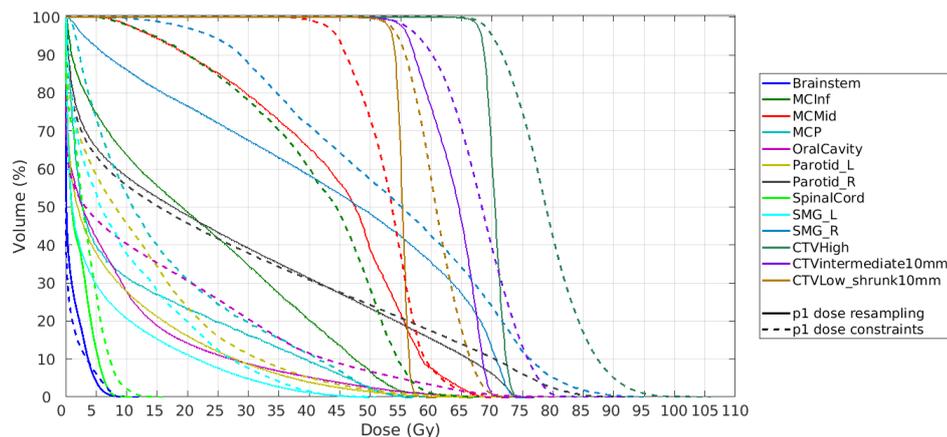
(d) Weight distribution 4: constraints and ℓ_1 -norm have weight 1, objectives have weight 0.

Figure 5.6: DVHs of four different dose distributions, compared with the dose distribution of the resampling method. All the dose distributions are based on the spot selection by the SISS method.

5.4. Voxel reduction

While this research focuses on the viability of obtaining a spot selection that is as good as the resampling spot selection with the use of sparsity inducement, another aspect that came up during the research is the computation time of resulting spot selection of the SISS method. If the spot selection takes too long to compute, the patient is required to hold still for a longer period of time, which is an unpleasant experience for the patient and also increases the chance of the patient moving more after a while, resulting in more uncertainty in the dose distribution. Furthermore, a longer treatment time per patient would also result in less treatments per day, which is undesirable for the hospital. The computation times of the different optimizations in the SISS method are seen below.

Table 5.5: Computation time of the optimization done in the SISS method.

Patient	ℓ_1 -norm optimization	2pcc optimization	Pareto projection	Total
1	1:59:59	0:05:04	0:00:17	2:05:20
2	1:55:04	0:03:37	0:00:08	1:58:49
3	2:03:47	0:04:18	0:00:15	2:08:20
4	1:53:10	0:04:09	0:00:13	1:57:32
5	1:40:43	0:03:33	0:00:12	1:44:28
6	1:31:28	0:04:37	0:00:09	1:36:14
7	1:37:00	0:05:10	0:00:11	1:42:21
8	1:58:44	0:03:52	0:00:13	2:02:49
9	2:04:00	0:03:45	0:00:09	2:07:54
10	1:30:24	0:03:02	0:00:07	1:33:33
Average	1:49:26	0:04:07	0:00:11	1:53:44

Table 5.5 shows that the ℓ_1 -norm optimization contributes the most to the total optimization time. One way to reduce this computation time, is by reducing the number of voxels on the patients volume. A voxel is a discretized piece of the patients volume, as seen in Figure 1.2. If a voxel is removed from the discretization, spots cannot radiate that voxel in the optimization, which can potentially save computation time, but also decrease the plan quality, as less spots will be selected in the solution.

Voxel reduction has been performed for three different reduction rates: 50%, 33% and 25%, removing re-

spectively every second, third and fourth row of voxels. The remaining voxel rows were used in the ℓ_1 -norm optimization to produce a new spot selection. After the ℓ_1 -norm optimization, the rest of the SISS method proceeds without the voxel reduction, keeping all the voxel rows. Since the computation times of the $2pec$ optimization and the Pareto projection are not influenced by the voxel reduction, only the computation times of the ℓ_1 -norm optimization are considered. A comparison is made in Table 5.6 between the ℓ_1 -norm optimization in the SISS method, and the ℓ_1 -norm optimization with the voxel reductions. Table 5.6 shows that the three voxel reductions have reduced the computation times with respectively 9%, 19% and 25%. The accompanying spot selections are seen in Table 5.7.

Table 5.6: Computation times of the ℓ_1 -norm optimization, performed with the regular SISS method, and with voxel reduction of 25%, 33% and 50%. The voxel reduction of 50% on patient 7 did not converge, thus producing a dose distribution with excessively high dose surpluses.

Patient	No voxel reduction	25% voxel reduction	33% voxel reduction	50% voxel reduction
1	1:59:59	1:35:51	1:26:38	1:23:17
2	1:55:04	1:49:02	1:37:50	1:14:25
3	2:03:47	2:00:58	1:54:28	1:33:33
4	1:53:10	1:43:31	1:33:08	1:29:42
5	1:40:43	1:35:08	1:21:14	1:18:31
6	1:31:28	1:22:34	1:12:48	1:04:09
7	1:37:00	1:03:49	1:10:50	~
8	1:58:44	1:56:42	1:27:47	1:25:48
9	2:04:00	1:55:03	1:45:39	1:32:03
10	1:30:24	1:28:19	1:14:05	1:13:08
Average	1:49:26	1:39:06	1:28:27	1:21:37

Table 5.7: Table showing the spot selection of the SISS method after performing a for voxel reductions of respectively 50%, 33% and 25%, in comparison to the spot selections from the resampling method and the SISS method without a voxel reduction. The voxel reduction of 50% on patient 7 did not converge, thus producing a dose distribution with excessively high dose surpluses.

Patient	Resampling solution	SISS solution	Voxel reduction 25%	Voxel reduction 33%	Voxel reduction 50%
1	1289	1397	1040	886	961
2	1064	1097	790	854	775
3	1225	1394	975	1098	644
4	990	1150	817	837	689
5	961	1026	883	722	731
6	1185	1189	950	869	736
7	1306	1384	1019	1077	~
8	1060	1213	998	831	813
9	873	916	699	715	701
10	793	825	680	494	496
Average	1074	1159	885	748	723

Table 5.7 shows that the average number of spots in the solution decreases as the number of voxels removed increases. Since less spots are selected in the solution, it is expected that this will decrease the plan quality of the dose distributions based on the voxel reductions. These dose distributions are visualized in Figure 5.7.

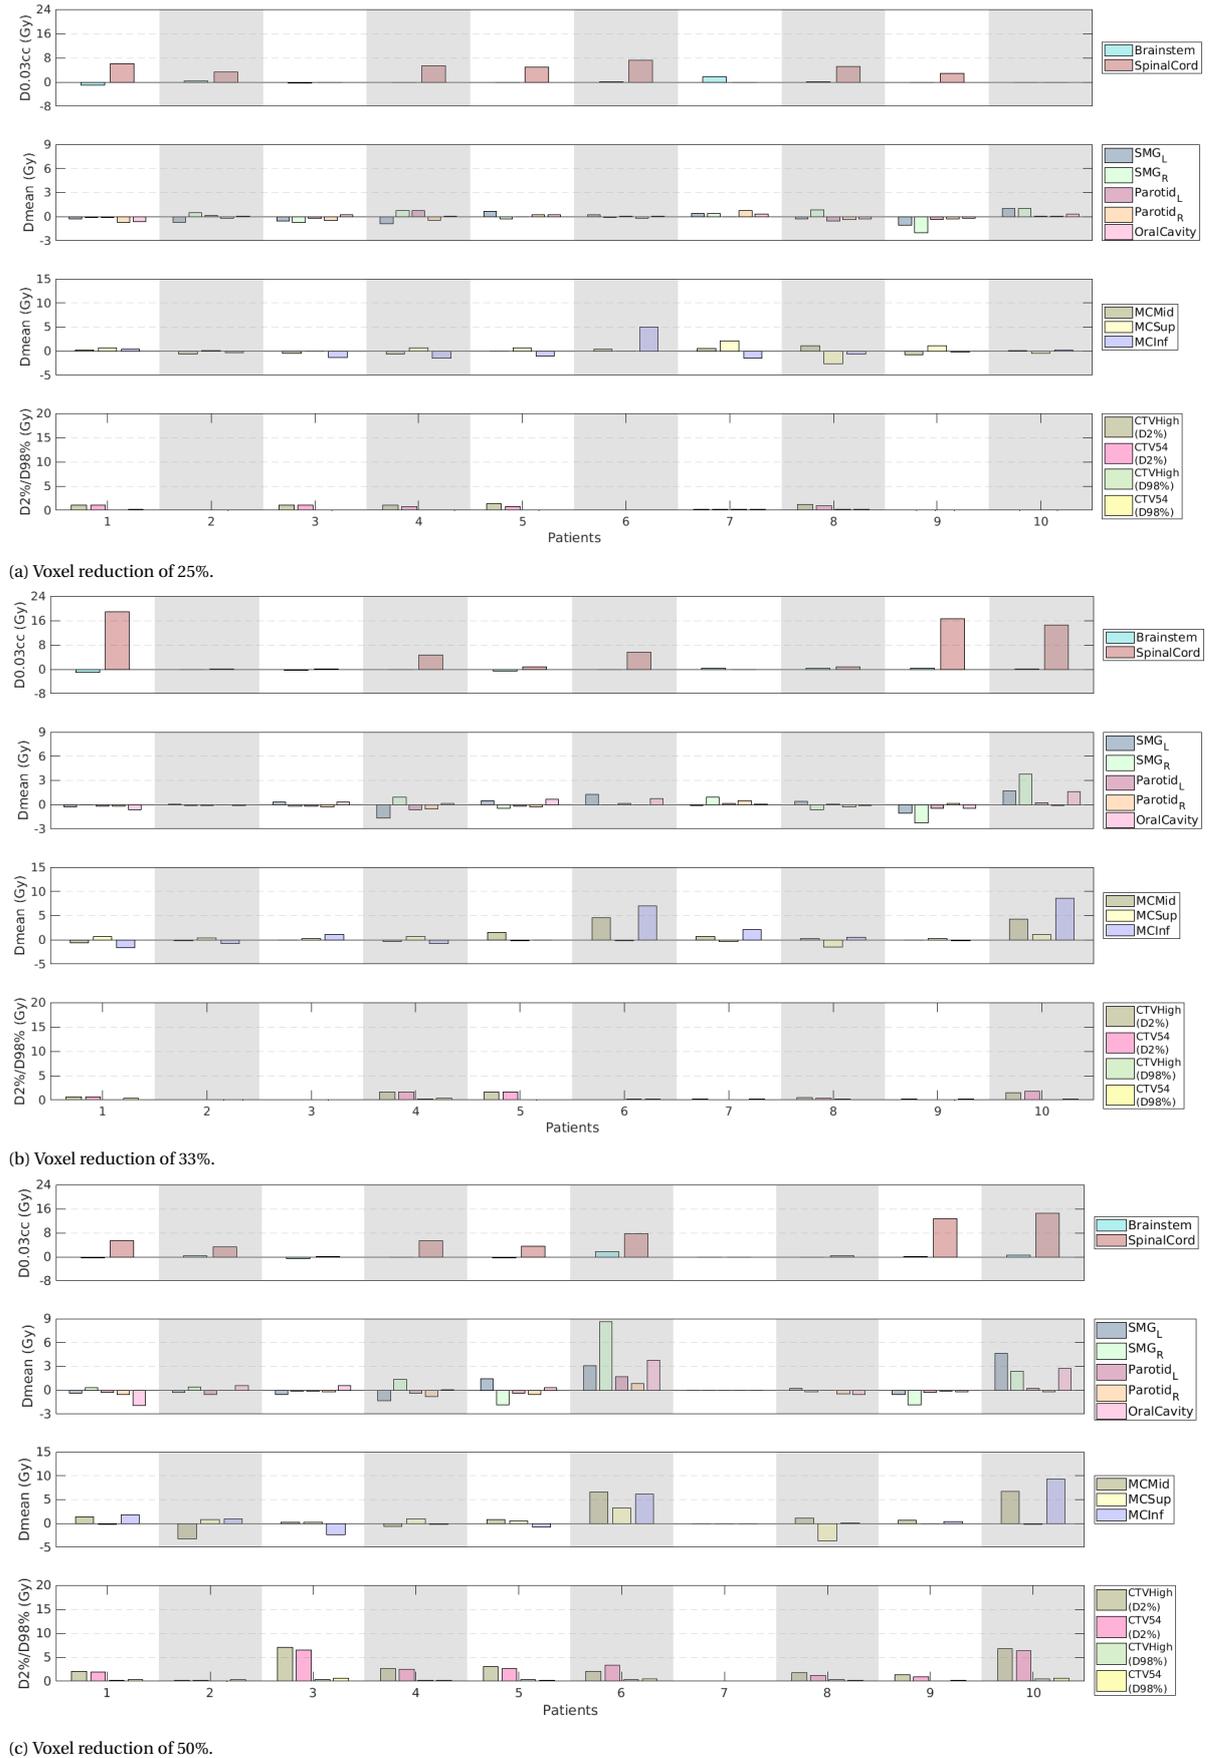


Figure 5.7: Bar charts of the effect of a voxel reduction of 25% (a), 33% (b) and 50% (c) on all 10 test patients. The bars represent the dose of the SISS method, subtracted from the dose of the voxel reduction. So for the D0.03cc and Dmean (OARs), if a bar is negative, it is in favor of the voxel reduction. For the D2% and D98% (CTV), the principle is reversed: a positive bar is in favor of the voxel reduction. Note that not all differences in dose show a visible bar. For these volumes, the difference in dose between the voxel reduction and the SISS method is marginal (≈ -0.1), and would barely be visible in the bar chart. It was decided to leave this negative bar out of the image in favor of showing the comparison between the voxel reductions on the same axes. The voxel reduction of 50% on patient 7 did not converge, thus producing a dose distribution with excessively high dose surpluses.

For the voxel reduction of 25%, the mean dose (D_{mean}) in the OARs shows similarity to the dose of the SISS method. On average, 8 out of 10 OARs received a dose that was within 1 Gy of the dose of the SISS method. Only the SpinalCord received a significantly higher dose with an average surplus dose of 3.5 Gy, with a maximum of 7.2 Gy for patient 6. The CTVs also received a similar dose to the SISS method, as the highest (absolute) deviation received a dose surplus of 1.3 Gy.

For the voxel reduction of 33%, the mean dose in the OARs also shows similarity to the dose of the SISS method, but more deviations occur and with a higher magnitude. Especially the SpinalCord receives significantly higher doses, with three patients receiving a dose surplus of 15 Gy. The CTVs are similar in dose to the SISS method.

Finally the voxel reduction of 50% shows further deterioration in the doses. Although patients 1-5 and 9 receive a similar mean dose to the SISS method, patients 6 and 10 receive an average dose surplus of 4.0 Gy. For the voxel reduction of patient 7, the solver did not converge, thus producing an unacceptable dose distribution. Therefore it is removed from the results.

6

Discussion

6.1. Norm weight

The SISS method starts by converting the resampling solution to a weighted-sum formulation, by extracting the weights of the resampling solution. Next the ℓ_1 -norm was added to the objective function of the weighted-sum problem with norm weight α . For the choice of a norm weight α , five values of different magnitude were evaluated. There are numerous other choices for α that might lead to a better spot selection, but due to time constraints, and the marginal difference in dose distributions, it was decided work with the five selected norm weight candidates.

The comparison of norm weights consisted of solving the weighted-sum optimization problem with different norm weights and comparing the resulting spot selections and dose distributions. The dose distributions for $\alpha = 5e-6$, $\alpha = 5e-7$ and $\alpha = 5e-8$ all were comparable in plan quality. The dose distribution for $\alpha = 5e-8$ showed a slightly better plan quality, but only with marginal differences. This small improvement over the other dose distributions can most likely be attributed to $\alpha = 5e-8$ having the lowest weight on the penalty in the objective function in the weighted-sum optimization problem, thus leaving the most spots in the solution.

6.2. Threshold

To choose a value for α , we first take the effect of a threshold θ on the solution into account. This threshold θ has the benefit of removing redundant spots from the solution and decreasing the solution size. After evaluating the frequency of different magnitudes of the spots, threshold $\theta = 0.1$ was chosen. This is not an exact solution for the optimal threshold, but an approximation based on testing. Similar to the norm weight α , due to time constraints and satisfying results with $\theta = 0.1$, it was decided to take $\theta = 0.1$ as a threshold. To compare the effect of the thresholds on the different norm weights, a *2pec* optimization was required to take the clinical trade-offs into account. The resulting dose distributions for the three norm weights contain marginal differences. It could be argued that due to this similarity of the dose distributions, only the number of spots should be taken into account. This would imply that norm weight $\alpha = 5e-6$ would be the best choice. It was however decided to choose a norm weight with a balance in the number of spots and the plan quality of the dose distributions. This resulted in a norm weight of $\alpha = 5e-7$.

6.3. Spot selection and dose distribution

After applying the norm weight $\alpha = 5e-7$ in the ℓ_1 -norm optimization and setting threshold $\theta = 0.1$, a *2pec* is performed, which results in a spot selection of 1766 spots on average. While this number is relatively higher than the average solution size of 1074 spots, it still requires spots to be removed for deliverability of the treatment plan.

Next, to create a clinically acceptable treatment plan, a stable beam during radiation is required, which removes spots with $MU < 1.33$. This results in the average number of spots in the solution of the SISS method being reduced to 1159, which deviates only 8% from the desired solution size of 1000 spots and less than 1% from the resampling solution. As a final step of the SISS method, a Pareto projection is performed to project

the resulting solution onto a Pareto front. The accompanying dose distributions again shows improvement over the dose distributions of the resampling method. For the average patient, 6 out of 10 OARs received a lower dose with the SISS method, and the remaining four patients had comparable dose distributions.

These numbers show that the quality of the solutions produced by the SISS method is comparable, and even slightly better, to the quality of the solutions produced by the resampling method. The improvement in dose distributions resulting in a superior spot configuration is expected due to the SISS method optimizing all spots at once, instead of only optimizing a subset of the spots, and then adding new spots, as done in the resampling method.

6.4. Voxel reduction

The voxel reduction of 50% resulted in reduction in computation time of 25%, but also in high dose surpluses in multiple OARs, thus making the resulting dose distributions unacceptable to use in practice. The voxel reductions of 25% and 33% showed a comparable dose distribution to that of the resampling solution, but the reduction in computation time was only 9% and 19% respectively. This shows that voxel reduction saves time during the weighted-sum optimization, but it could be argued that the ratio of (computation time reduced):(plan quality lost) is not beneficial enough.

To generate clinically acceptable treatment plans, treatment plans are made robust against patient positioning errors and proton range errors. Robust treatment planning takes significantly longer compared to non-robust treatment planning. It is expected that a voxel reduction of 25% will still reduce computation time of the SISS method, but the percentage by which the computation time is reduced could differ from the results found in Section 5.4.

Performing the voxel reductions did not only reduce the computation time of the weighted-sum optimization, as expected, but also the number of spots in the solution. During the selection of spots, the voxels on the volume of the patient should be covered as good as possible. A voxel reduction removes voxels from the discretization, which implies less spots need to cover those voxels. This could imply that less spots are required in the weighted-sum optimization.

6.5. Weight distribution

The selected weights in the weighted-sum optimization are extracted from the resampling solution. Because of this extraction, the current solution of the SISS method is still dependent on the resampling solution. In the future, it would be desirable to find a set of weights that produce a dose distribution with equally acceptable or improved plan quality, that can be generated without the dependence on the resampling method.

7

Conclusion

In IMPT, the current approach for a spot selection in treatment planning at the Erasmus MC is based on an iterative resampling method, which uses trial and error to compose a solution. Furthermore, the resampling method contains a number of complex steps, which decrease the general transparency of the method. To create a more mathematically substantiated and clearly defined method, we present the SISS method, a method for obtaining a spot selection through the use of sparsity inducing norms, as well as a potential method for reduction in computation time of the SISS method. With the SISS method, we aimed to find a spot selection of approximately 1000 spots, with comparable or improved plan quality of the accompanying dose distribution.

The SISS method adds the ℓ_1 -norm to the weighted-sum problem with norm weight $\alpha = 5e-7$ and solves the problem for the optimal MU per spot. A threshold $\theta = 0.1$ is imposed on the spot selection to remove redundant spots. Both α and θ have been tweaked until a clinically acceptable dose distribution was made. To take clinical trade-offs into account, the resulting spot selection is optimized with a *2pec* optimization. To account for deliverability, spots with < 1.33 were removed to guarantee a stable beam, and a Pareto projection was performed to project the solution on a Pareto front.

The iterative resampling method results in a spot selection of 1074 spots on average, for a dataset of 10 head and neck patients. The SISS method results in a spot selection of 1159 spots on average, for the same dataset. While this number is higher than the average spot selection of the resampling solution, it is still in the desired range of approximately 1000 spots.

The plan quality of the dose distributions showed improvement over the dose distributions of the resampling method. For the average patient, 6 out of 10 OARs received a lower dose with the SISS method, and the remaining four patients had comparable dose distributions.

To reduce the computation time of the SISS method, voxel reduction was applied during the weighted-sum optimization. Voxel reductions of 50% and 33% resulted in reducing the average computation time of the weight-sum optimization by respectively 25% and 19%, but also in too many cases of dose surpluses in OARs. A voxel reduction of 25% reduced the computation time by 9%, and resulted in an acceptable plan quality.

A limitation of this research is the weights in the ℓ_1 -norm optimization. These weights are extracted from the resampling solution, which makes the SISS method dependent on the resampling solution.

Future research could be done on constructing an appropriate set of weights that is independent of other distributions, and improving the computation time of the SISS method, for example by using a different solver.

A

Dose-Volume Histograms

A.1. Dose-Volume Histograms after $2pe$ optimization

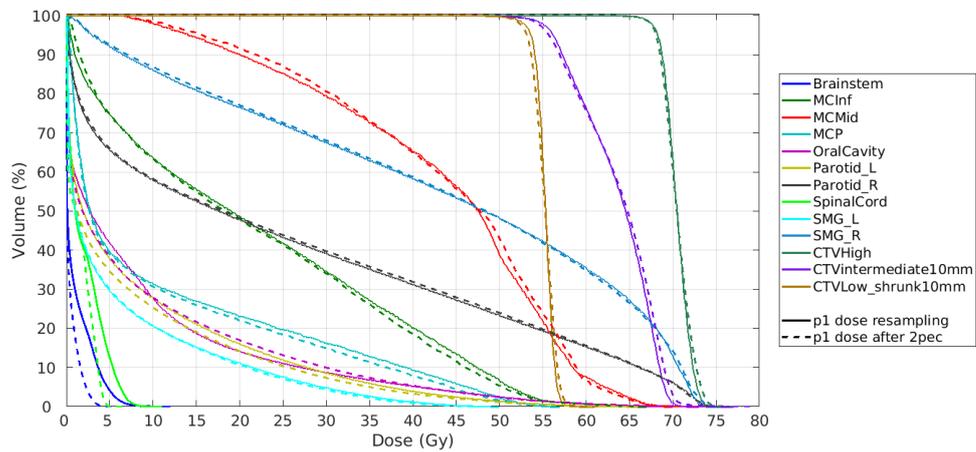


Figure A.1: Dose distribution for patient 1, generated with the SISS method (after the $2pec$ optimization).

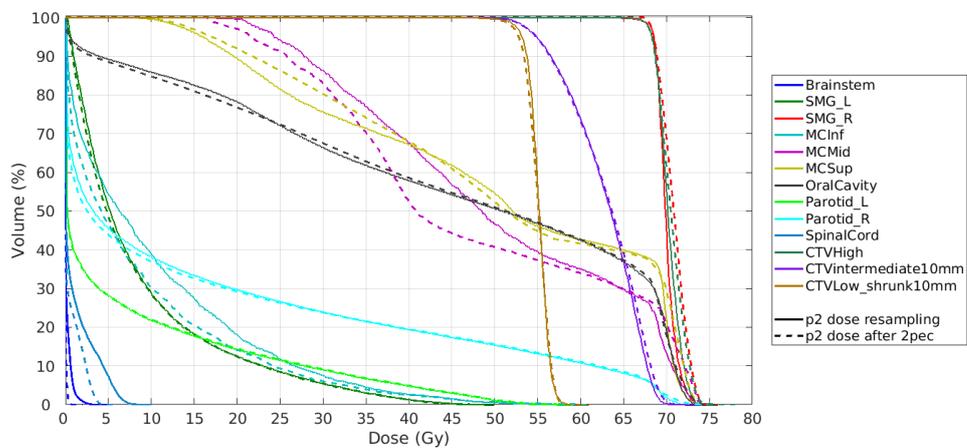


Figure A.2: Dose distribution for patient 2, generated with the SISS method (after the $2pec$ optimization).

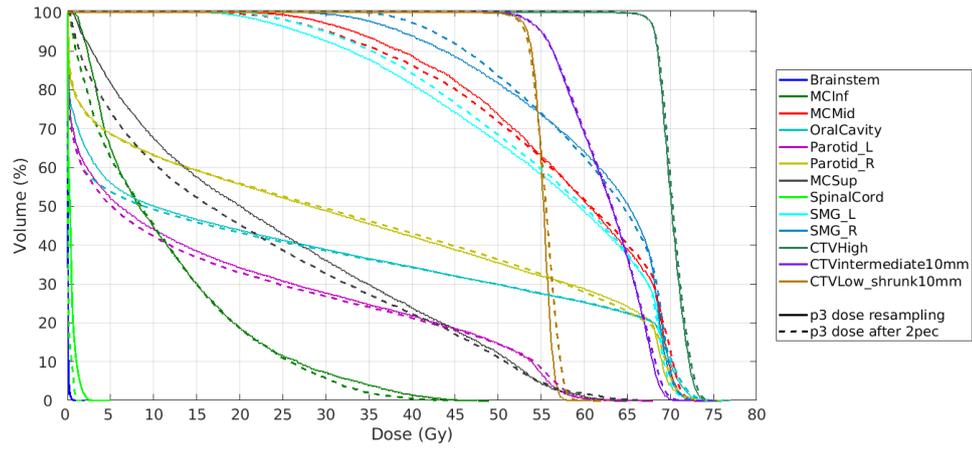


Figure A.3: Dose distribution for patient 3, generated with the SISS method (after the $2pec$ optimization).

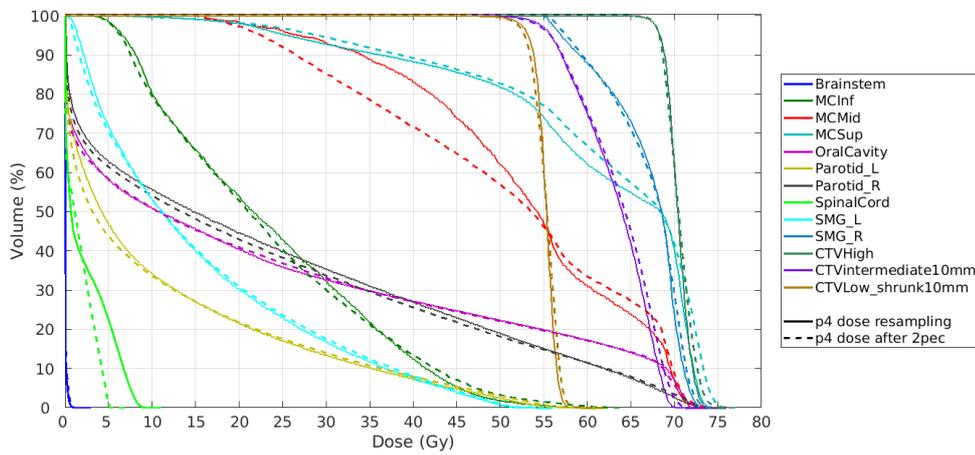


Figure A.4: Dose distribution for patient 4, generated with the SISS method (after the $2pec$ optimization).

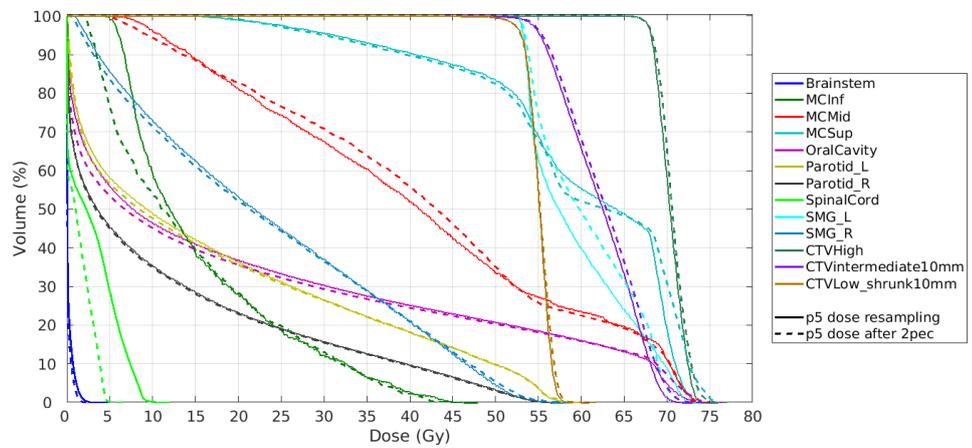


Figure A.5: Dose distribution for patient 5, generated with the SISS method (after the $2pec$ optimization).

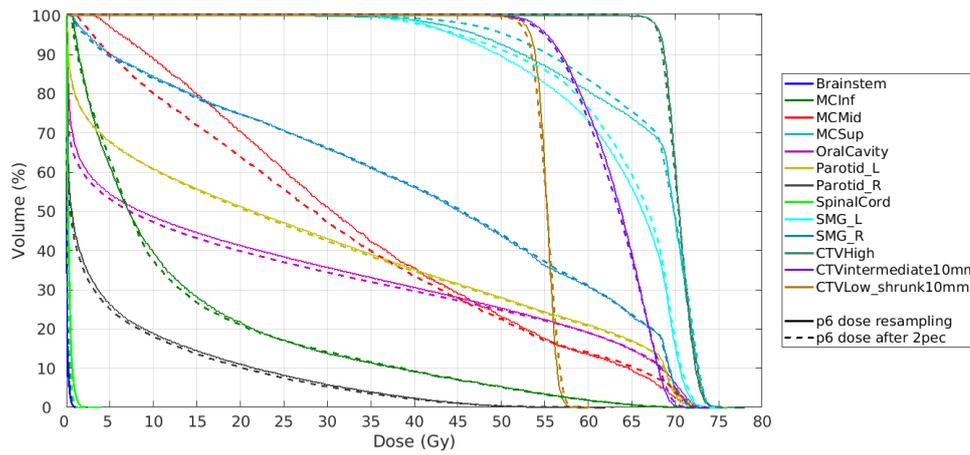


Figure A.6: Dose distribution for patient 6, generated with the SISS method (after the 2pec optimization).

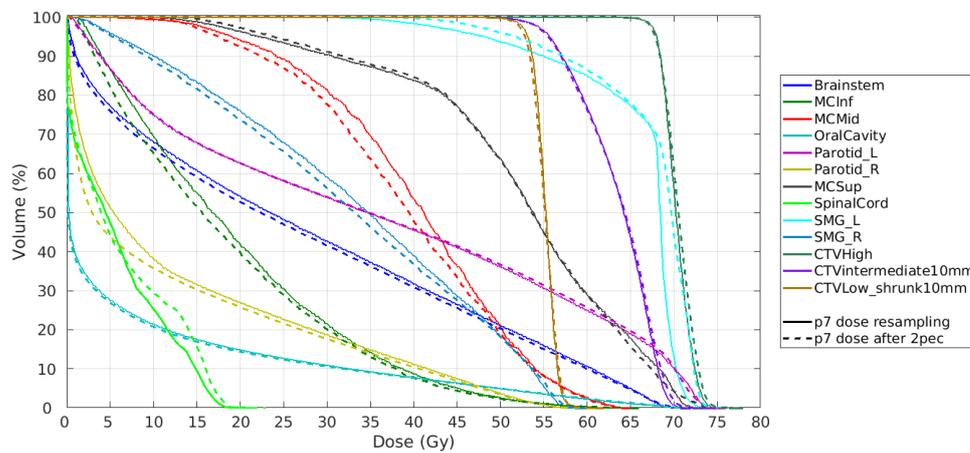


Figure A.7: Dose distribution for patient 7, generated with the SISS method (after the 2pec optimization).

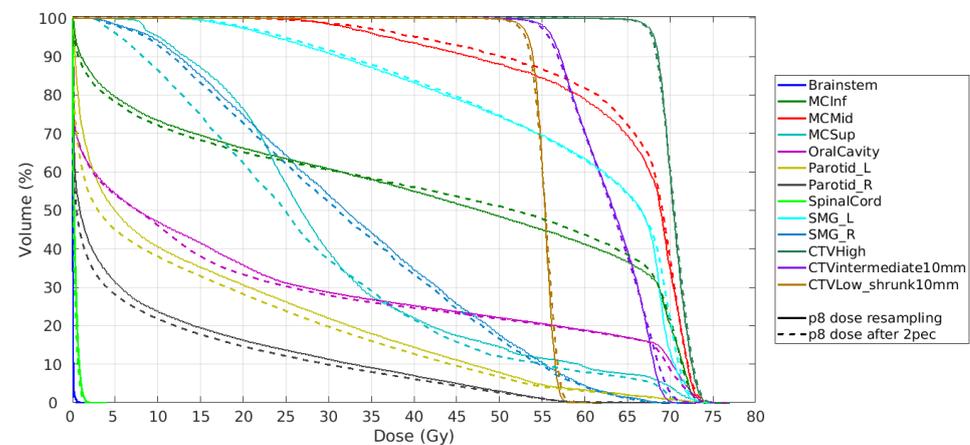


Figure A.8: Dose distribution for patient 8, generated with the SISS method (after the 2pec optimization).

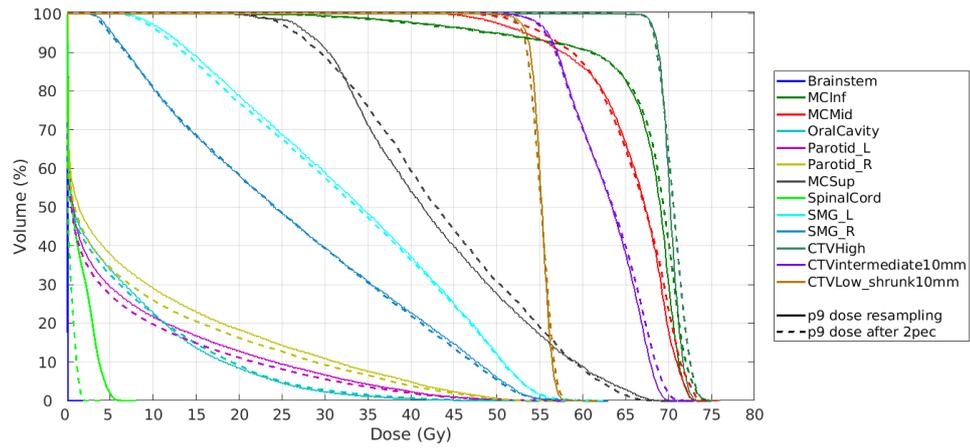


Figure A.9: Dose distribution for patient 9, generated with the SISS method (after the $2pec$ optimization).

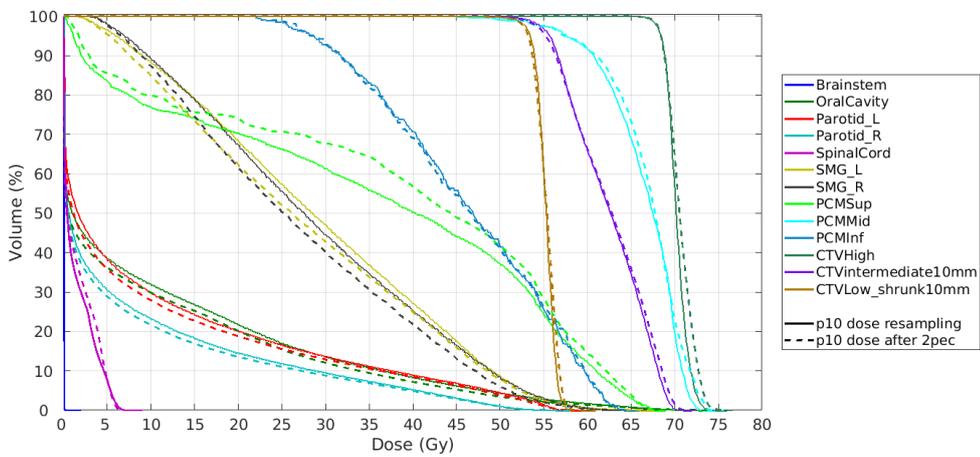


Figure A.10: Dose distribution for patient 10, generated with the SISS method (after the $2pec$ optimization).

A.2. Dose-Volume Histograms after Pareto projection

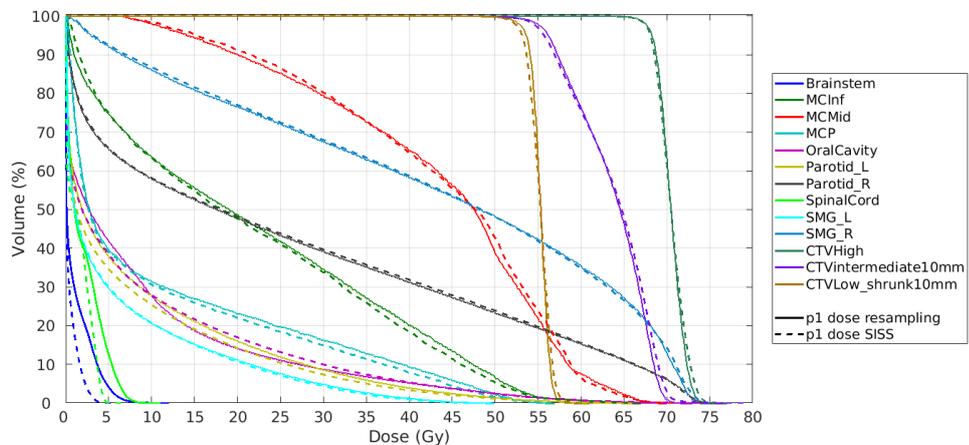


Figure A.11: Dose distribution for patient 1, generated with the SISS method (after the Pareto projection).

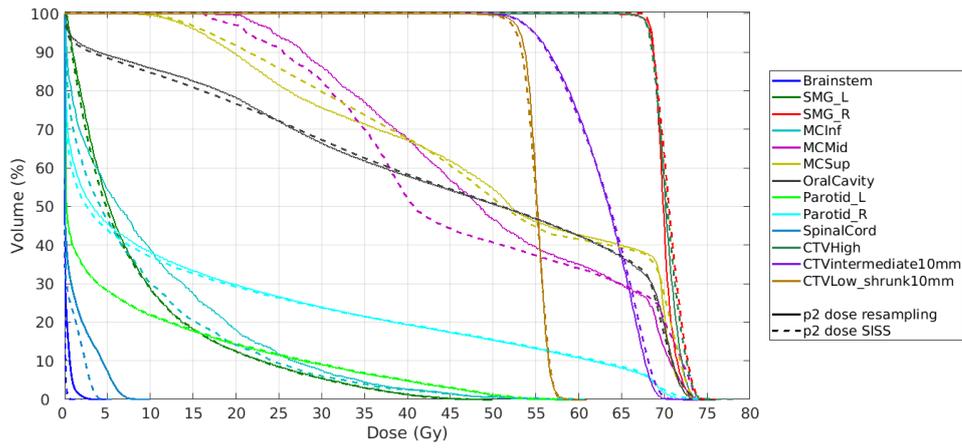


Figure A.12: Dose distribution for patient 2, generated with the SISS method (after the Pareto projection).

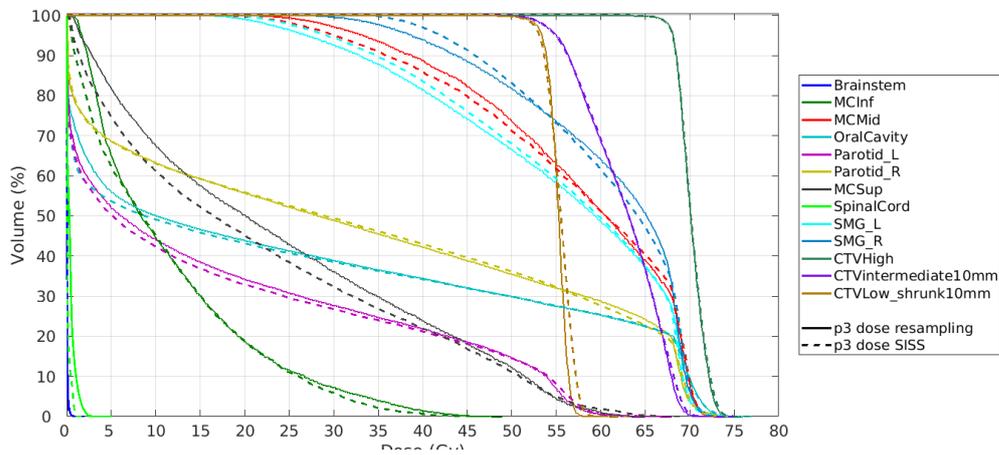


Figure A.13: Dose distribution for patient 3, generated with the SISS method (after the Pareto projection).

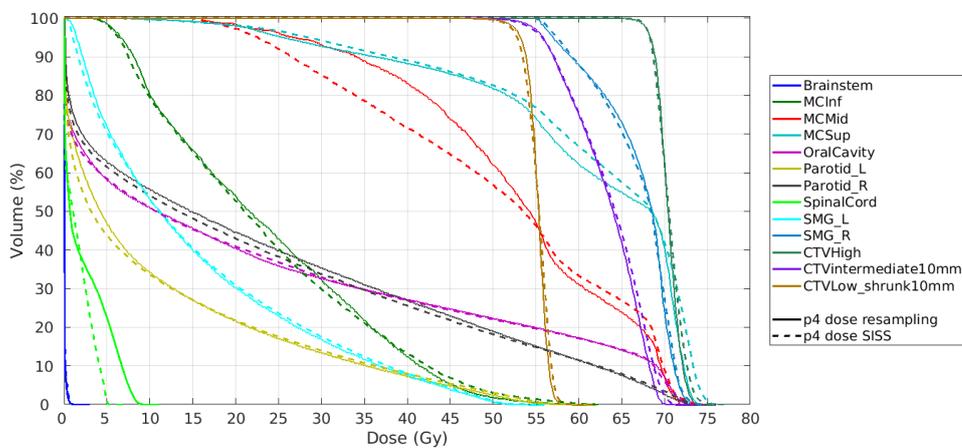


Figure A.14: Dose distribution for patient 4, generated with the SISS method (after the Pareto projection).

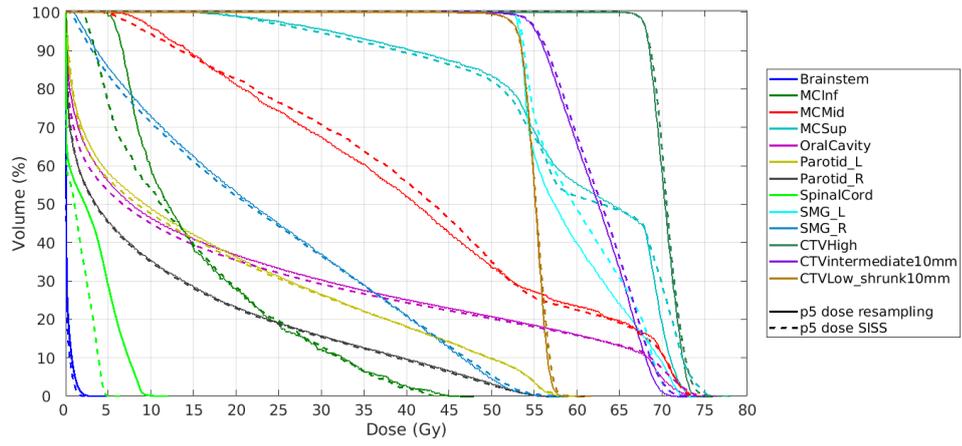


Figure A.15: Dose distribution for patient 5, generated with the SISS method (after the Pareto projection).

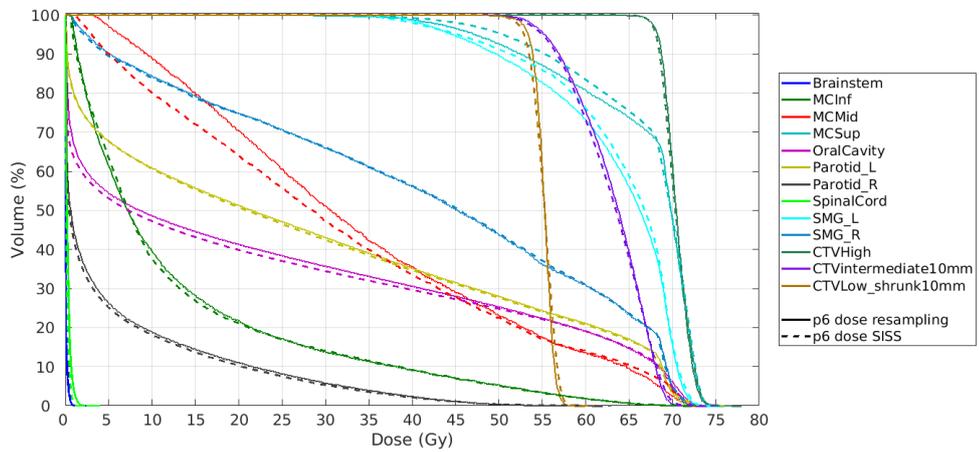


Figure A.16: Dose distribution for patient 6, generated with the SISS method (after the Pareto projection).

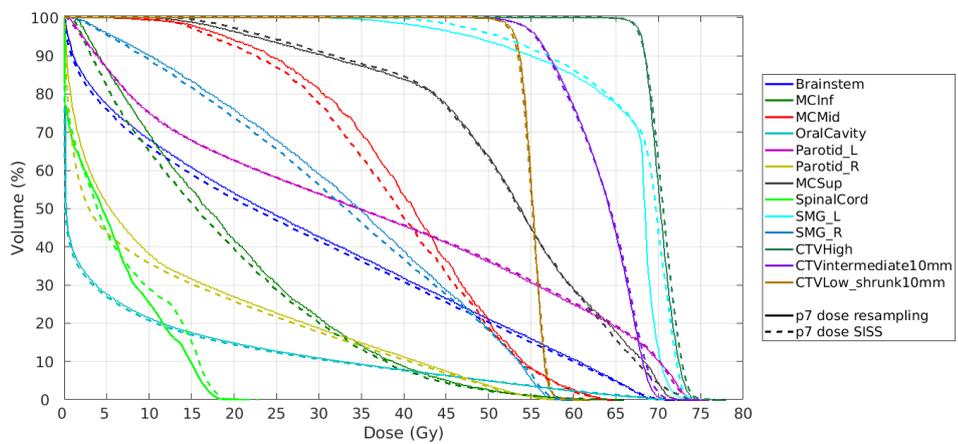


Figure A.17: Dose distribution for patient 7, generated with the SISS method (after the Pareto projection).

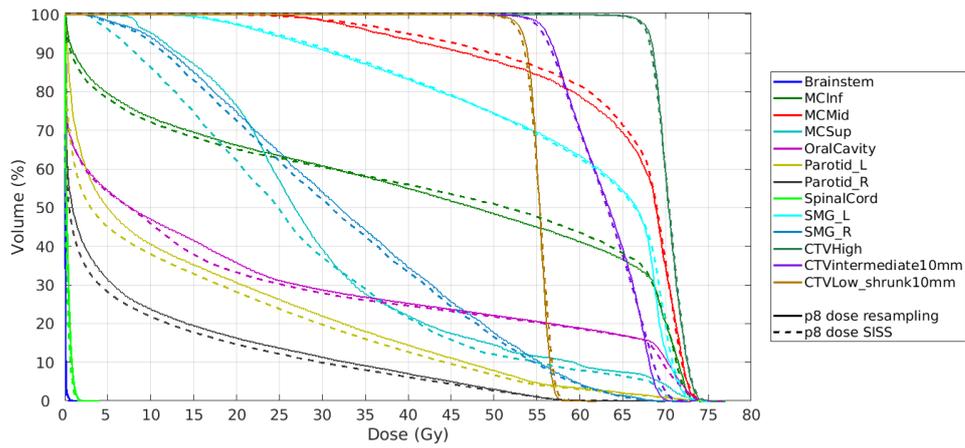


Figure A.18: Dose distribution for patient 8, generated with the SISS method (after the Pareto projection).

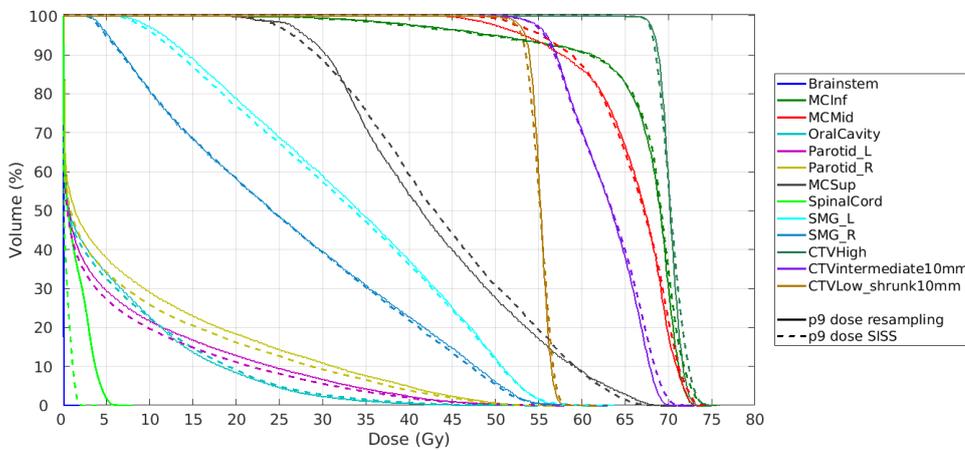


Figure A.19: Dose distribution for patient 9, generated with the SISS method (after the Pareto projection).

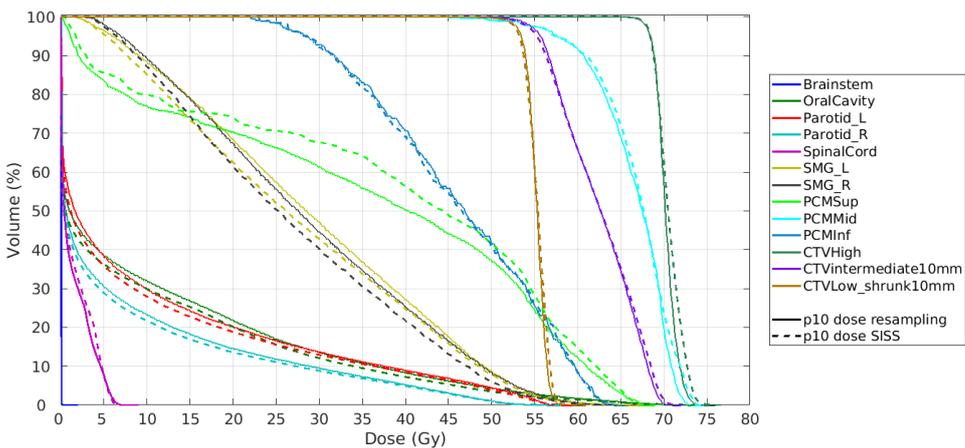


Figure A.20: Dose distribution for patient 10, generated with the SISS method (after the Pareto projection).

References

- [1] A. F. J. G. Janssen. “Beamlet selection and energy layer reduction in IMPT by sparsity induced optimisation”. In: *Bachelor Thesis* (2018).
- [2] H. Sung, J. Ferlay, R. L. Siegel, M. Laversanne, I. Soerjomataram, A. Jemal, and F. Bray. “Global Cancer Statistics 2020: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers in 185 Countries”. In: *CA: A Cancer Journal for Clinicians* 71 (3 2021). ISSN: 0007-9235. DOI: 10.3322/caac.21660.
- [3] M. J. LaRiviere, P.M. G. Santos, C. E. Hill-Kayser, and J. M. Metz. “Proton Therapy”. In: *Hematology/Oncology Clinics of North America* 33 (6 2019). ISSN: 15581977. DOI: 10.1016/j.hoc.2019.08.006.
- [4] Provision Cares Cancer Network. <https://provisionhealthcare.com/about-proton-therapy/advantages-of-proton/>.
- [5] J. Unkelbach and H. Paganetti. *Robust Proton Treatment Planning: Physical and Biological Optimization*. 2018. DOI: 10.1016/j.semradonc.2017.11.005.
- [6] S. van de Water, A. C. Kraan, S. Breedveld, W. Schillemans, D. N. Teguh, H. M. Kooy, T. M. Madden, B. J.M. Heijmen, and M. S. Hoogeman. “Improved efficiency of multi-criteria IMPT treatment planning using iterative resampling of randomly placed pencil beams”. In: *Physics in Medicine and Biology* 58 (19 2013). ISSN: 00319155. DOI: 10.1088/0031-9155/58/19/6969.
- [7] C. Ramirez, V. Kreinovich, and M. Argaez. “Why ℓ_1 Is a Good Approximation to ℓ_0 : A Geometric Explanation”. In: *Journal of Uncertain Systems* 7 (x 2013).
- [8] D. Ge, X. Jiang, and Y. Ye. “A note on the complexity of Lp minimization”. In: *Mathematical Programming* 129 (2 2011). ISSN: 00255610. DOI: 10.1007/s10107-011-0470-2.
- [9] S. Breedveld, P.R.M. Storchi, and B.J.M. Heijmen. “The equivalence of multi-criteria methods for radiotherapy plan optimization.” In: *Physics in medicine and biology* 54 (23 2009). ISSN: 13616560. DOI: 10.1088/0031-9155/54/23/011.
- [10] T. Jagt, S. Breedveld, R. Van Haveren, B. Heijmen, and M. Hoogeman. “An automated planning strategy for near real-time adaptive proton therapy in prostate cancer”. In: *Physics in Medicine and Biology* 63 (13 2018). ISSN: 13616560. DOI: 10.1088/1361-6560/aacaa7.
- [11] T. Z. Jagt, S. Breedveld, R. van Haveren, R. A. Nout, E. Astreinidou, B. J. M. Heijmen, and M. S. Hoogeman. “Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer”. In: *Acta Oncologica* 58 (10 2019). ISSN: 1651226X. DOI: 10.1080/0284186X.2019.1627414.
- [12] S. Breedveld, B. van den Berg, and B. Heijmen. “An interior-point implementation developed and tuned for radiation therapy treatment planning”. In: *Computational Optimization and Applications* 68 (2 2017). ISSN: 15732894. DOI: 10.1007/s10589-017-9919-4.