

# **Indication Specific Multi-Criteria Optimization in Radiotherapy Treatment Planning using Knowledge from previously delivered Treatment Plans**

**Christian Harrer**

Vollständiger Abdruck der von der Fakultät für Physik der Technischen Universität München zur Erlangung des akademischen Grades eines

**Doktors der Naturwissenschaften (Dr. rer. nat.)**

genehmigten Dissertation.

**Vorsitzender:**

Prof. Dr. Martin Zacharias

**Prüfende der Dissertation:**

1. Prof. Dr. Jan J. Wilkens
2. Prof. Dr. Franz Pfeiffer

Die Dissertation wurde am 01.04.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 01.07.2019 angenommen.



---

# **Indication Specific Multi-Criteria Optimization in Radiotherapy Treatment Planning using Knowledge from previously delivered Treatment Plans**

**Dissertation von Christian Harrer**

---



Technische Universität München

2019



# Abstract

Cancer is still one of the biggest challenges for medicine in modern societies. Therapy with ionizing radiation such as x-rays or protons, together with surgery and chemotherapy, has been one of the main treatment options for decades, and even the development of new methods such as immunotherapy is not likely to render it obsolete. The advent of more advanced delivery techniques, such as linear accelerators (linacs), and multi-leaf-collimators (MLCs), for precise shaping of the treatment beam has led to increasingly sophisticated radiation treatment options.

State-of-the-art treatment planning almost always involves the so-called inverse-planning approach: Dosimetric parameters crucial for the clinical outcome, like the dose prescribed to the tumor, or dose limits for important radiation sensitive organs at risk (OARs), are specified beforehand, while complex optimization algorithms try to find machine control parameters which lead to the desired dose distribution in the patient's body. Modern dose delivery techniques, like intensity modulated radiation therapy (IMRT) and volumetric modulated arc therapy (VMAT), are based on an inverse optimization approach.

This inverse optimization problem, however, almost always tries to satisfy highly contradictory optimization objectives, such as maximizing the dose in the tumor while doing minimal damage to the neighboring healthy tissue. The trade-off between these goals can be controlled by various adjustable parameters. In this type of problem, called multi-criteria optimization (MCO), an a-priori definition of a global optimum is not possible any more, up to an infinite number of equally 'optimal' solutions exist, and the final choice is often still left to a human decision-maker.

This thesis focuses on two aspects which could facilitate the MCO radiotherapy treatment planning workflow: First, the prediction of the outcome of the optimization process for a specific parameter setting without having to perform a time-consuming complete optimization. This is done via trilinear dose interpolation using pre-calculated results for specific parameter values. In contrast to existing research, this work is more focused on seeing whether useful predictions can be achieved without being able to make strong assumptions on the mathematical properties of the underlying optimizer. In addition, the optimization process studied here is governed by a small number of highly composite parameters only, each of these parameters controlling a number of important aspects of the optimization process simultaneously. The feasibility of this approach is investigated by comparison of actual result and approximation for a number of selected cranial and spinal cases and two different levels of interpolation resolution.

The second part consists of the application of machine learning models trained on existing optimization results and the underlying variation in the patient geometry for predicting the parameter settings which have the biggest potential impact on changing the optimization result. With this knowledge, a better approximation of the set of possible solutions can be achieved. The user can be provided with a more streamlined planning experience, because identification of the most relevant regions in the parameter space is possible. Comparisons are performed using support-vector-machine (SVM), random-forest (RF) and artificial neural network (ANN) classifiers. Again, a number of spinal and cranial cases are investigated and the feasibility and effectiveness of such a methodology is shown.

Finally, an outlook is given on how, when combined, the results of those two studies can be used to develop an optimized treatment planning workflow, which can be expected to be far more time-efficient than the existing current procedure.



# Zusammenfassung

Die Behandlung von Krebserkrankungen stellt nach wie vor eine der größten Herausforderungen in der modernen Medizin dar. Zusammen mit der chirurgischen Tumorentfernung und der Chemotherapie ist die Behandlung mit ionisierender Strahlung wie Röntgenstrahlen oder Protonen seit Jahrzehnten eine der Haupttherapiemöglichkeiten. Auch die laufende Entwicklung von neuen Therapieoptionen wie der Immunotherapie wird die Bedeutung der Strahlentherapie aller Voraussicht nach nicht schmälern, sondern trägt nur zum Trend zu immer differenzierteren und patientenspezifischeren Behandlungsmethoden bei.

Die Verfügbarkeit von technologisch zunehmend ausgereifteren Bestrahlungsgeräten wie z.B. klinischen Linearbeschleunigern mit Lamellenblende zur exakten Anpassung des Strahlungsfeldes, hat immer anspruchsvollere Möglichkeiten der Strahlenbehandlung ermöglicht. Modernste Bestrahlungsplanung für z.B. intensitätsmodulierte Radiotherapie (IMRT) und volumetrisch-modulierte Rotationsbestrahlung (VMAT) beinhaltet fast immer den sogenannten inversen Planungsansatz: Dosimetrische Parameter, die für das klinische Ergebnis entscheidend sind, wie z.B. die dem Tumor verordnete Strahlendosis oder Dosisgrenzwerte für besonders strahlungsempfindliche Risikoorgane (OARs), werden im Vorfeld spezifiziert. Komplexe computergestützte Optimierungsalgorithmen versuchen, die Maschinensteuerungsparameter, die zu der gewünschten Dosisverteilung im Körper des Patienten führen, zu ermitteln.

Eine besondere Eigenschaft eines derartigen Optimierungsproblems ist die gleichzeitige Verfolgung von teils stark gegensätzlichen Zielen, so setzt z.B. die Maximierung der Dosisabdeckung im Zielvolumen der gleichzeitigen bestmöglichen Schonung des unmittelbar umgebenden gesunden Gewebes technisch bedingt Grenzen.

Bei so einem, auch multikriterielle Optimierung (MCO) genannten Problem, ist in der Regel die a-priori Definition eines globalen Optimums nicht mehr möglich, es existieren teils unbegrenzt viele mathematisch gleichwertige Optima, und die letztendliche Entscheidung muss individuell getroffen werden.

Diese Dissertation konzentriert sich auf zwei Aspekte, die eine auf MCO basierende Bestrahlungsplanung erleichtern können:

Zum einen die Vorhersage des Ergebnisses des Optimierungsprozesses für eine spezifische Wahl der Startparameter, ohne daß jedes mal ein zeitaufwändiger kompletter Optimierungsprozess durchlaufen werden muss. Dies erfolgt hier durch trilineare Dosisinterpolation ausgehend von vorberechneten Ergebnissen für spezifisch ausgewählte Parameterwerte. Im Gegensatz zur bestehenden Forschung konzentriert sich diese Arbeit mehr darauf, zu untersuchen, ob nützliche Vorhersagen erreicht werden können, ohne stark einschränkende Annahmen bezüglich der mathematischen Eigenschaften des zugrunde liegenden Optimierers machen zu müssen. Darüber hinaus wird der hier untersuchte Optimierungsprozess nur durch eine kleine Anzahl von einstellbaren Parametern, die in sich jedoch vielfältige Aspekte des Optimierungsprozesses vereinen, gesteuert, während in vielen bisherigen Ansätzen eine unüberschaubare Zahl von Einzelparametern möglich ist. Die Durchführbarkeit solch einer Strategie wird durch den Vergleich von tatsächlichem Ergebnis und Approximation für eine Reihe ausgewählter kranialer und spinaler Patientengeometrien und trilinearere Interpolation mit zwei verschiedenen Auflösungsstufen gezeigt. Der zweite Teil besteht in der Entwicklung und Anwendung eines maschinellen Lernmodells, das an vorberechneten Optimierungsergebnissen trainiert wird, und unter Berücksichtigung der zugrundeliegenden Variation in der Patientengeometrie die Vorhersage der Parameter-Bereiche, in denen die Optimierungsergebnisse am meisten Variabilität aufweisen, ermöglicht. Hierdurch kann der Planungsprozess zielgerichteter und zeitsparender gestaltet werden.

Die folgenden maschinellen Lernverfahren werden eingesetzt und bezüglich der erreichten Vorhersagequalität verglichen: Support-Vektor-Maschine (SVM), Klassifikation mittels Random-Forest (RF) Entscheidungsbäumen und künstliche neuronale Netze (ANN). Auch hier wird die Effektivität des Verfahrens anhand der oben erwähnten spinalen und kraniellen Testgeometrien gezeigt. Schließlich erfolgt ein Ausblick, wie die Ergebnisse dieser beiden Studien kombiniert werden können, um den Prozess der Bestrahlungsplanung noch einheitlicher und zeiteffizienter zu gestalten.



# Contents

<b>Abstract</b>	<b>v</b>
<b>Zusammenfassung</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Introduction . . . . .	2
<b>2 Theoretical background</b>	<b>5</b>
2.1 Radiotherapy and treatment planning . . . . .	6
2.1.1 Radiobiology . . . . .	6
2.1.2 Radiation therapy . . . . .	6
Conformal radiation therapy . . . . .	8
Intensity-modulated radiation therapy . . . . .	8
2.1.3 RT treatment planning . . . . .	9
Dose-Volume-Histogram (DVH) . . . . .	9
Dosimetric planning criteria . . . . .	9
Additional treatment plan evaluation criteria used in this thesis . . . . .	10
2.2 Introduction to multi-criteria optimization . . . . .	11
2.2.1 Problem formulation . . . . .	11
2.2.2 Pareto surfaces, convexity . . . . .	11
2.2.3 Solution strategies . . . . .	12
Weighted sum method . . . . .	12
$\epsilon$ - constraint method . . . . .	13
Lexicographic order method . . . . .	13
2.3 MCO in RT treatment planning . . . . .	14
2.4 Continuous function value approximation in cubic affine parameter spaces . . . . .	14
2.4.1 Approximation with 7 anchor points and its limitations . . . . .	15
2.4.2 Trilinear interpolation with 8 anchor points . . . . .	16
2.4.3 Subdivision of feature space, 27 anchor points . . . . .	18
2.4.4 Adaptive-grid trilinear interpolation . . . . .	19
2.5 Mathematical characterization of patient geometry . . . . .	20
2.5.1 Overlap-volume-histogram (OVH) . . . . .	20
2.5.2 Inverse overlap-volume-histogram . . . . .	20
2.5.3 Concavity/convexity of anatomical structures . . . . .	22
2.6 Artificial intelligence and machine learning . . . . .	23
2.6.1 The supervised learning problem . . . . .	23
2.6.2 Regression vs classification . . . . .	24
2.6.3 Important concepts in model training and assessment . . . . .	24
Data preparation: regularization and oversampling . . . . .	24

	Binary classification terminology . . . . .	24
	Sensitivity/Specificity . . . . .	25
	Precision/Recall/Accuracy . . . . .	25
	ROC curves/AUC . . . . .	25
2.6.4	Machine learning algorithms used in this thesis . . . . .	26
	Support-vector-machines (SVM) . . . . .	26
	Random-forest-classifier (RF) . . . . .	26
	Artificial neural networks (ANN) . . . . .	27
<b>3</b>	<b>The VMAT planning Element from an MCO perspective</b>	<b>31</b>
3.1	The Brainlab Elements VMAT planning system (Cranial SRS/Spine SRS) . . . .	32
3.1.1	The Brainlab Elements VMAT planning system . . . . .	32
	General overview . . . . .	32
3.1.2	Configuration of planning objectives and constraints . . . . .	33
3.1.3	High-level composite optimization modification sliders and their func- tionality . . . . .	34
	The target-OAR-weighting control . . . . .	34
	The normal tissue sparing control . . . . .	35
	The dose modulation control . . . . .	35
	Slider range and discretization . . . . .	35
3.2	Elements SRS VMAT planning from an MCO perspective . . . . .	36
3.2.1	Weighting-, normal tissue sparing- and dose modulation slider values as base of a 3D-parameter space(3DPS) . . . . .	36
3.2.2	Analysis of the solution space generated in the 3DPS by means of the Brainlab VMAT optimizer . . . . .	36
3.3	Cranial data . . . . .	38
3.3.1	PTV coverage . . . . .	38
3.3.2	PTV maximum dose . . . . .	39
3.3.3	PTV minimum dose . . . . .	40
3.3.4	PTV mean dose . . . . .	41
3.3.5	Dose in OAR constraint point . . . . .	41
3.3.6	OAR mean dose . . . . .	42
3.3.7	OAR minimum dose . . . . .	42
3.3.8	OAR maximum dose . . . . .	43
3.3.9	Gradient index (GI) . . . . .	43
3.3.10	Conformity index (CI) . . . . .	44
3.3.11	Monitor units (MU) . . . . .	44
3.3.12	Double resolution . . . . .	45
3.4	Spinal data . . . . .	46
3.4.1	PTV coverage . . . . .	46
3.4.2	PTV maximum dose . . . . .	47
3.4.3	PTV minimum dose . . . . .	48
3.4.4	PTV mean dose . . . . .	48
3.4.5	Dose in OAR constraint point . . . . .	49
3.4.6	OAR mean dose . . . . .	49

3.4.7	OAR minimum dose . . . . .	50
3.4.8	OAR maximum dose . . . . .	50
3.4.9	Gradient index (GI) . . . . .	51
3.4.10	Conformity index (CI) . . . . .	51
3.4.11	Monitor units (MU) . . . . .	52
3.5	Approximation of Pareto surfaces by trilinear dose interpolation . . . . .	53
3.5.1	Considerations concerning the interpolation resolution . . . . .	54
3.5.2	Mathematical behavior of characteristic DVH values with respect to trilinear interpolation . . . . .	56
	Mean dose . . . . .	56
	Max dose . . . . .	57
	Min dose . . . . .	58
	General dose/volume values, e.g. 99% dose . . . . .	58
	Boundedness of min and max dose . . . . .	58
3.5.3	Trilinear interpolation with 8 anchor points . . . . .	59
	PTV coverage cranial . . . . .	59
	PTV coverage spine . . . . .	60
	OAR mean dose brainstem . . . . .	61
	OAR mean dose spinal canal . . . . .	62
	OAR constraint dose brainstem . . . . .	63
	OAR constraint dose spinal canal . . . . .	64
3.5.4	Trilinear interpolation with 27 anchor points . . . . .	65
	PTV coverage cranial . . . . .	66
	PTV coverage spine . . . . .	67
	OAR mean dose brainstem . . . . .	68
	OAR mean dose spinal canal . . . . .	69
	OAR constraint dose brainstem . . . . .	70
	OAR constraint dose spinal canal . . . . .	71
3.5.5	Limitations and advantages . . . . .	72
<b>4</b>	<b>Exploration of the solution space via machine learning</b>	<b>75</b>
4.1	Solution space exploration via machine learning: Motivation . . . . .	76
4.1.1	Regression vs classification in VMAT treatment planning . . . . .	76
4.2	Identification of relevant sections of the parameter space by prediction of achievable value ranges . . . . .	77
4.2.1	Feature matrix and target vector generation . . . . .	80
4.2.2	Generation of geometry vector G . . . . .	80
4.2.3	Analysis of cranial cases with brainstem as prioritized OAR . . . . .	83
	OAR constraint point dose w.r.t. weighting slider . . . . .	83
	PTV rel. coverage w.r.t. weighting slider . . . . .	84
	PTV 1% dose w.r.t. weighting slider . . . . .	85
	Gradient index (GI) w.r.t. normal-tissue-sparing slider . . . . .	87
	Monitor units (MU) w.r.t. modulation slider . . . . .	89
4.2.4	Analysis of spine cases with spinal canal as prioritized OAR . . . . .	90
	OAR constraint point dose w.r.t. weighting slider . . . . .	90

	PTV rel. coverage w.r.t. weighting slider . . . . .	92
	PTV 1% dose w.r.t. weighting slider . . . . .	93
	Gradient index (GI) w.r.t. normal-tissue-sparing slider . . . . .	94
	Monitor units (MU) w.r.t. modulation slider . . . . .	96
4.2.5	Results and discussion . . . . .	96
<b>5</b>	<b>Outlook &amp; Conclusion</b>	<b>101</b>
5.1	Combination of ML models . . . . .	102
5.1.1	Outline and structure of algorithm . . . . .	102
5.1.2	Additional enhancements . . . . .	103
	Use of big data/cloud-based data management . . . . .	103
	Learning of user preferences . . . . .	104
5.2	Summary & conclusion . . . . .	105
<b>6</b>	<b>Appendix</b>	<b>109</b>
6.1	Technical Background . . . . .	110
6.1.1	Geometry of test patient data . . . . .	110
6.1.2	Data generation and processing/software framework . . . . .	111
	<b>Bibliography</b>	<b>113</b>
	<b>List of Figures</b>	<b>121</b>
	<b>List of Tables</b>	<b>127</b>
	<b>List of Publications</b>	<b>129</b>
	<b>Acknowledgements</b>	<b>131</b>

**Chapter** **1**  
**Introduction**

## 1.1 Introduction

Cancer is still one of the leading causes of death in modern societies, and recent studies show that in western countries, cancer is even likely to replace heart disease as main cause of death within the next couple of years [29].

Therapy with ionizing radiation such as x-rays or protons, together with surgery and chemotherapy, has been one of the main treatment options for decades [3], and even the development of new methods such as immunotherapy is not likely to render it obsolete[80]. The advent of more advanced delivery techniques such as linear accelerators (linacs) and multi-leaf-collimators (MLCs) for precise shaping of the treatment beam has lead to increasingly sophisticated radiation treatment options[9].

In recent years, radiotherapy treatment planning in the clinical practice has changed tremendously. The cumbersome forward planning process, where the dosimetrist had to manually change treatment shapes and leaf positions, evaluate the resulting dose distribution, and iteratively change parameters again and again in order to achieve an acceptable treatment plan, is being abandoned more and more. State-of-the-art treatment planning almost always involves the so-called inverse-planning approach: Dosimetric parameters crucial for the clinical outcome, like the dose prescribed to the tumor or dose limits for important radiation sensitive organs at risk (OARs) are specified beforehand, while complex computerized optimization algorithms try to find machine parameters like leaf positions, table and gantry angles, which yield the desired dose distribution[35] [85]. Modern dose delivery techniques like intensity modulated radiation therapy (IMRT) and volumetric modulated arc therapy (VMAT) are based on the inverse optimization approach.

This inverse optimization problem, however, almost always consists of highly contradictory optimization objectives: Creating a homogeneously high dose inside the planning target volume (PTV) is very likely to inadvertently also deliver an unwanted radiation dose to OARs in the proximity to the tumor site. If, on the other hand, in order to keep toxic side effects at bay, the emphasis during the optimization process is placed on healthy tissue sparing, the chances of depositing the desired radiation dose in the complete target volume will be lowered as well.

The above type of optimization problem, where a trade-off between contradicting objectives has to be achieved, is generally known as multi-criteria optimization (MCO), and has been widely analyzed in operations research. A single global optimum does not exist any more and is replaced by an (often infinite) number of equally optimal (also called Pareto-optimal) solutions[24]. The term Pareto-optimal is credited to the Italian economist Vilfredo Pareto (1848-1923) and describes a state in certain optimization problems, where no single objective value can be improved without simultaneously worsening another, the set of all Pareto-optimal solutions is called Pareto surface[5]. In terms of radiotherapy treatment planning, this means that the operator is left with the task of deciding which of these Pareto-optimal plans is most likely to lead to the desired clinical outcome.

Since optimization of a radiosurgery treatment plan can take anywhere between a minute and a couple of hours, depending on the complexity of the plan and the available computer hardware, navigating the Pareto surface in real time without pre-calculated results is not practicable. The most common approach for dealing with this problem is to calculate a number of representative Pareto-optimal optimization results [15] and to provide the user with a navigable approximation of the Pareto surface[14] [49]. The feasibility of this approach is

vastly depending on the formulation of the optimization problem and the numerical algorithms applied to it. One method for tackling this challenging task is to formulate the problem as a convex optimization problem. This method has some tremendous advantages, like the fact that the theory of convex optimization problems is mathematically well understood, very fast off-the-shelf optimization algorithms exist, and any found optimum can be guaranteed to be a global optimum[17]. However, when applied to radiotherapy treatment planning, there are also a couple of downsides to this approach:

The concept of convex optimization does not really agree with the realistic circumstances of radiation therapy. Common optimization objectives, like dose volume constraints, break convexity and have to be re-formulated or approximated instead [32]. Optimization is often performed on dose fluences, and the result has to be transformed into a deliverable treatment plan in a post processing step [20]. Even with very efficient leaf-sequencing algorithms, the found optimal fluence pattern is not necessarily the one which is finally delivered to the patient. In commercially available solutions dose metrics like the equivalent uniform dose (EUD)[60] or the maximum dose delivered to a specific organ at risk can be directly specified as hard optimization constraints via slider control tools in the user interface. This allows for a vast number of variations in the optimization result and a highly customized dosimetric outcome. However, as soon as the number of organs at risk increases, modifications of one slider can severely impair the allowed range of the other, and the suitable slider positions to achieve the desired outcome are not always obvious.

The research presented in this work is based on the VMAT optimization algorithm used in the Brainlab Elements Cranial and Brainlab Elements Spine radiosurgery treatment planning software. This system follows a different approach: Instead of providing individual controls for each specific dosimetric criterion, the main optimization process is controlled by means of three main slider controls only. Each of these slider controls influences the optimization process in a number of composite non-trivial aspects, grouped into management of the OAR sparing vs PTV prescription fulfillment, general sparing of non-tumorous healthy tissue (even though it might be not explicitly delineated as organ-at-risk), and finally, the degree of dose modulation contained in the final treatment plan. The latter refers to features such as the size and complexity of allowed treatment beam-shapes and the modulation of the amount of radiation dose delivered per second over the time of the treatment.

The main purpose of this work is the exploration of this specific optimization approach using a low number of composite optimization objectives under the scope of multi-criteria optimization in cranial and spinal radiosurgery applications with the goal of improving the current treatment planning workflow in terms of speed and efficiency. This thesis is structured as follows:

First, after a short introduction to radiotherapy and the relevant treatment planning and delivery methods, all the theoretical concepts relevant for this work are introduced. These consist of:

- An introduction to the mathematical background of multi-criteria optimization and important solution methods
- Presentation of selected methods for approximation of function values in a multi-dimensional parameter space
- A mathematical method for describing the mutual spatial relationship of objects in a discretized 3-D space

- A short introduction to machine learning and important artificial-intelligence (AI) algorithms
- An overview of the VMAT planning capabilities of the optimizer used in this study

The main part focuses on two aspects which could facilitate the MCO treatment planning workflow:

First, the prediction of the outcome of the optimization process for a specific position of each control slider without having to perform a time-consuming complete optimization for each parameter configuration. This is done via approximation using pre-calculated results for specific slider positions. A lot of work has been done in that area for convex optimization problems and single dosimetric value objectives. This work, however, is more focused on seeing whether useful predictions can be also achieved without being able to make strong assumptions on the mathematical properties of the underlying optimizer. The feasibility of this approach is investigated by comparison of actual result and approximation for a number of selected cranial and spinal cases.

The second part consists in the application of a machine learning model trained on existing optimization results for predicting the slider positions which have the biggest potential impact on changing the optimization result. With this knowledge, a better approximation of the Pareto surface might be achieved and the user could be provided with a more streamlined planning experience. Again, a number of spinal and cranial cases are investigated.

Finally, an outlook is given on how, in combination, the results of those two studies could be combined to develop a streamlined treatment planning workflow which can be expected to be far more time-efficient than the existing current procedure.



Chapter **2**

Theoretical background

## 2.1 Radiotherapy and treatment planning

### 2.1.1 Radiobiology

Ionizing radiation consisting of particles such as photons, electrons or protons can seriously damage or even kill living tissue by two main effects:

- Damage caused by direct hits on the cell's DNA. Even though the DNA takes up only a tiny part of a cell's total volume, there is a slight chance that an ionizing particle directly hits the DNA helix contained in the nucleus. Biological cells possess sophisticated DNA repair mechanisms, but there is a slight chance that the damage surpasses these self-repair-capabilities, or only partial repair is possible. In severe cases this can lead to self-induced controlled cell death (apoptosis) or an inability to perform successful cell division (mitosis).
- Damage caused indirectly by ionizing events in the cell body. This is the predominant damage mechanism caused by ionizing radiation. The particles can change the chemical structure of the cell, especially creation of free radicals such as oxygen atoms can cause ionization and oxidation damage to the DNA. The consequences are comparable to direct DNA hits: if the damage cannot be repaired, the cell dies or is unable to divide.


During their life cycle cells are susceptible to radiation damage in a varying degree, but fast-dividing cells such as tumor cells are generally more vulnerable due to their higher rate of mitosis and impaired repair mechanisms. In radiation therapy this effect is used for *fractionated irradiation*: The tumor is irradiated in a number of consecutive fractions, giving healthy tissue time to repair possible damage and increasing the chance of hitting a cancer cell in a vulnerable state.

There exists a lot of literature dealing with radiobiology, a good overview can be found in *Hall*[34].

### 2.1.2 Radiation therapy


Radiation therapy (RT) has been used for more than a century to treat varying illnesses[30][51][67]. Nowadays one of the most important applications is treatment of cancers, especially in body regions which are very hard or even impossible to reach by surgical intervention. In other cases it is used as adjuvant therapy in combination with conventional surgery and chemotherapy. The application of radiation therapy can be distinguished in *external beam radiotherapy* where the irradiating source is placed outside of the body, and *brachytherapy*, where the source is temporarily directly placed in the tissue which is to be irradiated. This thesis covers topics related to external photon-beam radiotherapy only.

While some modern treatment units for external photon beam therapy use decaying radioactive isotopes such as  $Co^{60}$  for a photon source (such as the *Gamma Knife*), most of the time a clinical *linear accelerator* (figure 2.1) is used for producing the radiation. Electrons are accelerated to an energy of typically 4-18 MV and hit a tungsten target which releases photon bremsstrahlung. The photon beam is usually focused by a two pairs of primary collimators also called *jaws* while the final shape of the treatment beam is achieved via two stacks of movable tungsten leaves called *multileaf-collimator (MLC)*. The movable treatment unit containing the linear accelerator



**Figure 2.1:** Modern clinical Linac with MLC, patient positioning system and treatment couch [1]


and treatment head is called *gantry*. The Linac head also contains an ionization chamber where the delivery dose-rate in  $\frac{(MU)}{s}$  is measured (*MU* stands for *Monitor Unit*). For dose calculations in the treatment planning process a calibrated chamber in a water tank is used to measure the Linac calibration in  $\frac{Gy}{100MU}$  at a specific calibration-depth and -field size. Figure 2.2 shows typical relative-depth-dose curves for a 6MV photon beam and different field sizes in water. After an initial dose build-up the absorbed dose gradually declines with increasing depth.



**Figure 2.2:** Normalized relative depth dose (PDD) of a 6MV photon beam for 10x10mm(green), 60x60mm(red) and 100x100mm(blue) MLC square field size

Further information on the physics of radiation therapy can be e.g. found in *Khan*[45]. Treatment delivery is usually done from a number of different couch and gantry positions in order to minimize radiation exposure of healthy tissue. Delivery techniques can be categorized into

- *Conformal Forward-Planning* delivery methods



**Figure 2.3:** Schematic overview of target-outline (red), primary jaws (gray) and MLC leaves shaping the cross-section of the treatment beam (in reality, MLCs can have up to 120 leaf-pairs).

- *Intensity-modulated Inverse-Planning* delivery methods

### Conformal radiation therapy

During purely conformal treatments, the MLC leaves adapt to the projected outline of the target volume at the specific couch-, gantry- and collimator angle as seen in figure 2.3. The term 'forward-planning' refers to the corresponding treatment planning method. Gantry- and collimator angles are selected manually and the expected dose distribution in the tissue is calculated. In an iterative process angle adjustments and dose calculations are repeated until the user is satisfied with the resulting dose distribution.

Delivery can be either performed in static fields where gantry, MLC and couch remain fixed when the photon beam is enabled, or via dynamic conformal treatments, where the gantry moves around the treatment site in an arc and the MLC leaves dynamically adapt to the shape of the target.

### Intensity-modulated radiation therapy

*Intensity-modulated radiation treatments* (IMRT) and *volumetric modulated arc treatments* (VMAT) are characterized by the *inverse-planning* method. In this case, the operator first specifies the desired dose distribution and each MLC aperture is discretized into fluence *beamlets*. An optimization algorithm then tries to achieve this dose distribution by assigning each beamlet a certain intensity and determining deliverable MLC apertures realizing these beamlet intensities.

Figure 2.4 shows a sample IMRT treatment field and the discretization into beamlets with varying fluence intensity. A darker beamlet color means that the specific beamlet contributes more dose than a lighter colored one. In practice this is achieved by closing leaves over lower-intensity beamlets during the treatment, leading to MLC apertures which are not necessarily conformal to the target outline any longer. IMRT treatments are usually delivered from a number of different couch- and gantry angles which do not change while the beam is switched on. A lot of literature is available on IMRT, detailed information is provided e.g. by *Webb*[84][85] and *Bortfeld*[6].

During VMAT treatments there is an additional gantry movement during treatment and the gantry speed and MU delivery rate can vary, which adds additional degrees of freedom. While dosimetric outcomes of IMRT and VMAT are comparable in terms of tumor control and mitigation of radiation-induced toxicity in healthy tissue and both are generally superior to purely conformal delivery, VMAT treatments are more time efficient and have become state-of-the-art

for many treatment sites[65][77]. This thesis is covering topics concerning VMAT treatment planning exclusively.




Figure 2.4: Aperture of an IMRT beam with beamlet discretization and fluence

### 2.1.3 RT treatment planning

The treatment planning process generally consists of several planning steps. Image sets acquired by different imaging modalities are used to create a 3D tissue model where organ structures are delineated. The tissue model and all structures are discretized into tiny volume elements called *voxels*. The *planning target volume* (PTV) consists of the actual visible tumor tissue plus margins which take into account biological and technical uncertainties. Such uncertainties can arise due to positioning errors during treatment or due to issues when delineating the tumor in the image sets. Tissue areas and organs which are of special concern for radiation treatment, because toxic side-effects would occur if they received excessive dose, are also delineated as *organs-at-risk* (OAR).

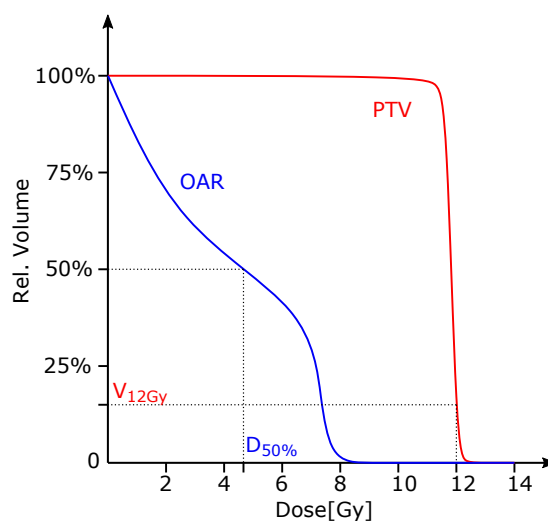
#### Dose-Volume-Histogram (DVH)

Once a 3D dose distribution is present after the optimization process, the dosimetric outcome for each organ can be reviewed. This is done by binning the radiation dose received by the voxels of each organ and plotting a cumulative histogram of the values which is called *dose-volume-histogram* (DVH). A sample DVH is shown in figure 2.5.

#### Dosimetric planning criteria

Desired dosimetric planning criteria can be specified by the user and are internally processed by the optimizer and used as optimization constraints. Such constraints can be

- *Absolute values*, such as the maximum, minimum or mean dose which should not be exceeded or which should at least be met, here denoted as  $D_{max}$ ,  $D_{min}$ ,  $D_{mean}$
- *Dose-at-volume-constraints*, meaning that a certain relative or absolute volumetric part of a structure shall at least receive or not receive more than a specific dose. This is written as e.g.  $D_{99\%}$  for relative volumes or  $D_{10ccm}$  for absolute volumes
- *Volume-at-dose-constraints*, referring to the volumetric part (relative or absolute) of a structure receiving a certain amount of dose. This is written as e.g.  $V_{10}$  for the volume which receives at least 10Gy.



**Figure 2.5:** Dose-Volume-Histogram (DVH) graphs for two objects: PTV (red) and OAR (blue)

The terms defined above are also illustrated in figure 2.5. In the Brainlab Elements treatment planning software a list of all organs taken into account for optimization and the desired specific DVH constraints, as well as desired dose coverage volumes for the PTV can be specified.

### Additional treatment plan evaluation criteria used in this thesis

In the studies conducted for this thesis the following additional plan evaluation criteria are examined:

- The *inverse conformity index* (CI). This widely used criterion was introduced by *Paddick*[62] and provides information how well the region receiving the prescribed dose correlates with the actual PTV. An ideal treatment plan would have  $CI = 1$ , in reality values greater than one are common.
- The *gradient index* (GI). This criterion provides information on the steepness of the dose fall-off outside the PTV. A steep fall-off ensures better sparing of the healthy tissue surrounding the PTV. A definition of the GI can be found at *Paddick*[63]. The CI value decreases with increased dose fall-off outside the PTV, thus treatment plans with low GI are usually preferred.
- The *modulation complexity score* (MCS) is a measure for gauging the delivery- and QA-complexity of a treatment plan. This takes into account number, sizes and shape (area and perimeter) of the MLC apertures, as well as the number of MUs to deliver. Such a measure was first introduced by *McNiven et al*[56] and later extended to VMAT delivery by *Younge et al*[88] and *Masi et al*[52]. Usually plans with low complexity are preferred because they promise a faster delivery and a more exact real-world reproduction of the planned dose distribution.
- The amount of *Monitor Units* (MU) needed to deliver the treatment plan. Here generally lower values are preferred since a lower MU value means faster delivery and thus reduced treatment time.

## 2.2 Introduction to multi-criteria optimization

Many real-world optimization problems deal with fulfilling conflicting criteria. A simple example is the construction of an internal combustion engine, where the wish to maximize power output is conflicting with the desire for high fuel-efficiency. In general, efficient production of any goods where high quality and low cost are desired impose conflicting optimization goals, since higher quality often comes at the cost of an expensive manufacturing process and high-quality (and thus more expensive) raw materials.

In these cases a single obvious optimal solution does not exist any more, there is always a trade-off to be made between contradicting objectives and the final decision on which solution is optimal in each specific setting is not uniquely defined [22]. The following sections are intended to give a basic understanding of the mathematical concept of multi-criteria optimization (MCO), introduce important terms and definitions and to present some basic solution strategies.

### 2.2.1 Problem formulation

A general MCO optimization problem can be formulated as follows [76], [24]:

$$\min_{x \in X} \{f_1(x), f_2(x), \dots, f_k(x)\} \quad (2.2.1)$$

where  $X$  is called *feasible set* and the  $f_i : \mathbf{R}^n \rightarrow \mathbf{R}, i \in (1, \dots, k), k \geq 2$  are conflicting *objective functions*. The feasible set consists of all  $x \in \mathbf{R}^n$  satisfying certain constraints imposed on allowed solutions, which can be user defined or caused by technical or physical limitations due to the nature of the underlying problem. In mathematical terms, this can be described as [5]

$$c_j(x) \leq 0, \quad (2.2.2)$$

$$Ax \leq b \quad (2.2.3)$$

Equation 2.2.2 denotes user-defined optimization constraints, while equation 2.2.3 limits the solution space to the technically and physically feasible set.

The objective functions  $f_i$  transform the feasible set  $X$  into the *objective space*  $f(X)$ . A solution  $x \in X$  is called *weakly Pareto optimal* if the vector  $f(x) \in \mathbf{R}^k := (f_1(x) f_2(x) \dots f_k(x))^T$  fulfills the following condition:

There exists no other  $x' \in X$  such that  $f_i(x') < f_i(x) \forall i = 1, \dots, k$ .

A solution  $x \in X$  is called *Pareto optimal* if the vector  $f(x) \in \mathbf{R}^k := (f_1(x) f_2(x) \dots f_k(x))^T$  fulfills the following condition:

There exists no other  $x' \in X$  such that  $f_i(x') \leq f_i(x) \forall i = 1, \dots, k$  and  $f_j(x') < f_j(x)$  for at least one index  $j$ .

### 2.2.2 Pareto surfaces, convexity

A set  $X \subset \mathbf{R}^n$  is called *convex* if for any  $x, y \in X, \alpha \in [0, 1]$  it holds that

$$\alpha x + (1 - \alpha) y \in X \quad (2.2.4)$$

A function  $f(x)$  is called *convex* on  $X$  if for any  $x, y \in X, \alpha \in [0, 1]$  it holds that

$$f(\alpha x + (1 - \alpha)y) \leq \alpha f(x) + (1 - \alpha)f(y) \quad (2.2.5)$$

The fact that generally the cost function components  $f_i$  of a MCO problem are contradicting implies that the exact meaning of 'min' in 2.2.1 is not clearly defined. Depending on the definition of 'min' are possibly several or even infinitely many *Pareto optimal* solutions. In the case of an infinite amount of *Pareto optimal* solutions  $x_p$ , they span a  $k$  dimensional hyperplane in the objective space which is called *Pareto front* or *Pareto surface*.

Figure 2.6 shows a simple two-dimensional example for the problem

$$\min_{x \in X} \{f_1(x), f_2(x)\} \quad (2.2.6)$$

The figure shows the two dimensional objective space defined by  $f_1, f_2$ .  $f(X)$  is the image of the feasible set  $X$  in the objective space. The image of the Pareto optimal points in  $X$  is shown in red. The dashed lines show that in the red area no further decrease in  $f_1(x)$  can be achieved without simultaneously increasing the value of  $f_2(x)$  and vice versa.

In the example the  $f_i$  are convex functions and  $X$  is a convex set, rendering the image  $f(X)$  a convex set in the objective space. Convexity is a very important and desirable feature when it comes to exploration of the Pareto front, however, many real-life problems are not inherently convex.




Figure 2.6: 2D example of Pareto optimality

### 2.2.3 Solution strategies

There exist numerous methods for dealing with the contradicting nature of the single components of the objective function vector  $f$ . In the following a couple of these strategies will be briefly outlined.

#### Weighted sum method

The multi-objective optimization problem 2.2.1 can be transformed into a single-objective optimization problem by replacing the objective function vector with a weighted sum of its



components as follows:

$$\min_{x \in X} \sum_k \lambda_k f_k(x), \lambda_k \in \mathbf{R}_0^+ \quad (2.2.7)$$

The weights  $\lambda_k$  can be used to put more emphasis on single components of the objective function and model the trade-off between single objectives. The new optimization problem 2.2.7 can be solved with a suitable optimizer, and it can be shown[24] that optimal solutions of 2.2.7 are always weakly Pareto optimal or (in case of strictly non-negative weights) Pareto optimal.

### $\epsilon$ - constraint method

Another widely used option is the  $\epsilon$ -constraint method. Optimization is only performed with respect to a single component  $f_j$  of the objective function, the remaining components are transformed into constraints:

$$\begin{aligned} & \min_{x \in X} f_j(x), \\ & f_i(x) \leq \epsilon_i \quad i = 1, \dots, k \quad i \neq j \\ & \epsilon_i \in \mathbf{R} \end{aligned} \quad (2.2.8)$$

It can be shown [24] that solutions of 2.2.8 are at least weakly Pareto optimal.

### Lexicographic order method

Another way of concretizing 2.2.1 is to define a lexicographic order  $<_{lex}$  on the  $f_k$ . 2.2.1 can then be re-written as follows:

$$\text{lexmin}_{x \in X} \{f_1(x), f_2(x), \dots, f_k(x)\} \quad (2.2.9)$$

A feasible solution  $x \in X$  is called *lexicographically optimal* if

$$f(x) \leq_{lex} f(x') \forall x' \in X, x' \neq x, \quad (2.2.10)$$

which means that

$$f(x_1) \leq_{lex} f(x_2) \implies f_k^*(x_1) < f_k^*(x_2) \vee f(x_1) = f(x_2). \quad (2.2.11)$$

$k^*$  denotes the smallest index  $k$  where  $f_k(x_1) \neq f_k(x_2)$ .

While scalarization methods like the weighted sum or the  $\epsilon$ -constraint method try to model a trade-off between contradicting objectives, the lexicographic order method implies that optimization with respect to  $f_k$  is only performed, if several competing solutions with respect to  $f_1, \dots, f_{k-1}$  have been found. While this means that a trade-off between competing  $f_i$  is not feasible, the strict order of priorities allows for an efficient iterative optimization process.

All three methods shown here have their own advantages and disadvantages when applied to specific real-life optimization problems. However, the choice of the weights  $\lambda_i$ , the constraints  $\epsilon_i$  or the definition of a lexicographic order on the  $f_k$  is still a highly individual process. If a number of equally Pareto optimal solutions exist, the final choice of the actual optimal solution of a MCO problem remains dependent on a human decision maker and the main task of MCO is to facilitate the decision maker's choice.

## 2.3 MCO in RT treatment planning

The optimization problem encountered in the inverse radiotherapy treatment planning approach presented in the previous chapter is a good example for an MCO task. Dose coverage of the PTV is often not possible without unwanted irradiation of OARs and vice versa. If the DVH constraints chosen and used as optimization objectives for each organ can be met, a high number of Pareto-optimal solutions exist.

In this case, a feasible strategy is to present these solutions in a graphical way and leave the final choice to the planner. This is usually done by providing slider controls where the value of each single DVH constraint can be adjusted in such a way, that when one value is changed, the remaining sliders also adjust and only Pareto optimal solutions are allowed. In reality that would mean obtaining the whole set of Pareto-optimal treatment plans in advance, a process which can take days or even weeks, which is not possible in clinical practice.

The established method of overcoming this challenge is the pre-computation of a number of specific treatment plans, often called 'anchor plans' or 'anchors'. These plans have to be chosen in a sophisticated way in order to get a good representation of the available solution range for the chosen treatment objectives. If the user navigates to a choice of objective values not contained in the set of pre-calculated plans, an approximation is performed by means of a mathematical combination of the existing results. Various studies have shown that such an approach facilitates the treatment planning process[47] and leads to higher plan quality and improved clinical outcomes[11][38][40][18]. A general introduction into the mathematical background application of MCO techniques in that specific context can be found at *Küfer et al*[49] *Hamacher* [35] and *Bokrantz*[5]

In the case of more atomic planning objectives, such as dose volume constraints or max/min dose limits, and a clear (ideally convex) mathematical formulation of the optimization problem, numerous solutions exist[49][32]. Research on Pareto-surface approximation[57][16][15] and dealing with approximation errors has been published[17][4][12]. MCO treatment planning has been developed for IMRT[49][50][68][58] and VMAT delivery[19][11]. A good general introduction into the current state of research is given by *Craft*[14].

## 2.4 Continuous function value approximation in cubic affine parameter spaces

This section gives some theoretical background on different methods for approximating a continuous function in a 3-dimensional affine parameter space. In the scope of this work, *affine* is defined as follows:

**Definition 1.** Let  $x_1, \dots, x_m \in \mathbf{R}^n$ . The affine space  $X_{aff}$  spanned by the  $x_i$  is defined as

$$X_{aff} = \left\{ x : x = \sum_{i=1}^m \lambda_i x_i; 0 \leq \lambda_i \leq 1 \forall i \in \{1, \dots, m\} \right\} \quad (2.4.1)$$

Let

$$D : \mathbf{R}^n \rightarrow \mathbf{R} \quad (2.4.2)$$

define a function which is continuous on  $X_{aff}$ .

The problem consists of approximating unknown function values  $D(x)$  for any  $x \in X_{aff}$  as a weighted sum of a limited amount of known function values  $D_{x_k}$ :

$$D_{approx}(x) = \sum_k \lambda_k D(x_k); \tag{2.4.3}$$

### 2.4.1 Approximation with 7 anchor points and its limitations

The first choice for an approximation base set under investigation consists of a subset of seven pre-calculated anchor points. Function values  $D$  on the three-dimensional affine parameter space spanned by the three coordinate axes:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \tag{2.4.4}$$

with  $x, y, z \in [0, 2]$  are approximated using the exact values  $D_{xyz}$  at the following coordinates:

$$\begin{aligned} D_1 &:= D \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}; D_2 := D \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}; D_3 := D \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix}; D_4 := D \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}; \\ D_5 &:= D \begin{pmatrix} 1 \\ 1 \\ 2 \end{pmatrix}; D_6 := D \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}; D_7 := D \begin{pmatrix} 2 \\ 1 \\ 1 \end{pmatrix} \end{aligned} \tag{2.4.5}$$

The approximation base points given in 2.4.5 is illustrated in figure 2.7. The unknown value  $D$




Figure 2.7: Illustration of 7D-approximation base


at position

$$\begin{pmatrix} x_1 \\ y_1 \\ z_1 \end{pmatrix} \tag{2.4.6}$$

is approximated as follows, depending on the octant where  $D$  is located (in the case of the example shown in figure 2.7 it is the octant defined by  $D_2, D_3, D_7, D_5$ ):

$$\begin{aligned}
 D = \frac{1}{3} [ & ((x_1 - 1) D_2 + (2 - x_1) D_3) \\
 & + ((y_1 - 1) D_2 + (2 - y_1) D_7) \\
 & + ((z_1 - 1) D_2 + (2 - z_1) D_5)]
 \end{aligned}
 \tag{2.4.7}$$

This is the equivalent of a linear interpolation in the direction of each coordinate axis and averaging all three results. This process is illustrated in figure 2.8. The first summand of 2.4.7



**Figure 2.8:** Illustration of the approximation steps with a 7D-approximation base

which represents the interpolation in  $x$  direction is represented by the pink circle, the second summand representing the approximation in  $y$  direction is shown as a green circle and the final summand in  $z$  direction is drawn in blue. While this approach seems to be plausible if the desired parameter coordinates lie on one of the coordinate axes close to  $D_2$ , the generalization leads to an unacceptable loss in accuracy.

This is due to the fact that for a proper balanced (in this case linear) approximation the function values at the points

$$D_8 := D \begin{pmatrix} 2 \\ 2 \\ 1 \end{pmatrix}; \quad D_9 := D \begin{pmatrix} 2 \\ 2 \\ 2 \end{pmatrix}; \quad D_{10} := D \begin{pmatrix} 1 \\ 2 \\ 2 \end{pmatrix};
 \tag{2.4.8}$$

have to be taken into account, too. In this case, however, this is by no means the case and it is clear that an approximation which is functional in the complete parameter space has to be performed differently. This is presented in the following section.

### 2.4.2 Trilinear interpolation with 8 anchor points

Trilinear interpolation is a mathematical technique often used in scientific applications where exact values only exist on the nodes of a specific parameter grid. These values can e.g. be color intensities in computer graphics, or temperatures during the simulation of physical processes

inside a component of some machine. The underlying task always consists of creating an approximation of the unknown value at a specific coordinate by combining the known values of the nearest neighboring grid points in a meaningful way. [41]

Figure 2.9 serves as an illustration of this task. Like in the previous section  $D_1 \dots D_8$  span a three-dimensional affine parameter space and continuous function  $D$  is defined. The values  $D_i$  at the corner points of the parameter space are known and we want to approximate the value

$$D(x \ y \ z)^T; \quad 0 \leq x, y, z \leq 1 \tag{2.4.9}$$

utilizing the known values. For a more detailed description of the approximation process,




Figure 2.9: Illustration of the general trilinear interpolation problem

which consists of three consecutive interpolations along each parameter axis, we define

$$\begin{aligned} D_1 &:= D \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}; D_2 := D \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}; D_3 := D \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}; D_4 := D \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}; \\ D_5 &:= D \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}; D_6 := D \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}; D_7 := D \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}; D_8 := D \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \end{aligned} \tag{2.4.10}$$

The first step illustrated in figure 2.10 is composed of four simultaneous linear interpolations in  $x$ -direction, resulting in the four value approximations  $D_{00}, D_{01}, D_{10}$  and  $D_{11}$  drawn in blue, all having the same  $x$ -value as  $D$ . The second step illustrated in figure 2.11 is composed of two simultaneous linear interpolations in  $y$ -direction, resulting in the two value approximations  $D_p$  and  $D_q$  drawn in pink, both having the same  $x$ - and  $y$ -value as  $D$ . Figure 2.12 shows the final linear interpolation along the  $z$ -axis which yields the wanted value  $D$ .

Re-substitution and mathematical simplification gives raise to the following:

$$D = \sum_{i=1}^8 \alpha_i D_i \tag{2.4.11}$$

The coefficients  $\alpha_i$  are given by

$$0 \leq \alpha_i \leq 1 \quad \forall i \in \{1, \dots, 8\}; \quad \sum_{i=1}^8 \alpha_i = 1 \tag{2.4.12}$$




Figure 2.10: First interpolation step




Figure 2.11: Second interpolation step

with

$$\begin{aligned}
 \alpha_1 &= (1-x)(1-y)(1-z) \\
 \alpha_2 &= x(1-y)(1-z) \\
 \alpha_3 &= xy(1-z) \\
 \alpha_4 &= (1-x)y(1-z) \\
 \alpha_5 &= (1-x)(1-y)z \\
 \alpha_6 &= x(1-y)z \\
 \alpha_7 &= xyz \\
 \alpha_8 &= (1-x)yz
 \end{aligned} \tag{2.4.13}$$

Simple re-calculation shows that a different order of the parameter axes  $x, y, z$  in the interpolation process yields the same coefficients, the order of interpolation directions is not important for the final result.

### 2.4.3 Subdivision of feature space, 27 anchor points

Approximation of function values over the whole parameter space with only eight known function values as shown in 2.9 has the property that the approximation error gets smaller with increasing proximity to a position of known function value. This means that the best




Figure 2.12: Final interpolation step

approximation quality will be achieved close to the corner points of the parameter space. However, in real-world applications it is often more likely that the function value to be approximated has a balanced parameter configuration, which means that the point defined by its parameter set in the parameter space is located more towards the center of the parameter cube. In this case, it would be more suitable to have pre-calculated function values in the center region available to improve the approximation quality.

Figure 2.13 shows a suitable parametrization taking central positions into account.




Figure 2.13: Parameter space sub-division

The first step when approximating  $D$  by trilinear interpolation is to determine the sub-space where its parameters  $x, y, z$  are located and its eight boundary data points  $D_i$ . The remaining steps are identical to 2.4.11. While this procedure requires 27 data points  $D_i$ , the approximation accuracy, especially in the central area of the parameter space can be expected to be much better.

### 2.4.4 Adaptive-grid trilinear interpolation

The method presented in the previous section can be generalized, since it is sufficient for the eight nodes used for interpolation to be aligned in a cuboid manner instead of a perfect cube. If a higher approximation accuracy with respect to one parameter is needed, the parameter space can be discretized more densely in that direction or even whole partitions of the parameter

space can be discretized with a higher resolution. According to the parameter coordinates of the desired function value, the eight relevant nodes encompassing the desired coordinate can be determined, and the value is approximated as shown in 2.4.11. An example for such an adaptive discretization of the parameter grid is shown in figure 2.14.




Figure 2.14: Adaptive resolution of the parameter space discretization

## 2.5 Mathematical characterization of patient geometry

### 2.5.1 Overlap-volume-histogram (OVH)

The overlap-volume-histogram (OVH) is a way to describe the spatial relationship of 3D objects by means of a 1D function. It has seen various uses in radiotherapy planning applications [87] [43] [86] [39]. Sometimes it is also referred to as distance-to-target-histogram (DTH) [91]. The OVH for a certain OAR with respect to a planning target  $T$  and the radius ( $r$ ) describes the relative fraction of the total volume  $V_{OAR}$  of the OAR located within the distance  $r$  of the surface of  $T$ :

$$OVH(r) = \frac{V(p_{OAR} | d(p_{OAR}, T) \leq r)}{V_{OAR}} \quad (2.5.1)$$

where  $p_{OAR}$  denotes a subset of of the OAR and  $d(p_{OAR}, T)$  denotes the distance between the points contained in  $p_{OAR}$  and the target. In case of overlapping OAR and PTV negative radius values are applied until  $OVH(r) = 0$ .

Figure 2.15 illustrates this method for three examples in 2D. The red shape is the target  $T$  and the blue shapes represent OARs. A qualitative sketch of the corresponding OVH graph is plotted next to each object constellation.

The whole OVH graph can be used as a feature to describe the spatial relationship between the two objects, as well as characteristic properties such as the mean slope, number of inflection points, or the penumbra width, only.

### 2.5.2 Inverse overlap-volume-histogram

In certain cases it makes more sense to switch object roles and consider the relative volume of the PTV with respect to the OAR object. While the original OVH graph is suitable for use-cases



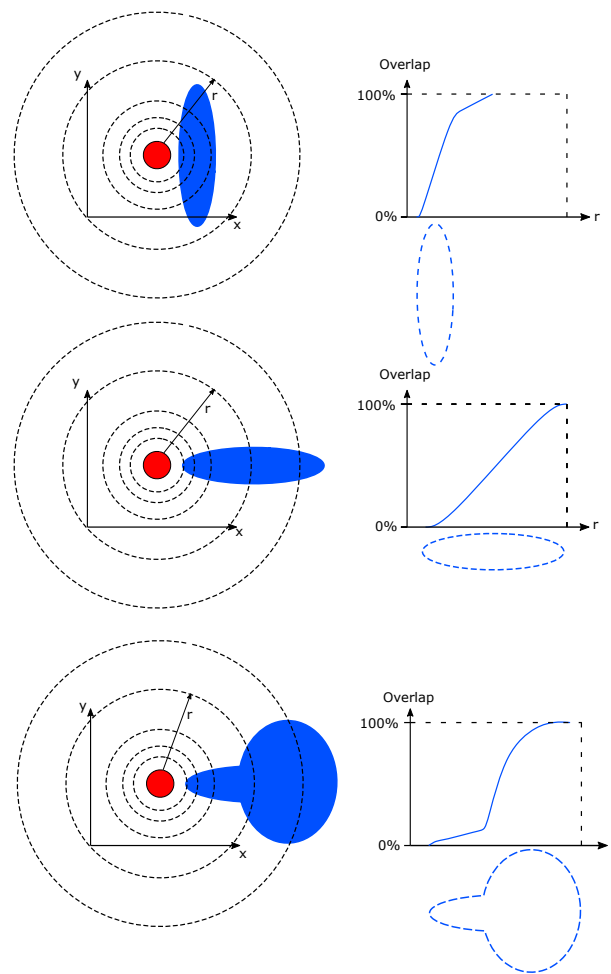
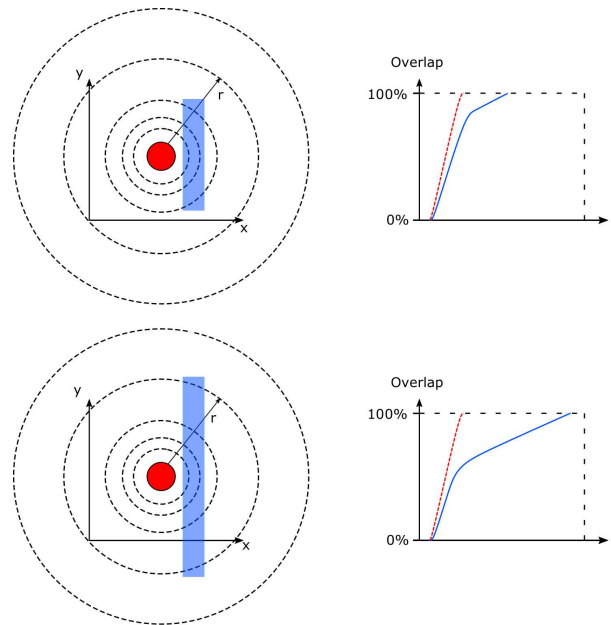


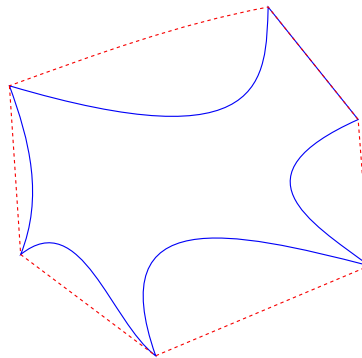
Figure 2.15: Generation of Overlap-Volume-Histograms(OVH)

with rather uniform PTV object and a number of OAR objects varying a lot in size and shape, with increasing radius  $r$  the shape characteristics of the central PTV object become increasingly insignificant for the shape of the OVH graph. In some cases the shape of the PTV object is more complicated and of more interest for characterizing the geometry than the OAR object, such as in spinal radiosurgery cases, where the most important OAR consists of the spinal cord or the spinal canal, which are both more or less cylindrical structures, while the PTV is a vertebral body or a section of it, which is a geometrically rather complex object.

This overlap histogram with inverted object roles will be referred to as *inverse OVH*. It also solves another problem arising in spinal cases. Usually not the complete spinal cord is delineated, since only the section in the proximity of the tumor site is clinically significant. For different delineation lengths of the OAR object the resulting OVH graphs look very different, even though the important area close to the PTV object is identical. The inverse OVH graph is not affected by this issue. Figure 2.16 shows a comparison of PTV object (red) and the same OAR object (blue) in two different delineation versions.



**Figure 2.16:** OVH (blue) and inverse OVH (red, dashed) for OAR objects of different length with otherwise identical geometrical setup



**Figure 2.17:** 2-D object(blue) and its convex hull(red)

### 2.5.3 Concavity/convexity of anatomical structures

The shape of an individual 3D object can be mathematically characterized by its *concavity*  $C$ , which describes the relationship between the surface area  $a$  of an object  $O$  and the surface area of its *convex* hull  $H_{conv}(O)$ :

$$C = \frac{a(O)}{a(H_{conv}(O))} \quad (2.5.2)$$

The *convex hull* of an object  $O$  is defined as the smallest convex object which contains  $O$  completely (for a definition of convexity see equation 2.2.4). This is illustrated in figure 2.17 in a 2D example. The dashed red line shows the convex hull of the blue object. Concavity of the involved organ structures can have a huge impact on whether the desired PTV dose coverage and OAR sparing can be achieved, since concave objects require complicated MLC shapes.

## 2.6 Artificial intelligence and machine learning

This section shall provide an introduction into the principles of artificial intelligence (AI) and machine learning (ML), present the specific algorithms applied in this thesis and define some important quality metrics. More detailed information can be found e.g. in *Goodfellow et al*[31], *Hastie et al*[36] and *Nielsen*[59]. While these methods and algorithms are basically a part of applied statistics, and some of the theoretical background has already been published some decades ago, the recent rise of artificial intelligence (AI) due to the enormous amount of readily available computation power, has lead to increased activity in this area, especially in the sub-field called machine learning.

### 2.6.1 The supervised learning problem

The task of predicting output values  $y^*$  of an unknown function  $f : \mathbf{R}^n \rightarrow \mathbf{R}$  with  $y^* = f(x^*)$  by using information contained in a given set of examples  $x_1, \dots, x_m; x^* \notin \{x_1, \dots, x_m\}$  and the corresponding function values  $y_1, \dots, y_m$  by means of a mathematical algorithm is called *supervised learning problem*.

In this case the term *supervised* is used because the algorithm has access to a number of corresponding input-value output-value data pairs which are guaranteed to be correct.

Typically the  $x_i$  are arranged in a  $m \times n$  matrix  $X$ , each row representing one sample  $x_i$ , each described by  $n$  numerical values called *features*. This matrix is often called the *input matrix* or *feature matrix*. The corresponding function values  $y_i$  are grouped in the  $m$ -dimensional *output vector* or *target vector*  $y$ . The set  $X, y$  is often referred to as *training data* or *training set*, since it is used to fit an existing prediction model  $P(x)$  to the specific training data. During the training process weights and parameters of  $P(x)$  are adjusted with the goal of optimizing the prediction with respect to a quality metric or cost function  $C$ .

Let  $y_p := P(X)$  be the predicted output values of the training data  $X$  and

$$C(y, y_p) := \frac{1}{m} \sum_{i=1}^m (y_i - P(x_i))^2 \quad (2.6.1)$$

denote the mean squared error as a measure of deviation between predicted and actual values in  $y$ . The goal of the training process is the minimization of this cost function. This minimization process is often done by applying some sort of *gradient descent method*, where the gradient of the cost function is calculated and  $P$  is updated iteratively by performing steps in the direction of the gradient.

In order to prevent *overfitting*, a case where the model fails to generalize well to the actual problem but merely exactly replicates the information contained in the training data, usually a partition of the training data is left out of the training process and is later used to asses the quality of the model. This partition is often referred to as *test set*.

Many machine learning algorithms can be adjusted by certain meta-parameters. A third partition of the data called *validation* data is therefore often set aside. The complete model development process consists of training the model on the training set, with different meta-parameter configurations. After each training process the model is tested on the validation set. The version performing best on the validation set is then applied to the test set for final assessment of the generalization quality of the model. Often many ( $k$ ) different sub-partitions

of the same available data into training set, validation set and test set are evaluated in order to mitigate overfitting, this process is called *k-fold cross-validation*.

### 2.6.2 Regression vs classification

If the function  $f$  which shall be approximated is a continuous function in  $\mathbf{R}$ , the problem of predicting the unknown value  $f(x)$  is called *regression problem*. If only a limited amount of function values is possible, e.g.  $f(x) \in \{y_1, \dots, y_j\} \forall x \in \mathbf{R}^n$  the input values are assigned a specific *label* or *class*, the problem is called *classification problem*. The biggest part of the chapters of this thesis which are related to machine learning deal with such *classification* tasks.

### 2.6.3 Important concepts in model training and assessment

#### Data preparation: regularization and oversampling

Since individual features in real life problems (the coordinates of the  $x_i$ ) can exist in value ranges which differ significantly, single features might completely dominate the evaluation of the cost function and the information contained in other features can be obscured. Therefore some sort of regularization of the data has to be performed. This is achieved by rescaling  $X$  so that the values of each column range from -1 to 1 with expectation 0.

The training data in classification problems often suffers from a high data imbalance. This means there are a lot more samples belonging to the one category than the other categories available. This can lead to a model where the algorithm is biased towards the numerically overrepresented category. A reliable technique to mitigate this problem is called *oversampling*. Additional samples of the underrepresented class are artificially generated from existing ones or, which is an extremely simple yet still very effective way, samples of the underrepresented class are randomly selected and duplicated until the classes are numerically balanced. The latter technique is called *random oversampling* (RO), while examples for artificial data generation are the *SMOTE*[10] and *ADASYN*[37] techniques.

#### Binary classification terminology

For evaluation of the model quality we consider a binary classification problem. The data instances are assigned exactly one of only 2 possible classes/labels:

- class  $N$  (hugely overrepresented, e.g. a high percentage of data points are of this class)
- class  $P$

The following terms are important:

- *True Positive (TP)*: Sample classified as having label  $P$  and its real label is  $P$
- *False Positive (FP)*: Sample classified as having label  $P$  even though its real label is  $N$
- *True Negative (TN)*: Sample classified as having label  $N$  and its real label is  $N$
- *False Negative (FN)*: Sample classified as having label  $N$  even though its real label is  $P$

The cardinality of each class or classification partition is denoted by vertical strokes  $||$ , such as e.g. the number of samples with label  $P$   $|P|$  or the number of false positives  $|FP|$ .

### Sensitivity/Specificity

A high-quality binary classification predictor aims to maximize correct detection of both classes and thus maximize both

$$sensitivity = \frac{|TP|}{|TP| + |FN|} \quad (2.6.2)$$

and

$$specificity = \frac{|TN|}{|TN| + |FP|} \quad (2.6.3)$$

### Precision/Recall/Accuracy

The terms *precision* and *recall* are often also used in the literature. *Recall* is identical to *sensitivity*(2.6.2) while *precision* is defined as

$$precision = \frac{|TP|}{|TP| + |FP|} \quad (2.6.4)$$

The general *accuracy* of a model can be calculated as

$$accuracy = \frac{|TP| + |TN|}{|P| + |N|} \quad (2.6.5)$$

While this has the advantage of describing the prediction quality of the model in a single value, it can give misleading information in the case of highly imbalanced data.

### ROC curves/AUC

A graphical illustration of the prediction quality of binary classifiers can be obtained by means of the *receiver-operator-characteristic (ROC)* curve. Usually the classification output of a model is not directly the label  $N$  or  $P$ , but the probability that the current sample has this label. Depending on the selected decision threshold, the trade-off between *sensitivity*(2.6.2) and the *false-positive-rate (FPR)*

$$FPR = \frac{|FP|}{|N|} \quad (2.6.6)$$

can be visualized. This is shown in figure 2.18. The ROC curve of a high-quality model tends to rise steeply and flatten out at a value close to one, the ideal ROC curve would immediately rise to 1 and remain constant at that value, while a ROC curve close to the diagonal means that the prediction quality is not much better than a coin toss.

Calculating the *area-under-curve (AUC)* allows for a numerical description of the quality of the ROC curve, usually a higher AUC value means a better model.

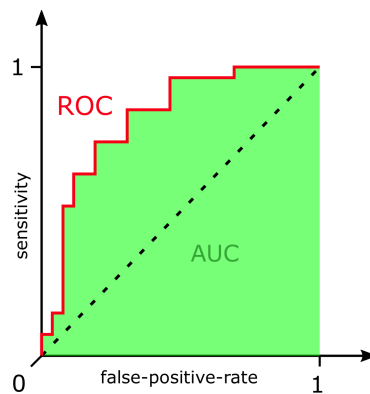


Figure 2.18: ROC-curve (red) and AUC (green)

## 2.6.4 Machine learning algorithms used in this thesis

### Support-vector-machines (SVM)

Support-vector-machines are a type of classifier trying to establish an optimally (with as big a margin as possible) separating hyperplane between the data [13]. A basic illustration is given in figure 2.19. Most of the time the data is not perfectly separable by a linear hyperplane. In this case the algorithm tries to create as few wrong classifications as possible, or have wrongly classified samples as close to the separating hyperplane as possible.

Application of the *kernel-trick* [21] allows for complex non-linear separation planes. However, a too complex kernel-based SVM model increases the risk of overfitting. SVMs are basic yet powerful classifiers and therefore a useful first tool for data exploration.




Figure 2.19: 2D-example of SVM with separable data

### Random-forest-classifier (RF)

The random-forest-classifier is a learning algorithm based on decision trees. It was first presented by Breiman in 2001[8]. A detailed overview about recent developments is given by *Fawagreh et al*[26]. Common decision trees suffer from the problem that they can represent the existing data really well but fail to generalize on unknown data. RF classifiers bypass this problem by generating a large number of trees generated from a randomly selected partition of the available data. In addition, each of these trees only uses a random subset of the available features, which mitigates the risk of overfitting, or one feature dominating the whole decision

process.

The RF decision process is schematically illustrated in figure 2.20. In reality a much higher




Figure 2.20: structure of the Random-Forest(RF) method

number of trees is generated (*Mayumi Oshiro et. al*[61] suggest between 64 and 128 trees), each using a random selection of the available features. In the figure the current sample would receive the red label. It is important that both the choice and order of features in each tree is random. The final quality of the predictor is evaluated with the remaining data which has not yet been chosen for training.

The generation of a large number of decision trees in the RF training process yields other important additional information about the input data: For each feature of the input data, the RF classifier can calculate its relative importance for the decision process. This can be very helpful when evaluating the efficiency of the current features or trying to engineer additional features to replace or supplement existing ones.

### Artificial neural networks (ANN)

Artificial neural networks (ANN), also often referred to as neural networks (NN) have become the predominant type of machine learning algorithm for many applications in recent years, achieving competition-winning results in image classification[48] and super-human performance in highly complex games such as Go [74]. Recently NNs have been also used successfully in radiotherapy planning[72][71] and patient QA[79]. Only a short introduction into the mechanism of a neural network shall be given at this place, more detailed information can be found in *Nielsen*[59] *Goodfellow et al*[31] and *Hastie et al*[36].

The smallest functional unit of an ANN is called a *neuron*. It is characterized by a *weight*  $w$




Figure 2.21: Input and output of a single neuron  $n$

and a *bias*  $b$  (see figure 2.21). The neuron receives a single or multiple numerical input values  $x_i$  and produces a single numerical output value  $y$  via its *activation function*  $f$  as follows:

$$y = f \left( w \left( \sum_i x_i \right) + b \right) \quad (2.6.7)$$




Figure 2.22: Sigmoid activation function  $f(x)$

The input variable of  $f$  is a weighted sum of the input values  $x_i$  shifted by the bias  $b$ . For the activation function  $f$  various options exist, the most natural choice being some sort of *sigmoid* function (figure 2.22). However, for numerical reasons one of the most popular variety is the *Rectified Linear Unit (ReLU)* depicted in figure 2.23. In a NN, as shown in figure 2.24, a high




Figure 2.23: ReLU activation function  $f(x)$

number of neurons are interconnected in a number of *layers*. The first layer, often called *input layer* serves as input channel for the features  $x_i$  (without weighting, bias of activation function). The following layers are called *hidden layers*. The example in figure 2.24 has one hidden layer consisting of 5 neurons  $n_i$ , each working as shown in equation 2.6.7. The final layer which is called *output layer* usually has as many output units  $y_i$  as there are classes in the classification task. The output value of each  $y_i$  is the predicted likelihood of the sample characterized by the features  $x_i$  having the correct label  $y_i$ .

The objective of training the NN is to find weights and biases for each neuron minimizing the error between the predicted label of the input data  $x$  and the correct output label  $y$ . As with activation functions, lots of different ways for defining this *cost-* or *loss-function* exist. The iterative process used for training NNs is called *backpropagation algorithm*. The loss value for a certain number or *batch* of training data is calculated and the gradient of the loss function with respect to changes in specific weights and biases is evaluated. Using a gradient descent method, weights and biases are updated in the direction of the gradient in order to decrease the loss value. Running the whole training data through this process once is called an *epoch*, and the training process usually consists of a high number (e.g. hundreds) of epochs.

Due to the high computational demand of this process, calculations are often run on powerful computation clusters in the cloud or on graphics processing units (GPUs) which work highly parallelized. In contrast to the expensive training procedure, prediction of the output value for an unknown sample with the final model is a straight forward and very fast procedure.

In contrary to this simple example, NNs for solving complicated tasks, such as image recognition,



consist of thousands of input values and a high number of large hidden layers, this type of NN is often referred to as *deep neural networks* and their application is called *deep learning*.




Figure 2.24: Three-layer neural network with 5 neurons in the hidden layer



Chapter **3**


The VMAT planning Element from an MCO  
perspective

## 3.1 The Brainlab Elements VMAT planning system (Cranial SRS/Spine SRS)

### 3.1.1 The Brainlab Elements VMAT planning system

#### General overview

The Brainlab Elements suite is an indication specific RT treatment planning software solution providing image processing and contouring functionality as well as radiotherapy treatment planning and DICOM export. Figure 3.1 shows the graphical user interface (GUI) for planning



**Figure 3.1:** Screenshot of the main planning GUI of the Cranial Element

a cranial VMAT treatment. Three of the four views on the right show axial, coronal and sagittal views of the currently selected image set, including the contours of previously delineated organs. The remaining section of the central screen area is used to display a 3D view of the geometric structures as well as a 3D representation of the VMAT arc trajectories.

The left side of the screen shows a list of all structures including relevant dosimetric information concerning the optimization process. The organ listed directly after the PTV (in this case the brainstem) is explicitly selected by the user as the clinically most relevant OAR and receives a special treatment during optimization, which will be explained later.

The duration of the first optimization run is significantly longer than the duration of subsequent optimizations, since important pre-calculation steps, such as discretization of the fluence beamlets, have to be performed. Once this is achieved, further optimizations are fairly quick (about 1-2 minutes for a typical cranial case, depending on the planning hardware). However, in spinal cases, where the PTV consists of a section or the total body of a vertebra, a single

optimization can take up to one hour or longer. This is due to the complicated PTV shape which is addressed with a customized version of the optimization algorithm.

After an optimization run the calculated dose distribution is shown in the slice views and the DVH histograms for each organ and the PTV can be reviewed (figure 3.2).

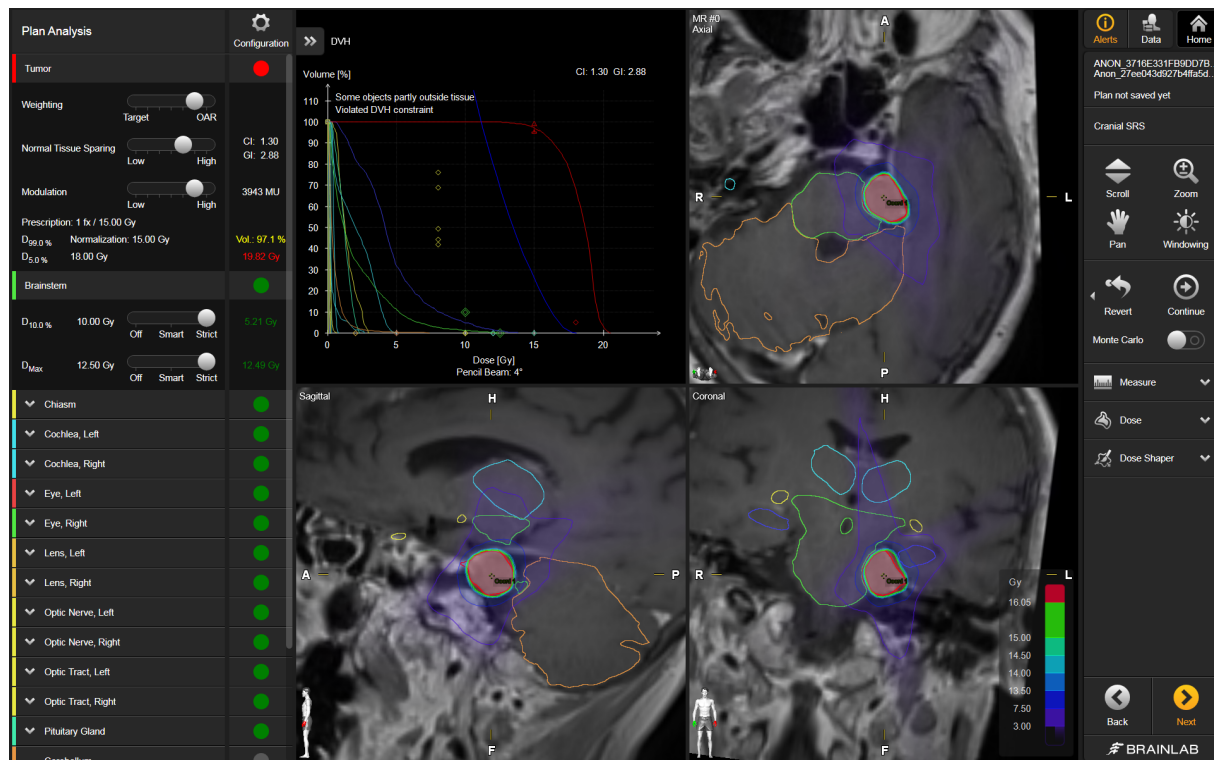


Figure 3.2: Screenshot after optimization showing dose distributions and DVH

### 3.1.2 Configuration of planning objectives and constraints

Constraints are externally imposed mandatory dosimetric properties a treatment plan has to fulfill in order to be clinically acceptable. Optimization objectives are internally used to generate the optimizer's cost function in order to meet the constraints. Constraints and objectives can be identical, but it is also possible that objectives are stricter or that the number of optimization objectives exceeds the number of constraints. More strict objectives can be used to force the optimizer towards fulfillment of the clinical constraints.

The organs which are to be considered for optimization, as well as their individual DVH constraint types (such as maximum/minimum dose, dose received by a volume percentage) and -values, as well as the dose prescription to the PTV, are described in a *clinical protocol file*, which can be prepared outside the planning application and is selected upon entering the planning workflow. The state of each OAR constraint can be adjusted by a slider control to one of the three options:

1. *Off*: This dose-volume (DV) value is ignored during optimization
2. *Smart*: This DV value is considered for optimization but violations are allowed to a certain extent if needed.

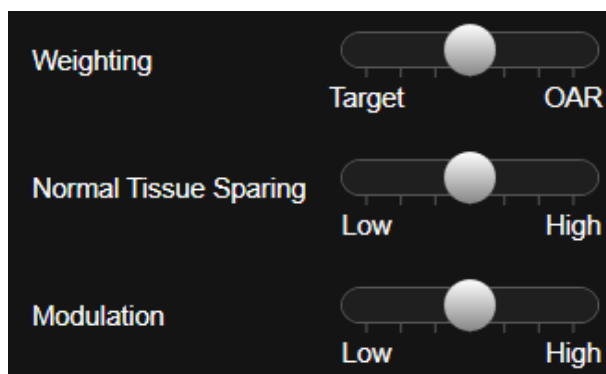
3. *Strict*: This DV value is considered for optimization as a non-negotiable constraint, any violation will be punished in the optimizer's cost function.

### 3.1.3 High-level composite optimization modification sliders and their functionality

The main part of this thesis is related to the functionality of the slider controls in the upper left corner of the screen (figure 3.3) and their influence on the generated treatment plans.

In many existing solutions there exist similar slider controls enabling the user to perform some sort of multi-criteria optimization (MCO) workflow. The tolerated threshold for user-specific DVH values can be adjusted for one value, and by means of information taken from some model of the Pareto surface of the solution space, the slider value ranges for the remaining DVH values are adjusted so that only feasible solutions remain.

In contrast to this approach, the VMAT optimizer used in the Elements planning system has



**Figure 3.3:** The three slider controls for adjusting the composite meta-objectives

only three main controls for individualizing the outcome of the optimization process. Each of them is dedicated to the modification of a certain quality aspect of the resulting treatment plan. In contrast to the often used approach mentioned above, all possible constellations of slider positions can be used as input for the optimizer and each slider control is not dedicated to a single DVH value only, but modifies a whole group of dosimetric aspects. Pareto optimality in this context does not mean strict dosimetric fulfillment of all optimization goals, but producing the best feasible result with the current slider positions.

#### The target-OAR-weighting control

The target-OAR-weighting slider generally controls the planning compromise between the PTV and the OARs, e.g.:

- The importance of the PTV versus all existing OARs
- The special emphasis on a user-defined most important OAR during the optimization process
- The importance of fulfilling the upper dose constraint of the PTV

- The tolerated PTV volume prescription reduction in case a constraint of the most important OAR is violated, and in which region of the PTV this reduction shall be allowed

### **The normal tissue sparing control**

The normal tissue sparing slider controls the importance of sparing the normal tissue surrounding the PTV by

- Reducing the dose in explicitly specified OARs in the vicinity of the tumor
- Reducing the dose in tissue surrounding the tumor, even if it is not part of a delineated organ
- Adjusting size and handling of artificially generated helper structures which are internally used during optimization
- Maximizing the conformity of the prescription dose region with the PTV (CI) and the steepness of the dose fall-off outside the PTV (GI)

### **The dose modulation control**

The modulation slider controls the amount of dose modulation in the resulting treatment plan. The term 'modulation' refers to

- The overall number of monitor units (MU)
- The variability of the number of MU per arc segment
- The amount of leaf movement between control points
- The amount of leaf movement within one segment
- The size of the segments (smaller field sizes correspond to higher modulation)

### **Slider range and discretization**

The internal numerical range of the sliders is internally fixed to numerical ranges suited best for most typical planning situations.

While the numerical value space for each slider is a continuum, for the sake of usability each slider range was discretized into seven equidistant positions.

In the commercially available version of the software, the leftmost weighting slider position automatically changes the mode of the OAR constraints from 'Strict' to 'Smart' to put even more emphasis on fulfillment of the PTV coverage. For continuity reasons this functionality has been disabled in the research version of the optimizer used for this thesis.

## 3.2 Elements SRS VMAT planning from an MCO perspective

### 3.2.1 Weighting-, normal tissue sparing- and dose modulation slider values as base of a 3D-parameter space(3DPS)

In the context of multi-criteria optimization, the three previously described sliders can be considered as follows:

Each dose distribution achieved after optimization for a certain configuration of slider positions is optimal in the sense, that with the chosen settings, there is no better result possible. Therefore it can be conceived as Pareto optimal in the context of this specific optimization problem. Thus, all three sliders span a space of Pareto optimal solutions. The user can explore this solution space interactively in order to find the preferred trade-off between OAR sparing, PTV coverage, normal tissue sparing and dose modulation (e.g. complexity of the resulting treatment plan in terms of quality assurance (QA), delivery and duration). A brute-force implementation would simply calculate the individual optimized treatment plans instantaneously, once the slider positions are changed. In reality, such a procedure is possible but far too time consuming. Even in straight-forward cranial cases the waiting time for each optimization process makes a fluid exploration of the solution space too cumbersome to be of practical relevance.

Therefore, within the scope of the composite 'meta-objectives', in combination with the non-convex nature of the optimization problem and the heuristic brute-force optimization approach used in the Brainlab Elements optimizer, a feasible approach for Pareto surface approximation has to be found.

A good set of anchor plans should have the following properties:

- Low number of pre-calculated plans in order to allow for an expeditious planning process
- Good and fast approximation of points on the approximated Pareto surface
- Possibility to increase approximation quality by subsequent addition of optimized plans

### 3.2.2 Analysis of the solution space generated in the 3DPS by means of the Brainlab VMAT optimizer

In order to find an approximation method satisfying the items mentioned in 3.2.1, it is necessary to perform an exploratory evaluation of the feasible solution space. To this purpose a dataset comprised of cranial and spinal patient geometries with varying PTV sizes, shapes and positions was chosen. The sample data consisted of a total of 18 spinal and 44 cranial test geometries. An overview of the most relevant organ structures available in the data is given in Table 3.1.

For all cranial cases the brainstem was labeled as most crucial OAR, while for the spine data two studies were performed. In one case the spinal canal was specified as the crucial OAR, in the other the spinal cord. Selecting the spinal canal makes the task more challenging for the optimizer, due to its greater volume and direct connection to the vertebral body. PTV volumes ranged from  $1\text{cm}^3$  up to  $113.2\text{cm}^3$  in the cranial geometries and  $29.4\text{cm}^3$  to  $61.3\text{cm}^3$  in the spinal data. PTV concavity ranged from 1.12 - 2.03 (cranial) and 1.34 - 2.92 (spine). More detailed information on the properties of the patient geometries used for data acquisition can be found



Table 3.1: Geometric organ structures in the data

Indication	OAR objects available in all geometries	additional OAR objects	OAR specified for special treatment
Cranial	brainstem, optic nerves, optic tract, chiasm	hippocampus left, hippocampus right [42]	brainstem
Spinal	spinal cord, spinal canal	-	spinal cord, spinal canal

in the appendix of this thesis 6.1.1.

Figure 3.4 illustrates some of the spinal patient geometries. PTV concavity and complexity of the optimization problem decline from left to right. Some examples for the cranial geometries used throughout this thesis are shown in figure 3.5. In the left case it is much easier for the optimizer to satisfy the optimization constraints since the PTV is of limited size and within a safe distance from the brainstem.




Figure 3.4: Patient geometries of spine cases #7-#9. PTV (orange) and cropped spinal canal (OAR, purple).




Figure 3.5: Patient geometry images of cranial cases #5 and #36 including PTV (orange) and brainstem (OAR, purple). Other objects shown are the optical structures including the eyes. Case #5 is a much harder task for the optimizer due to PTV shape and size. In addition the PTV is directly engulfing large parts of the OAR.

Detailed information on planning constraints and machine setup used for optimization can be found in table 3.2. In terms of the constraint mode options introduced in section 3.1.2, for

Table 3.2: Planning constraints

Indication	PTV Dose Prescription	Planning constraints for special OAR	OAR specified for special treatment
Cranial	15Gy 1 Fraction 99% PTV coverage, 70% min. allowed PTV coverage, max. Dose 18Gy <i>delivery method:</i> 3 arcs symmetric setup <i>machine model:</i> 6MV flattened Varian HD120 MLC	max.Dose 12.5 Gy, $D_{10\%} < 10Gy$	Brainstem
Spinal	16Gy 1 Fraction 95% PTV coverage, 70% min. allowed PTV coverage, max. Dose 20Gy <i>delivery method:</i> 1 360° arc <i>machine model:</i> 6MV flattened Varian HD120 MLC	max. Dose 14 Gy, $V_{10Gy} < 0.25cm^3$ , $V_{7Gy} < 1.2cm^3$	Spinal Canal/Spinal Cord

consistency reasons all OAR constraints were set to *Strict* for generation of any result data used in this thesis.

### 3.3 Cranial data

#### 3.3.1 PTV coverage

Figure 3.6 shows the relative coverage of the PTV for all cranial cases. Values range from 0.7 up to the desired value of 0.99 with median values ranging from 0.78 - 0.99. Values do not go beyond the minimum coverage of 0.7 because this user adjustable minimum value is enforced by the optimizer. The achievable value range per case varied from 0.00 (case #26, #29) to 0.29 (case #1, #5, #42). The variation in the results can be explained by the fact that the PTV in case #1 has a volume of  $29.8cm^3$ , a concavity of 1.61 and is at some points directly touching the brainstem (OAR) which makes for a challenging optimization task.

Geometries like in case #26, with a PTV volume of  $1cm^3$ , a PTV concavity of 1.18 and a minimum distance of  $9mm$  between PTV and brainstem (OAR) are much simpler and the optimizer is able to achieve the desired coverage for almost any parameter configuration.




Figure 3.6: Boxplots for relative PTV coverage, 343 optimizations per case

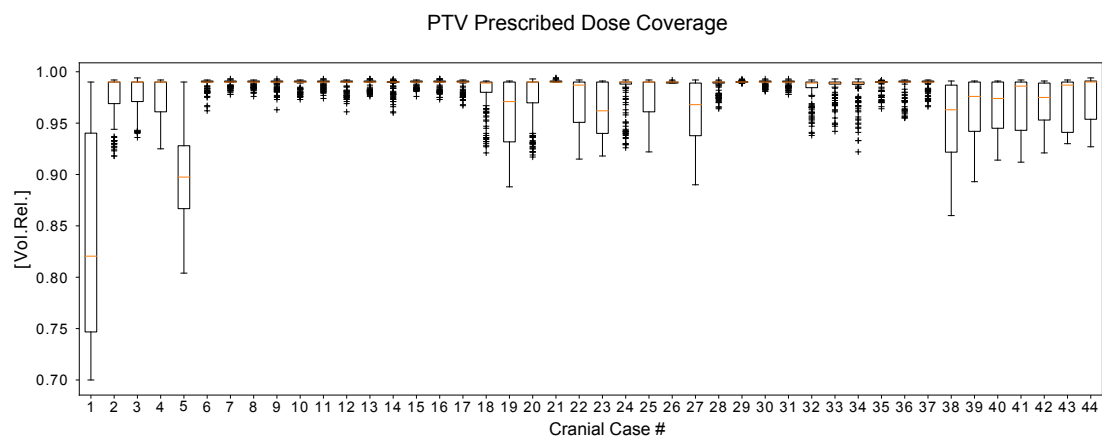



Figure 3.7: Boxplots for relative PTV coverage, only 'Modulation' values of 3 or higher allowed

A lot of the outliers with low coverage can be explained by the 'Modulation' parameter value. While some cases like case #1 are so complex that a harsh trade-off is almost unavoidable, allowing for more complicated MLC shapes can vastly improve the results. Figure 3.7 shows the coverage results if only 'Modulation' values greater than 2 are allowed. Especially in cases #36 - #42 the outliers disappear, which is also visible in figure 3.8 showing the Pareto surfaces of achievable coverage values in case of *Modulation* = 0 (left), *Modulation* = 3 (middle) and *Modulation* = 6 right. Even for high 'Normal Tissue Sparing' values and maximum OAR sparing the PTV coverage values in the middle and right figure are much higher than on the left surface.


### 3.3.2 PTV maximum dose

Figure 3.9 shows the maximum dose in the PTV object for all cranial cases. Values range from 15.81 Gy - 20.38 Gy with the median values ranging from 17.66 Gy - 18.69 Gy. The achievable value range per case varied from 2.07 Gy (case #23) - 3.84 Gy (case #1).

The data shows that the optimizer exploits the full range up to the specified desired max. value of 18 Gy. The cases where this value is massively exceeded have a complicated geometry (case #1) and a very big PTV ( $113.2\text{cm}^3$ , case #5).



**Figure 3.8:** Pareto surface of feasible PTV coverage in a selected cranial VMAT case (case #40), *Modulation* slider in minimum (left image) middle (center image) and maximum position (right image)



**Figure 3.9:** Boxplots of PTV Max. Dose, 343 optimizations per case

### 3.3.3 PTV minimum dose

Figure 3.10 shows the minimum dose in the PTV object for all cranial cases. Values range from 5.91 Gy to 14.92 Gy with the median values ranging from 9.11 Gy - 14.53 Gy. The achievable value range per case varied from 1.1 Gy (case #16) to 7.90 Gy (case #1).

Again, the effect of the three meta-parameters is biggest for complicated cases (high PTV concavity, huge PTV volume close to the OAR) and smallest for cases with a small convex PTV within a safe distance from the brainstem.




Figure 3.10: Boxplots of PTV min. Dose, 343 optimizations per case

### 3.3.4 PTV mean dose

Figure 3.11 shows the mean dose in the PTV object for all cranial cases. Values range from 15.31 Gy to 17.75 Gy with the median values ranging from 16.24 Gy - 17.21 Gy. The achievable value range per case varied from 1.62 Gy (case #5) to 2.34 Gy (case #42).

Again, the effect of the three meta-parameters is biggest for complicated cases (high PTV concavity, huge PTV volume close to the OAR) and smallest for cases with a small convex PTV within a safe distance from the brainstem.

The values show that in the cranial sample cases the PTV mean dose ranges did not vary a lot, even though PTV shape and volume differed a lot for certain cases.




Figure 3.11: Boxplots of PTV Mean Dose, 343 optimizations per case

### 3.3.5 Dose in OAR constraint point

The first OAR value in this analysis of the cranial example cases is the  $D_{10\%}$  in the brainstem. The boxplots in figure 3.12 imply that the optimization goal of  $D_{10\%} < 10$  Gy is safely met in most cases, however, cases #1, #5 and #27 show severe violations of this constraint value. Due to their complex geometry, it was to be expected that cases #1 and #5 might need more sophisticated planning, going beyond simple modifications of the three main adjustment sliders

only, to produce acceptable treatment plans. The extreme value in case #27 can be explained by the fact that this is the only case in the data, where the PTV not only touches the brainstem, but also is a complete subset of this OAR. Therefore, there is no way for a treatment beam to reach the target without traversing a certain depth of OAR tissue.




Figure 3.12: Boxplots of Brainstem  $D_{10\%}$ , 343 optimizations per case

### 3.3.6 OAR mean dose

Figure 3.13 shows the mean dose in the OAR object 'Brainstem' for all cranial cases. Values range from 0.10 Gy to 15.17 Gy with the median values ranging from 0.31 Gy - 11.51 Gy. The achievable value range per case varied from 6.44 Gy (case #1) to 0.10 Gy (case #29).

Again the more complicated cases show more effectiveness of the slider value changes, while in the easy cases, where the optimization constraints are already satisfied, further optimizations with different slider settings are not likely to yield significant improvements.




Figure 3.13: Boxplots of Brainstem Mean Dose, 343 optimizations per case

### 3.3.7 OAR minimum dose

Figure 3.14 shows the minimum dose in the OAR object 'Brainstem' for all cranial cases. Values range from 0.0 Gy to 6.87 Gy with the median values ranging from 0.0 Gy - 3.0 Gy. The

achievable value range per case varied from 0.0 Gy (case #19) to 4.83 Gy (case #1).

As seen before, cases #1 and #5 are again the cases where even the minimal dose received by the brainstem is significant and the sliders can create the greatest dose variations. In the remaining cases however, significant parts of the brainstem receive virtually no radiation.




Figure 3.14: Boxplots of Brainstem Min. Dose, 343 optimizations per case

### 3.3.8 OAR maximum dose

Figure 3.15 shows the maximum dose in the OAR object 'Brainstem' for all cranial cases. Values range from 0.16 Gy to 19.38 Gy with the median values ranging from 0.62 Gy - 17.59 Gy. The achievable value range per case varied from 0.48 Gy (case #8) to 6.93 Gy (case #19).

The cases where the desired maximum dose of 12.5 Gy in the brainstem is significantly surpassed, are geometries where the PTV is in direct contact with the OAR, so that a certain amount of OAR voxels inevitably receive a high radiation dose.




Figure 3.15: Boxplots of Brainstem Max. Dose, 343 optimizations per case

### 3.3.9 Gradient index (GI)

Figure 3.16 shows the gradient index (GI) for the cranial sample cases. Values range from 2.45 up to 5.71 with the median values ranging from 2.56 up to 4.08. The achievable value range per

case varied from 0.46 (case #19) to 2.28 (case #42) and up to 2.66 (case#1).

In most cases  $GI < 3$  could be achieved, with the mean value for the smallest achievable GI for all cases being 2.69.




Figure 3.16: Boxplots of Gradient Index (GI), 343 optimizations per case

### 3.3.10 Conformity index (CI)

Figure 3.17 shows the conformity index (CI) for the cranial sample cases. Values range from an excellent 1.02 up to 2.69 with the median values ranging from 1.06 up to 1.83 (this occurred in case #41 where the PTV is of a u-shaped structure and in sub-optimal treatment position for the fixed three-arc-setup). The achievable value range per case varied from 0.11 to 1.37.

On average, the minimal CI achievable for all cases was 1.11.




Figure 3.17: Boxplots of Conformity Index (CI), 343 optimizations per case


### 3.3.11 Monitor units (MU)

Figure 3.18 shows the amount of monitor units (MU) for the cranial sample cases used when trying to achieve the desired  $D_{99\%}$  of 15Gy in the PTV. All cases were planned using the same 6MV flattened machine model using a *Varian HD120* MLC. Values range from 1962 up to 5040 with the median values ranging from 2467 up to 3942. The achievable value range per case




varied from 424 to 2812.

Especially complex cases with big concave PTVs show a great variation in MU values, since a better coverage and conformity has to be achieved by smaller, more complicated MLC shapes, resulting in higher overall MU values for the treatment plan.



**Figure 3.18:** Boxplots of monitor units (MU), 343 optimizations per case

The MU value is highly influenced by the setting of the 'Modulation' slider, especially the lower settings create a spread in the results. Once the 'Modulation' slider is at least at the central position, its influence on the amount of MU decreases. This can be seen in figure 3.19 which shows the MU values in the cases where 'Modulation'  $\geq 3$ . Compared to figure 3.18 the values are not that spread out any more, the middle percentiles are grouped noticeably closer.



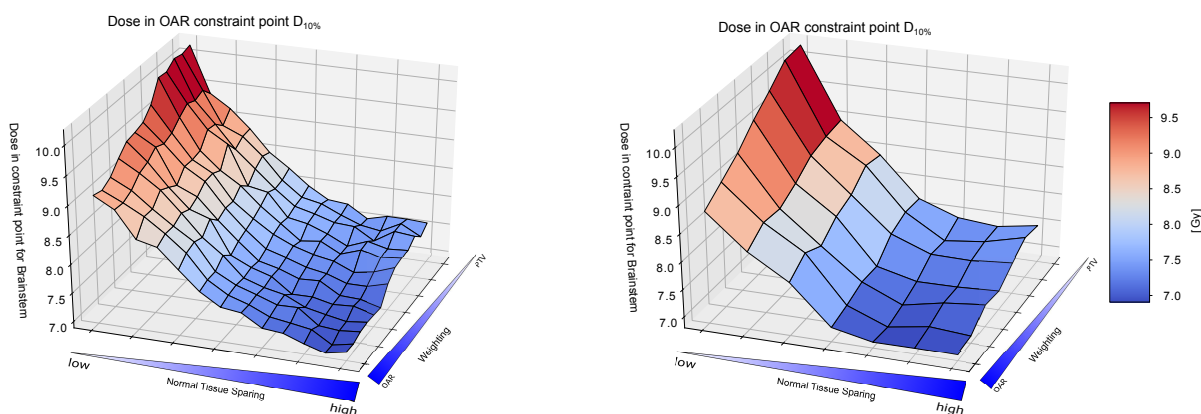
**Figure 3.19:** Boxplots of monitor units (MU), only 'Modulation' values of 3 or higher allowed

### 3.3.12 Double resolution

The complete range of each slider is discretized into 7 equidistant positions in order to make the user interface consistent and streamline the workflow. Since the parameters change continuously, it seems unlikely that a finer resolution, or intermediate slider positions, will produce optimization results turning out to be extreme outliers in terms of their dosimetric properties

when compared to the treatment plans produced with the 7 step discretization.

This hypothesis could be confirmed by choosing an experimental discretization with 13 equidistant positions per slider which means an increase of the resolution by factor 2. For a selected cranial case all 2197 (instead of 343) possible optimizations were performed and checked for consistency. As an example, figure 3.20 shows the Pareto surface of the constraint point dose  $D_{10\%}$  in the brainstem with respect to all possible options for the 'Weighting' and 'Normal Tissue Sparing' slider with the 'Modulation' parameter fixed in the central position. The right picture shows the surface with the standard 7x7 resolution while the left image displays the same situation including all intermediate results obtained by the increased parameter resolution. Other than a certain amount of noise, which is to be expected in a stochastic optimization process, the left picture does not contain any additional information.



**Figure 3.20:** Pareto surface of feasible  $D_{10\%}$  values in the Brainstem of a cranial VMAT case, *Modulation* slider in middle position, same plan optimized with standard (7 equidistant positions, right) and double (14 equidistant positions, left) slider resolution

## 3.4 Spinal data

### 3.4.1 PTV coverage

Figure 3.21 shows the relative coverage of the PTV for all spinal cases where the section of the spinal canal in the vicinity of the relevant vertebra was specified as OAR. Values range from 0.7 up to the desired value of 0.95, with median values ranging from 0.79 - 0.89. Similarly to the cranial examples, values do not go beyond the minimum coverage of 0.7 because this user adjustable minimum value is enforced by the optimizer. In all cases the achieved coverage values spanned the complete possible value range of 25%.




Figure 3.21: Boxplots for relative PTV coverage, 343 optimizations per case

### 3.4.2 PTV maximum dose

Figure 3.22 shows the maximum dose in the PTV object for all spinal cases. Values range from 17.12 Gy to 37.92 Gy with the median values ranging from 19.97 Gy - 21.57 Gy. The achievable value range per case varied from 4.90 Gy (case #10) to 20.43 Gy (case #9).

The data shows that the median results generally stay close to the max. value of 20 Gy. Slight violations of this value frequently still occur frequently. The infrequent cases where this value is massively exceeded (cases #3,6,9,12,15,18) are the cases where the PTV consists of the complete vertebral bodies which have a complicated structure (concavity values ranging from 2.07 – 2.92). The greatest maximum PTV dose correlates with the largest convexity value. Not only the 'Modulation' value is responsible for the extreme outliers concerning the PTV maximum dose, but also the value of the 'Normal Tissue Sparing' parameter. All the data points marked by a '+' in figure 3.22 are results of optimizations performed with both 'Normal Tissue Sparing'  $\leq 3$  and 'Modulation'  $\leq 3$ .




Figure 3.22: Boxplots of PTV Max. Dose, 343 optimizations per case

### 3.4.3 PTV minimum dose

Figure 3.23 shows the minimum dose in the PTV object for all spinal cases. Values range from 3.03 Gy to 14.10 Gy with the median values ranging from 5.86 Gy - 9.68 Gy. The achievable value range per case varied from 6.68 Gy (case #2) to 10.70 Gy (case #13).

The maximum values all occurred when the 'Weighting' parameter was completely shifted towards the PTV. Since PTV and OAR share a direct boundary area, even including some slight overlaps, these values are very sensitive to small shifts in the dose gradient.




Figure 3.23: Boxplots of PTV Min. Dose, 343 optimizations per case

### 3.4.4 PTV mean dose

Figure 3.24 shows the mean dose in the PTV object for all spinal cases. Values range from 15.16 Gy to 19.20 Gy with the median values ranging from 17.01 Gy - 17.52 Gy. The achievable value range per case varied from 6.30 Gy (case #5) - 3.94 Gy (case #3).

It is noticeable that the median PTV dose is in all cases higher than the desired prescribed dose of 16 Gy.




Figure 3.24: Boxplots of PTV Mean Dose, 343 optimizations per case

### 3.4.5 Dose in OAR constraint point

Figure 3.25 shows the  $D_{0.25cm^3}$  constraint dose in the OAR object 'Spinal Canal' for all spinal cases. Values range from 7.59 Gy to 16.64 Gy with the median values ranging from 9.18 Gy - 9.55 Gy. The achievable value range per case varied from 5.58 Gy (case #10) to 8.41 Gy (case #5). The median values are below the defined constraint value of 10 Gy. The range of achievable values for this constraint increases with increasing convexity of the PTV. This can be explained due to the fact that concave PTVs make the planning situation so challenging that that the maximum allowed OAR dose has to be exploited in order to be able to maximize PTV coverage.




Figure 3.25: Boxplots of Spinal Canal  $D_{0.25cm^3}$ , 343 optimizations per case

### 3.4.6 OAR mean dose

Figure 3.26 shows the mean dose in the OAR object 'Spinal Canal' for all spinal cases. Values range from 2.4 Gy to 9.21 Gy with the median values ranging from 2.94 Gy - 5.33 Gy. The achievable value range per case varied from 2.11 Gy (case #10) to 4.44 Gy (case #3). The maximum values all occurred with both 'Weighting' and 'Modulation' slider in the leftmost position, e.g. the simplest MLC shapes possible and full focus on PTV dose coverage.




Figure 3.26: Boxplots of Spinal Canal Mean Dose, 343 optimizations per case

### 3.4.7 OAR minimum dose

Figure 3.27 shows the minimum dose in the OAR object 'Spinal Canal' for all spinal cases. Values range from 0.19 Gy to 2.74 Gy with the median values ranging from 0.21 Gy - 2.12 Gy. The achievable value range per case varied from 0.05 Gy (case #11) to 1.11 Gy (case #9). In the spinal geometries this value is directly correlated to the cropping length of the OAR and does not contain much information, since it will become virtually zero if the OAR length is chosen large enough.




Figure 3.27: Boxplots of Spinal Canal Min. Dose, 343 optimizations per case

### 3.4.8 OAR maximum dose

Figure 3.28 shows the maximum dose in the OAR object 'Spinal Canal' for all spinal cases. Values range from 8.71 Gy to 18.93 Gy with the median values ranging from 10.83 Gy - 14.08 Gy (case #4). The achievable value range per case varied from 5.84 Gy (case #11) to 8.49 Gy (case #3).

With the slight exception of case #4, all median values are well under the desired maximum dose value of 14 Gy. The high maximum values cannot be avoided if a good PTV coverage is desired, with PTV and OAR being so close together, some OAR voxels are bound to receive the prescription dose.




Figure 3.28: Boxplots of Spinal Canal Max. Dose, 343 optimizations per case

### 3.4.9 Gradient index (GI)

Figure 3.29 shows the Gradient Index (GI) for the spinal sample cases. Values range from 2.70 up to 6.52 with the median values ranging from 3.12 up to 4.64. The achievable value range per case varied from 1.22 (case #2) to 3.44 (case #6).

In general GI values were not as good as the results in the cranial sample cases. The results fall into three distinct clusters, depending on PTV concavity. In the cases where the PTV consists of the main vertebral body only, without the vertebral arch, and thus has the lowest concavity, the resulting GI values are lowest as well (case #2, #5, #8, #11, #14, #17). If the whole vertebra (body and vertebral arch) are irradiated (case #3, #6, #9, #12, #15, #18) and the concavity (and volume) of the PTV structure is highest, the GI values are highest, too.




Figure 3.29: Boxplots of Gradient Index (GI), 343 optimizations per case

### 3.4.10 Conformity index (CI)

Figure 3.30 shows the Conformity Index (CI) for the spinal sample cases. Values range from 1.07 up to 4.78 with the median values ranging from 1.19 up to 1.55. The achievable value range per case varied from 0.73 (case #11) to 3.63 (case #6).

The results fall into the same but not as distinct cluster structure as the GI values. The cases with the largest PTV volume and -concavity all show a group of outliers with extremely bad CI results.




Figure 3.30: Boxplots of Conformity Index (CI), 343 optimizations per case

### 3.4.11 Monitor units (MU)

Figure 3.31 shows the amount of monitor units (MU) used for the spinal sample cases when trying to achieve the desired  $D_{95\%}$  of 16Gy in the PTV. All cases were planned using the same 6MV flattened machine model as in the cranial cases, using a *Varian HD120* MLC. This time MU values range from 2985 up to 14307 with the median values ranging from 5313 up to 10262. The achievable value range per case varied from 3031 to 9010.

The results fall into the same PTV-shape-driven clustering as the CI and GI values and some of the DVH constraints. In all spinal sample cases the MU values seem to be depending on the PTV shape to a higher degree than on its volume. The PTV volume in case #9 is  $61.3\text{cm}^3$  and  $36.4\text{cm}^3$  in case #12 (a ratio of 1.68) while the median MU values are 9749 MU vs 8864 MU (a ratio of 1.10) and the concavity values are 2.81 vs 2.26.

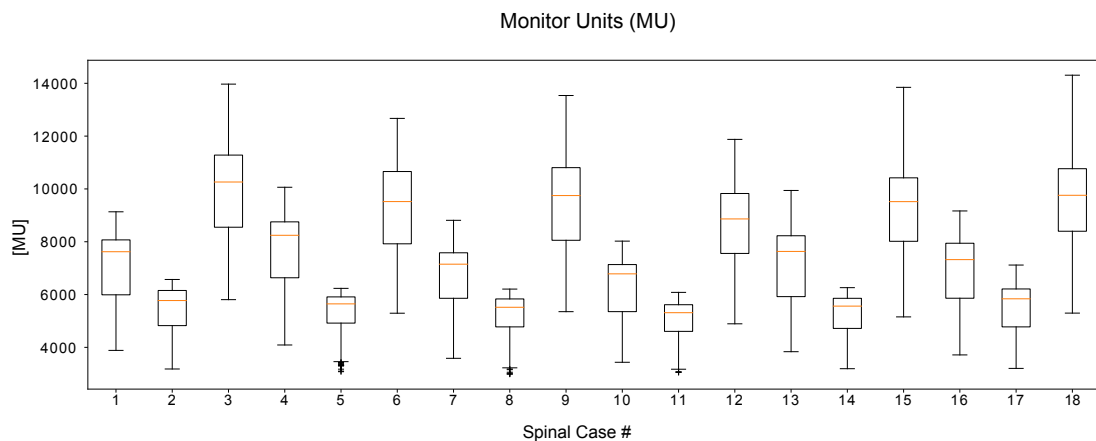
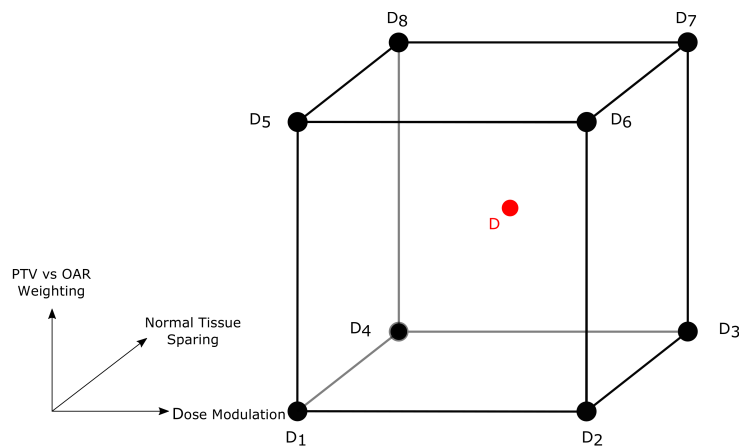


Figure 3.31: Boxplots of monitor units (MU), 343 optimizations per case





**Figure 3.32:** Illustration of the trilinear interpolation problem mapped onto the 3D space spanned by the of VMAT optimization sliders

### 3.5 Approximation of Pareto surfaces by trilinear dose interpolation

The previous analysis showed that variation of the three composite parameters 'Weighting', 'Normal Tissue Sparing' and 'Modulation' can have huge influence on the dosimetric characteristics of the resulting treatment plans in some cases, while in other cases (e.g. some of the cranial examples) the results didn't change much.

The optimization results in cases with not very challenging geometry (convex PTV in a reasonable distance from any OAR), such as cranial case #8, only change marginally with parameter variation, and given the continuous nature of the parameter space, a low-resolution approximation of the Pareto surfaces of the solution space might be a completely sufficient strategy for getting a good overview of the solution space.

In case of more complicated geometries and more challenging optimization constraints a higher resolution is likely to be necessary.

In extremely challenging cases like cranial case #1, a lot of the results might be ruled out beforehand (such as e.g. plans with  $CI \geq 2$ ). In these cases it might be convenient to identify the corresponding regions of the solution space in order to make the planning process more time efficient. This issue will be addressed later.

In this section a solution-space approximation approach using trilinear dose interpolation in two different resolutions is presented, and the approximation accuracy in the cranial and spinal test cases is evaluated. The theory of the underlying methodology has been introduced in section 2.4. The anchor points  $D_i$  and the value  $D$  now correspond to three-dimensional dose distributions in a specific patient geometry as a result of optimization with a certain choice of the three composite optimization parameters 'Weighting', 'Normal Tissue Sparing' and 'Modulation', which make up the dimensions  $x, y, z$  of the parameter space. Due to the composite nature of these three optimization parameters, where some optimization aspects are affected by more than a single slider, the three axes can not be automatically assumed to be completely independent. Nevertheless, for easier comprehension the coordinate axes in the following illustrations are drawn orthogonally. Figure 3.32 shows the mapping of the approximation of dose distributions onto the affine trilinear interpolation space. The three slider




Figure 3.33: Linear interpolation along *Dose Modulation* axis

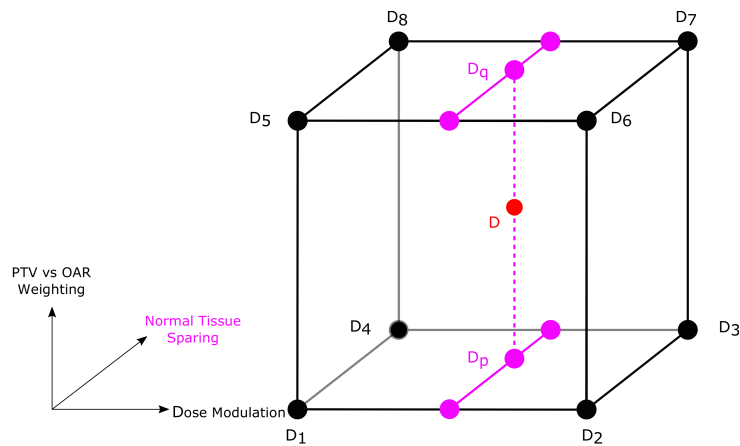


Figure 3.34: Linear interpolation along *Normal Tissue Sparing* axis

positions are mapped onto three coordinate axes. The order of the three interpolation processes described below is arbitrary. Mathematically all sequences lead to the same coefficients  $\alpha_i$  given in 2.4.13, with  $x, y, z$  replaced with the normalized slider values for *Target-OAR-Weighting*, *Normal Tissue Sparing* and *Dose Modulation*, respectively.

In the first step, 4 linear interpolations along the *Dose Modulation* axis are performed, yielding  $D_{00}, D_{01}, D_{10}, D_{11}$ , as shown in figure 3.33.

Starting from those 4 intermediate results, two additional interpolations are performed, this time strictly in *Normal Tissue Sparing* direction, which yields  $D_p, D_q$  illustrated in figure 3.34. Finally, figure 3.35 shows that a single linear interpolation process, this time in *Target-OAR-Weighting* direction yields the sought after approximation of the dose distribution  $D$ .

### 3.5.1 Considerations concerning the interpolation resolution

The approach shown in section 3.32, where only the dose distributions obtained by optimizing for the extremal slider positions are considered as anchor points for interpolation, is likely to produce rough estimates for the central slider range only. This potential issue is aggravated by the observation that the dose distributions obtained for the extremal slider positions often

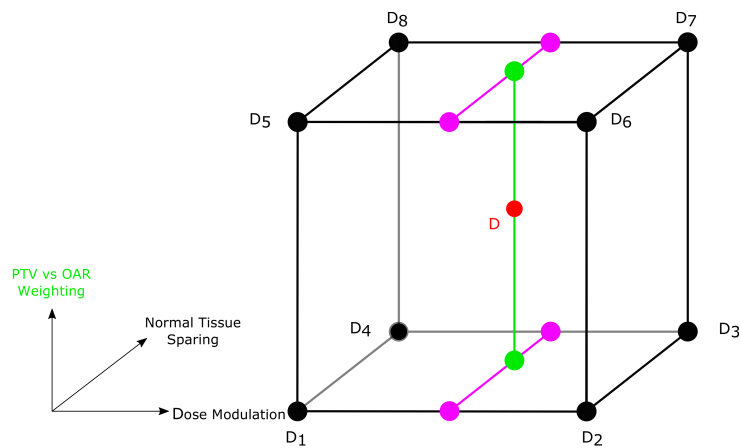


Figure 3.35: Linear interpolation along *PTV vs OAR Weighting* axis




Figure 3.36: Pattern of additionally pre-calculated slider positions (red) embedded in already available results (black). 2D slice of the parameter space

differ strongly from any results obtained for parameters within the central range. However, given the fact that only 8 optimizations have to be performed, taking approximately 2 min each, this has to be accepted as the price to pay for being able to get a grasp of achievable optimization results quickly. In a lot of the cranial example cases, accepting this fact seems justified. Still, in order to achieve better results, it seems appropriate to take at least some additional optimization results into account. This is done by choosing a finer sampling rate of the slider positions, as shown in figure 3.36.

The number of initial optimization runs is increased from 8 to 27, dividing the initial interpolation cube into 9 sub-cubes. Only optimization results from the 8 pre-optimized slider position configurations closest to the one to be approximated are taken into account. This is shown in figure 3.37




Figure 3.37: Optimization results taken into account for approximation at cross shown in green. 2D slice of the parameter space

### 3.5.2 Mathematical behavior of characteristic DVH values with respect to trilinear interpolation

Even though the concept of trilinear interpolation of dose distributions might seem plausible and straightforward, the effects on dose volume metrics requires further investigation.

In the following sections, an experimental study of the suitability of trilinear dose interpolation for the prediction of inverse optimization dose results will be presented. For a realistic approximation of a physical phenomenon such as the distribution of ionizing radiation dose in biological tissue, it is important to make sure that the chosen mathematical approach is methodologically sound. Interpolation with higher order polynomials, for example, often has unwanted side effects. It can create high-amplitude oscillations not related to any physical phenomenon. This effect is called Runge phenomenon, it always occurs in equidistant interpolation with higher order polynomials[23] [28]. In terms of dose distributions this means that we have to make sure that e.g. the mean dose of an interpolation  $D_{int}$  of two doses  $D_1, D_2$  is not unexpectedly much higher or lower than the mean doses of the two initial dose distributions. Let

$$D_i = \left( V_1^{(i)} \dots V_n^{(i)} \right)^T \quad (3.5.1)$$

denote the set of dose values in a three dimensional tissue model discretized by an arbitrary grid fixed for all  $i$ . Defining  $V_j^{(i)}$  as the radiation dose received by voxel  $j$  in  $D_i$  yields for the dose value in voxel  $V_j$  of the interpolated dose distribution:

$$V_j = \sum_{i=1}^8 \alpha_i V_j^{(i)} \quad (3.5.2)$$

The dosimetric criteria for evaluating the treatment plan quality analyzed in this thesis can be categorized into convex criteria, such as mean dose values, and non-convex criteria, such as dose/volume constraints and minimum or maximum dose values. Both categories and their behavior with respect to trilinear interpolation have to be considered. With the definitions given above, we can make the following observations.

#### Mean dose

It holds the following:

**Theorem 1.** *The mean dose of the interpolation result is equal to the interpolation of the mean doses of the dose distributions used for the interpolation.*

$$\frac{1}{n} \sum_{j=1}^n V_j = \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^8 \alpha_i V_j^{(i)} = \sum_{i=1}^8 \alpha_i \frac{1}{n} \sum_{j=1}^n V_j^{(i)} \quad (3.5.3)$$

*Proof.*

$$\begin{aligned}
\frac{1}{n} \sum_{j=1}^n V_j &= \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^8 \alpha_i V_j^{(i)} = \frac{1}{n} \left( \sum_{i=1}^8 \alpha_i V_1^{(i)} + \sum_{i=1}^8 \alpha_i V_2^{(i)} + \dots + \sum_{i=1}^8 \alpha_i V_n^{(i)} \right) = \\
&= \sum_{i=1}^8 \frac{1}{n} \alpha_i V_1^{(i)} + \sum_{i=1}^8 \frac{1}{n} \alpha_i V_2^{(i)} + \dots + \sum_{i=1}^8 \frac{1}{n} \alpha_i V_n^{(i)} = \\
&= \left( \frac{1}{n} \alpha_1 V_1^{(1)} + \frac{1}{n} \alpha_2 V_1^{(2)} + \dots + \frac{1}{n} \alpha_8 V_1^{(8)} \right) \\
&+ \left( \frac{1}{n} \alpha_1 V_2^{(1)} + \frac{1}{n} \alpha_2 V_2^{(2)} + \dots + \frac{1}{n} \alpha_8 V_2^{(8)} \right) + \\
&+ \dots + \left( \frac{1}{n} \alpha_1 V_n^{(1)} + \frac{1}{n} \alpha_2 V_n^{(2)} + \dots + \frac{1}{n} \alpha_8 V_n^{(8)} \right) = \\
&= \frac{1}{n} \alpha_1 \left( V_1^{(1)} + \dots + V_n^{(1)} \right) + \dots + \frac{1}{n} \alpha_8 \left( V_1^{(8)} + \dots + V_n^{(8)} \right) = \\
&= \sum_{i=1}^8 \alpha_i \underbrace{\frac{1}{n} \sum_{j=1}^n V_j^{(i)}}_{\text{mean dose of } D_i}
\end{aligned} \tag{3.5.4}$$

□

### Max dose

**Theorem 2.** Let the maximal dose value received by a structure consisting of  $n$  voxels be defined as

$$\max_{j=1\dots n} V_j \tag{3.5.5}$$

Then in general

$$\max_{j=1\dots n} \left( \sum_{i=1}^8 \alpha_i V_j^{(i)} \right) \neq \sum_{i=1}^8 \alpha_i \max_{j=1\dots n} V_j^{(i)} \tag{3.5.6}$$

*Proof.* Let  $D_i = \left( V_1^{(i)} \dots V_n^{(i)} \right)^T$ ;  $i = 1 \dots 8$  be the affine base used for interpolation with  $D_{int}$  being the interpolated dose distribution at the affine coordinates  $\left( \frac{1}{2}, 0, 0 \right)$ . Then

$$D_{int} = \frac{1}{2} (D_1 + D_2) \tag{3.5.7}$$

Let  $p, q \in \{1, \dots, n\}$ ;  $p \neq q$  be two indices,

$$D_1 := \left( V_1^{(1)} \dots V_n^{(1)} \right)^T \text{ with } V_p^{(1)} := a \in \mathbb{R}^+ := \max_{j=1\dots n} \left( V_j^{(1)} \right)$$

$$V_j^{(1)} := b \in \mathbb{R}^+ < a \forall j \neq p$$

$$D_2 := \left( V_1^{(2)} \dots V_n^{(2)} \right)^T \text{ with } V_q^{(2)} := c \in \mathbb{R}^+ := \max_{j=1\dots n} \left( V_j^{(2)} \right) \gg a$$

$$V_p^{(2)} = 0; V_j^{(2)} := d; 0 < d < b \forall j \neq p, q \Rightarrow d < b < a < c$$

Then the maximum dose value in the interpolated dose distribution

$$D_{int} = \frac{1}{2} \left( (b+d) \cdots (b+d) \underbrace{a}_{\text{index } p} (b+d) \cdots (b+d) \underbrace{(b+c)}_{\text{index } q} \right)^T$$

$$\text{is } \max_{j=1 \dots n} D_{int} = \frac{1}{2} (b+c).$$

$$\text{However, } \frac{1}{2} \left( \max_{j=1 \dots n} (V_j^{(1)}) + \max_{j=1 \dots n} (V_j^{(2)}) \right) = \frac{1}{2} (a+c) > \frac{1}{2} (b+c). \quad \square$$

### Min dose

**Theorem 3.** Let the minimal dose value received by a structure consisting of  $n$  voxels be defined as

$$\min_{j=1 \dots n} V_j \quad (3.5.8)$$

Then in general

$$\min_{j=1 \dots n} \left( \sum_{i=1}^8 \alpha_i V_j^{(i)} \right) \neq \sum_{i=1}^8 \alpha_i \min_{j=1 \dots n} V_j^{(i)} \quad (3.5.9)$$

*Proof.* The proof is done analogously to theorem 2 with  $\min_{j=1 \dots n} V_j := \max_{j=1 \dots n} (-V_j)$ .  $\square$

### General dose/volume values, e.g. 99% dose

**Theorem 4.** Dose/Volume values like the 99% Dose  $D_{99}$  generally do not behave linearly for the interpolated dose distribution, e.g.

$$D_{99} \left( \sum_{i=1}^8 \alpha_i D_i \right) \neq \sum_{i=1}^8 \alpha_i D_{99}(D_i) \quad (3.5.10)$$

*Proof.* Proof by contradiction with artificially constructed dose distributions  $D_1$  and  $D_2$  similar to the proof of theorem 2.  $\square$

### Boundedness of min and max dose

$D_{int} = \sum_{i=1}^8 \alpha_i D_i$ ;  $0 \leq \alpha_i \leq 1 \wedge \sum_{i=1}^8 \alpha_i = 1$  (convex combination)

$$\max \left( \sum_i \alpha_i D_i \right) \leq \max \left( \sum_i \alpha_i \max(\max D_i) \right) \stackrel{\sum \alpha_i = 1}{\leq} \max(\max(D_i)) \quad (3.5.11)$$

$$\max \left( \sum_i \alpha_i D_i \right) \geq \max \left( \sum_i \alpha_i \min(\max D_i) \right) \stackrel{\sum \alpha_i = 1}{\geq} \min(\max(D_i)) \quad (3.5.12)$$

Thus, max and min are bounded by the points used for interpolation and max and min show up on the boundary of the interpolation area. In terms of the physical background of radiation dose distributions, this is a reasonable expectation. Higher order interpolation might cause oscillation effects leading to unrealistic behavior.

### 3.5.3 Trilinear interpolation with 8 anchor points

The following section provides a comparative overview of the results obtained from approximation with 8 anchor plans obtained for the extremal slider positions only. According to the categorization and mathematical analysis of DVH values with respect to trilinear dose interpolation, the PTV dose coverage, OAR mean dose, and the OAR DVH constraint point for both cranial and spine example cases were selected for comparison. Additionally, Pareto surfaces for selected cases (cranial case #1 and spine case #4) are presented in order to give a visualization of the approximation results. For a specific DVH value/dosimetric criterion  $C$ , the boxplots show the difference

$$\Delta(C) := C_{\text{approximation}} - C_{\text{actualOptimization}} \quad (3.5.13)$$


A negative deviation means the approximation is smaller than the actual value, a positive deviation means the approximated value exceeds the actual value.

The example Pareto surface plots for each criterion clearly show the strong smoothing effect of such a coarse anchor grid, as well as another important property: The set of anchor plans only contains results for either maximum or minimum value of the 'Modulation' parameter. The plots show approximations for the central value in the 'Modulation' range. As a consequence, even the corner points of the approximated surfaces do not match the exact values.


#### PTV coverage cranial

Figure 3.38 shows  $\Delta(V_{99\%})$ , e.g. the deviation between approximated and actual prescribed dose coverage for the cranial sample cases. The maximum deviation occurred in case #1 with an overestimation of the actual value by 26% in a case where both 'Modulation' and 'Normal Tissue' sparing were set to the smallest possible value. Cranial case #1 has already been identified as being extremely challenging for the optimizer in the previous chapter, especially with a low degree of 'Modulation'. In general, the distribution of overestimation and underestimation turned out to be quite balanced, with the median deviation of 1.18% for case #1 and a general slight underestimation of  $-1.31\%$  (median) for all 44 cranial cases.

This is also visible in figure 3.39 showing the approximated (left) and actual (right) Pareto surfaces with respect to the 'Weighting' and 'Normal Tissue Sparing' parameters ('Modulation' is fixed at the central position). Even with 8 anchor plans the general shape is reproduced, however, the information provided by only 8 pre-calculated dose distributions is not sufficient to accurately reflect the steep gradient in the middle of the 'Normal Tissue Sparing' range.



**Figure 3.38:** Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases.



**Figure 3.39:** Pareto surfaces for the PTV dose coverage in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position


### PTV coverage spine

Figure 3.40 shows  $\Delta(V_{95\%})$ , e.g. the deviation between approximated and actual prescribed dose coverage for the spinal sample cases. The maximum deviation occurred in case #5 with an overestimation of the actual value by almost 27%, again in a case where both 'Modulation' and 'Normal Tissue' sparing were set to the smallest possible value. In case #5 the PTV consists of a complete lumbar vertebra with a high volume and high concavity value. For all spine cases, the median of all deviations for the PTV coverage was  $-2.27\%$  which means a slight underestimation.


Figure 3.41 shows the approximated (left) and actual (right) Pareto surfaces with respect to the 'Weighting' and 'Normal Tissue Sparing' parameters ('Modulation' is fixed at the central position) for spine case #4. Compared to the cranial example, the surfaces clearly show that 8 anchor plans do not provide enough information for a suitable approximation of the solution space for complex patient geometries. Since the approximation is done for cases where the 'Modulation' slider is in the medium position and all available optimization results have obtained



for a minimal or maximal degree of modulation only, even the corner points of the solution surface are not corresponding to the actual calculation results.




**Figure 3.40:** Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases




**Figure 3.41:** Pareto surfaces for the PTV dose coverage in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position.

### OAR mean dose brainstem

Figure 3.42 shows  $\Delta(D_{mean})$  for the OAR 'Brainstem', e.g. the deviation between approximated and actual OAR mean dose for the cranial sample cases. Since this value is very small for both actual calculation and approximation in a lot of the cranial cases a comparison of the absolute deviation is more reasonable than looking at percentages. Cranial case #1 has already been identified as being extremely challenging for the optimizer in the previous chapter, especially with a low degree of 'Modulation', and also this time it turned out that in some cases, the mean dose for the brainstem was overestimated by more than 2 Gy. In general, especially for e.g. cases #6-17 the deviation between estimated and actual value was only marginal. Still the nature of how mean values behave in the context of trilinear interpolation has an enormous smoothing effect and erases a lot of information, which can be observed in figure 3.43.




**Figure 3.42:** Boxplots of difference between mean dose in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases.




**Figure 3.43:** Pareto surfaces for the Mean Dose of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position

### OAR mean dose spinal canal

Figure 3.44 shows  $\Delta(D_{mean})$  for the OAR 'Spinal Canal', e.g. the deviation between approximated and actual OAR mean dose for the spinal sample cases. Due to the higher complexity of the patient geometries the deviations in the case of 8 anchor plans were significantly higher than in the cranial cases. Again interpolation of the mean values creates a very smooth evenly sloped surface, which does not reflect the actual solution space shown in figure 3.45, in some cases the actual value was overestimated by more than 80%. The median deviation for all 18 spinal cases was still 19.4%.



**Figure 3.44:** Boxplots of difference between mean dose in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases.




**Figure 3.45:** Pareto surfaces for the Mean Dose of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position.

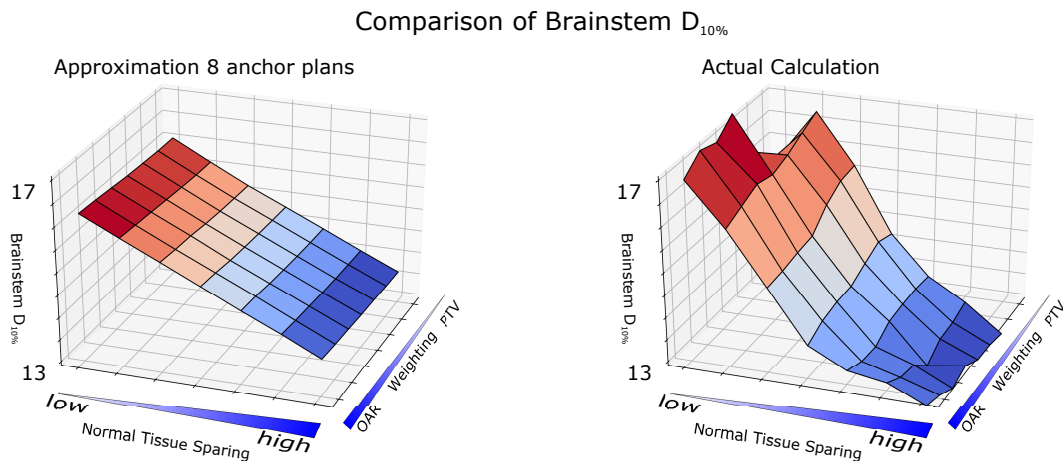
### OAR constraint dose brainstem

Figure 3.46 shows  $\Delta(D_{10\%})$  for the OAR 'Brainstem', e.g. the deviation between approximated and actual prescribed dose at the DVH constraint point for the cranial sample cases. The maximum deviation in terms of absolute values occurred in case #5 with an overestimation of the actual value by 15.2% of the actual value. In general the mean median deviation was 1.8% for all 44 cranial cases.

The Pareto surfaces in figure 3.47 show that most significant deviations occur in the range of low 'Normal Tissue Sparing' values.



**Figure 3.46:** Boxplots of difference between  $D_{10\%}$  in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases.




**Figure 3.47:** Pareto surfaces for  $D_{10\%}$  of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position.

### OAR constraint dose spinal canal


Figure 3.48 shows  $\Delta(D_{0.25cm^3})$  for the OAR 'Spinal Canal', e.g. the deviation between approximated and actual dose at the DVH constraint point for the spine sample cases. The maximum deviation in terms of absolute values occurred in case #3 with an overestimation of the actual or almost 6Gy, corresponding to a relative deviation 54.8%. In general the mean median deviation was 17.5% for all 18 spine cases.

The Pareto surfaces in figure 3.49 show that the most significant deviations occur in the range of low 'Normal Tissue Sparing' values. This information is lost if only results for extremal parameter values are taken into consideration for approximation.

In summation, the analysis showed that for very simple cranial cases, approximation using 8 anchor plans only might be sufficient, however, for more complicated geometries existing e.g. in the spinal cases, more information and a higher resolution of the solution parameter space is needed.



**Figure 3.48:** Boxplots of difference between  $D_{0.25cm^3}$  in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases.



**Figure 3.49:** Pareto surfaces for  $D_{0.25cm^3}$  of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position.

### 3.5.4 Trilinear interpolation with 27 anchor points


Originally the sliders for the three composite objectives 'Weighting', 'Normal Tissue Sparing' and 'Modulation' are conceptualized in a way that the chosen values will be somewhere in the middle of the range and the extremal positions represent rather extreme cases. The approximation approach with 8 extremal anchor plans has the advantage of only a limited number of necessary optimizations before navigation of the Pareto surface can be performed and approximations can be made for the complete available parameter range. The evaluation of the approximation quality achieved with such an approximation base (as performed in the previous chapter) showed that for non-trivial challenging patient geometries, more information about the characteristics of the optimization results in the center of the parameter range is needed. Section 2.4 and 3.5 introduced a method using 27 pre-optimized treatment plans, including results obtained for the central slider positions. This discretization of the parameter

space is shown in figure 3.37 as a 2D projection, the 3D-subdivision can be seen in figure 2.13. The approximation results using this approach are presented in the following section.


### PTV coverage cranial

Figure 3.50 shows  $\Delta(V_{99\%})$ , e.g. the deviation between approximated and actual prescribed dose coverage for the cranial sample cases. The maximum deviation occurred again in case #1 with an overestimation of the actual value by 17.6% in the same case where both 'Modulation' and 'Normal Tissue' sparing were set to the smallest possible value. This is a big improvement compared to the 26% deviation observed for the approximation with 8 plans. For not so extreme patient geometries the approximation was even better, and the median deviation for all cranial cases was only -0.33%.

This is also visible in figure 3.51 showing the approximated (left) and actual (right) Pareto surfaces with respect to the 'Weighting' and 'Normal Tissue Sparing' parameters ('Modulation' is fixed at the central position), only for 'Weighting' values focusing completely on the PTV some deviations can still be observed.



**Figure 3.50:** Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases.




**Figure 3.51:** Pareto surfaces for the PTV dose coverage in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.


### PTV coverage spine

The approximation results for the PTV coverage in the spinal cases shown in figure 3.52 still show a lot of outliers, but the general deviation error improved. For single parameter configurations values still deviated up to a maximum of 21% (case #5) for certain regions of the parameter space, but the mean median deviation for all 18 patient geometries was reduced to -0.68%.

The regions with the highest approximation errors can be seen on the left of the Pareto surfaces in figure 3.53 with low 'Normal Tissue Sparing' values where the approximated surface is still unable to fully reproduce the steepness of the original slope.




**Figure 3.52:** Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases.



**Figure 3.53:** Pareto surfaces for the PTV dose coverage in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.

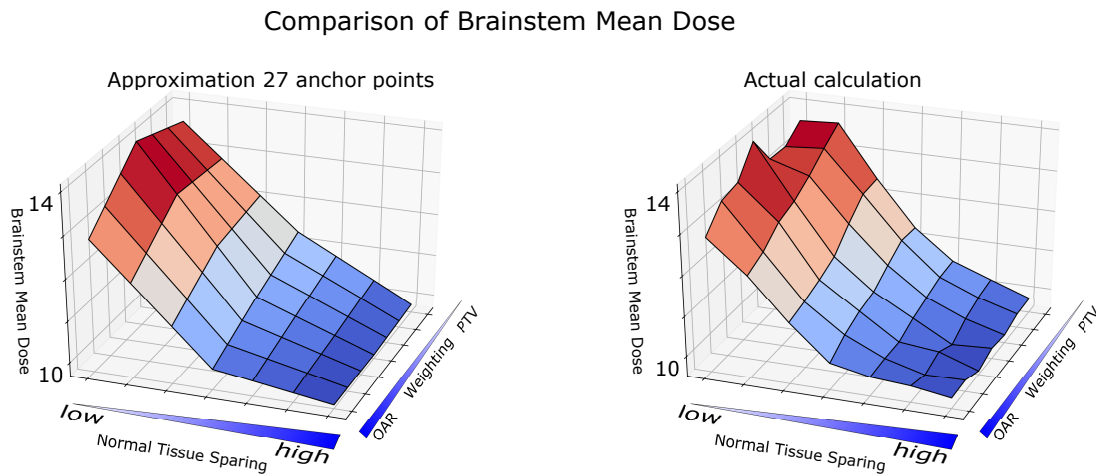
### OAR mean dose brainstem

Figure 3.54 shows  $\Delta(D_{mean})$  for the OAR 'Brainstem', e.g. the deviation between approximated and actual OAR mean dose for the cranial sample cases for approximation with 27 anchor plans. Since this value is very small for both actual calculation and approximation in a lot of the cranial cases, a comparison of the absolute deviation is more reasonable than looking at percentages. This time cranial case #5 showed the maximum absolute approximation error, overestimating the dose by about 1.5 Gy, or 21% when the weighting was fully focused on OAR sparing and only marginal levels of modulation and normal tissue sparing were allowed. This is a rather unusual and not very realistic combination, since it is plausible that in a complex patient geometry good OAR sparing is only possible with complex MLC shapes, and normal tissue sparing is basically another means of OAR sparing on its own. In general, in more than 80% of the cases the deviation between estimated and actual value was less than 0.5 Gy. The smoothing effect of mean dose approximation is just not able to accurately reproduce the detailed structure of the actual Pareto surface, which can be observed in figure 3.55 in the region with low 'Normal Tissue Sparing' value.



**Figure 3.54:** Boxplots of difference between mean dose in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases.






**Figure 3.55:** Pareto surfaces for the Mean Dose of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.

### OAR mean dose spinal canal

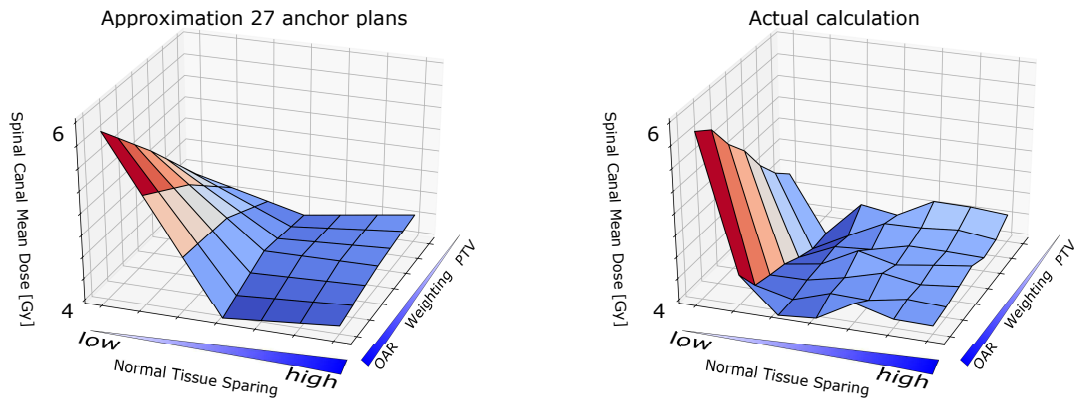
Figure 3.56 shows  $\Delta(D_{mean})$  for the OAR 'Spinal Canal', e.g. the deviation between approximated and actual OAR mean dose for the spinal sample cases. Due to the higher complexity of the spinal patient geometries, even in the case of 27 anchor plans, the deviations were significantly higher than in the cranial cases.

Even with information from results obtained in the central parameter range, the interpolation of the mean values could not completely reproduce the extreme slope observed between the two lowest possible 'Normal Tissue Sparing' values in the actual solution space shown in figure 3.57. In some cases the actual value was still overestimated by more than 60%. The nature of the solution space suggests that in such cases the narrowing of the slider range with a finer resolution might cause a higher variation in the solution space. However, the mean median deviation for all 18 spinal cases still decreased from 19.4% to 2.76%.



**Figure 3.56:** Boxplots of difference between mean dose in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases.

## Comparison of Spinal Canal Mean Dose




**Figure 3.57:** Pareto surfaces for the Mean Dose of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.


## OAR constraint dose brainstem

Figure 3.58 shows  $\Delta(D_{10\%})$  for the OAR 'Brainstem', e.g. the deviation between approximated and actual prescribed dose at the DVH constraint point for the cranial sample cases and 27 pre-optimized treatment plans. The maximum deviation in terms of absolute values occurred in case #1 with an underestimation of the actual value by 10.9% of the actual value. In general, the mean median deviation was  $-0.31\%$  for all 44 cranial cases.

The Pareto surfaces in figure 3.59 show that most significant deviations occur in the range of low 'Normal Tissue Sparing' values. In addition, for 'Weighting' values focusing on the PTV, the original shape of the surface in 'Normal Tissue Sparing' direction is concave while the same region of the approximated Pareto surface is convex due to the initial dose increase.



**Figure 3.58:** Boxplots of difference between  $D_{10\%}$  in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases.




**Figure 3.59:** Pareto surfaces for  $D_{10\%}$  of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.


### OAR constraint dose spinal canal

Figure 3.60 shows  $\Delta(D_{0.25cm^3})$  for the OAR 'Spinal Canal', e.g. the deviation between approximated and actual dose at the DVH constraint point for the spine sample cases and 27 pre-optimized treatment plans. The maximum deviation in terms of absolute values occurred again in case #3 with an overestimation of the actual or more than 4Gy, corresponding to a relative deviation 52%. In general, the mean median deviation decreased from 17.5% to 2.36% for all 18 spine cases.

The Pareto surfaces in figure 3.61 show that the most significant deviations still occur in the range of low 'Normal Tissue Sparing' values due to the extremely steep dose fall-off.



**Figure 3.60:** Boxplots of difference between  $D_{0.25cm^3}$  in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases.



**Figure 3.61:** Pareto surfaces for  $D_{0.25cm^3}$  of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position.

### 3.5.5 Limitations and advantages

In this chapter the optimization of cranial and spinal VMAT radiotherapy treatment plans, using a low number of three partly contradicting composite objectives, focusing on the trade-off between PTV dose coverage and sparing of specific OARs, the general protection of non-malignant tissue surrounding the PTV, and the delivery- and quality assurance complexity of the resulting treatment plans, was considered from a multi-criteria optimization perspective. The continuous 3D parameter space was discretized into 343 points. The three dimensions comprise the numerical values of the composite objectives 'Weighting', 'Normal Tissue Sparing' and 'Modulation', the range of each dimension being represented by 7 equidistant positions, including the extremal and central positions. Thus, for each individual patient geometry, this setup provides 343 optimization parameter sets leading to different dose distributions. For 44 cranial and 18 spine patient geometries optimizations for all possible parameter sets were performed, and the resulting dose distributions were analyzed. Exploratory optimization for selected geometries with a parameter resolution twice as high was performed, and it could be observed that the generated high resolution Pareto surfaces did not reveal any previously obscured discontinuities (see figure 3.20). The characteristics of the solution space were described by the achievable variability of a selection of dosimetric criteria such as the coverage of the PTV with the desired radiation dose, or the maximal dose delivered to a certain partition of the chosen OARs.

This analysis revealed that the dose distributions did not vary much, if the patient geometry (distance, volume and concavity/convexity of the involved organ structures) did not create a very taxing problem for the optimizer. Such cases e.g. consist of convex objects with moderate volumes and a distance of 1cm or more from each other. In this case, if the optimizer finds a dose distribution satisfying all imposed dosimetric constraints, further optimizations with different parameter sets are not likely to yield more favorable dose distributions.

For large objects of complex shape (and thus a high concavity value) located close to each other or even touching, the resulting dose distributions were influenced by the choice of the optimization parameter sets to a much higher degree. However, despite this great variation in some cases, none of the resulting dose distributions was able to fulfill all dosimetric constraints

simultaneously. This situation tends to occur e.g. if the PTV consists of a complete vertebra and the most critical OAR is the Spinal Canal. In a clinical setting more detailed treatment planning strategies have to be applied in such a case, of course, and the three slider controls analyzed in this study are by far not the only means for customizing the optimization process in the Brainlab Elements treatment planning system. Here, since one objective of this study was to observe the behavior of the optimizer in extremely challenging situations where the modification effects of the different parameter settings were expected to be most pronounced, clinical acceptability of the resulting dose distributions is not the primary objective.

Since calculation of the complete solution space is much too time consuming even on powerful hardware, a faster method for exploring the solution space was presented and applied to the cranial and spinal sample patient geometries, using trilinear dose interpolation with 8 and 27 pre-calculated dose distributions. Some dosimetric properties such as the relative dose coverage of the PTV in approximated and actual dose distributions matched pretty well even with only 8 pre-optimized plans available for interpolation. Increasing the number of pre-optimized plans to 27 improved the approximation accuracy, even for dosimetric properties where interpolation with 8 dose distributions had performed very poorly. The observed huge deviations mostly occurred in the following situations:

- Extremely contradicting parameter choice (e.g. maximal sparing of an irregularly shaped OAR without allowing the needed degrees of freedom in the MLC shapes)
- The transition from extremal parameter values to the next adjunct discretization step, e.g. minimal 'Modulation' value to the second smallest 'Modulation' value
- Any Pareto surface with a very inhomogeneous slope

Without additional information on the expected solution space, the fixed choice of the 27 equidistant parameter sets is a straight-forward, but obviously too rigid approach. Fortunately, this problem might be mitigated by the fact that the trilinear interpolation approach allows for an adaptive parametrization grid, where a higher parametrization resolution can be chosen where needed, and areas without much variability in the resulting dose distributions can be sufficiently described by a low number of grid points. The theoretical background of such an approach has already been presented in section 2.4.4. This leaves the problem of identifying parameter ranges with an expected high variation in the optimization results. A possible method for dealing with this problem will be presented in the next chapter.



# Chapter 4

## Exploration of the solution space via machine learning

## 4.1 Solution space exploration via machine learning: Motivation

The three composite parameters 'Weighting', 'Normal Tissue Sparing' and 'Modulation' in the Brainlab VMAT optimizer can be adjusted by three individual slider controls in the user interface and allow for a variety of different optimization results. For a quick exploration of the feasible solution space, a discretization of the available parameter ranges was performed, and a method for approximation of the to-be-expected dose distributions resulting from a choice of not-yet-optimized parameter values was presented. This approximation technique relies on a certain number of pre-optimized dose distributions, and the work presented in the previous section brought up the problem of how to choose the amount and specific parameter values of those so-called 'anchor plans', or 'anchor points'.

The analysis of the cranial and spinal test cases indicated that the characteristics of the solution space, and the Pareto surfaces of the corresponding MCO problem, apparently change with respect to the geometric properties of the patient model, such as object concavity/convexity, -volume, and -distances. Based on these observations, this chapter will introduce a method for predicting important parameter space regions using machine-learning models generated from previously obtained optimization data.

### 4.1.1 Regression vs classification in VMAT treatment planning

Section 2.6.2 introduced the two major classes of machine-learning problems: classification and regression. In the scope of predicting characteristics of optimization results in radiosurgery treatment plans by making inference based on a large number of previously obtained results, the most obvious approach for an application of machine learning technology in this area would be the direct numerical prediction of the desired dosimetric quality criteria (e.g. PTV dose coverage), in other words: a regression problem [7]. However, for the following reasons, the research presented in this work takes a classification approach:

Using regression for direct value prediction would demand a really high amount of available optimization results well spread over the complete range of patient geometries (in the scale of tens of thousands and above). Obtaining such a massive amount of data is an extremely challenging task and the methods and standards for the acquisition of such a massive database and the required data protection issues present a topic of current research on its own [25]. Generalization based on the data available for this study, to a degree that a reliable prediction of actual numerical values is possible, is an unrealistic expectation. This assumption was quickly confirmed after some exploratory research, which showed that the accuracy of the predicted values was often not significantly better than random generation of values in the expected range. Even extremely powerful learning methods, such as artificial neural networks (ANN) presented in section 2.6.4, with a relatively large number of hidden units, can not mitigate this problem since they are prone to extreme over-fitting (e.g. learning to recognize the training data, but being unable to generalize well on unknown data) in the absence of a sufficient amount of training- and test data.

However, even without being able to predict exact values and with only a limited amount of data available, knowledge on the behavior of the optimizer in previous planning procedures can help during the exploration of the solution space:



- Machine-learning driven identification of regions in the parameter space, where the expected dose distributions are likely to vary strongly, can be used to save the planner from performing optimizations in a region of parameter values, where the results are not likely to improve.
- If such regions of large result variability can be identified, the window in the parameter value range represented by each slider could be adjusted, so that the seven possible positions really focus on the interval where the optimizer is most sensitive to this parameter. This could e.g. mean spreading the whole slider range over half the initial numerical interval, if it is likely that any slider position exceeding the central position is going to have only a marginal effect on the resulting dose distribution.
- In terms of the approximation method presented in section 3.5 the spacing and number of the pre-calculated anchor plans could be adjusted to achieve an optimized representation of the solution space.

## **4.2 Identification of relevant sections of the parameter space by prediction of achievable value ranges**

In the following section, various approaches for the prediction of the effectiveness of the slider controls for 'Weighting', 'Normal Tissue Sparing' and 'Modulation' will be presented.

An important first step in each AI model generation process is the acquisition and pre-processing of the training data. Later on, this data will be used to create the feature matrix  $X$  and the target vector  $y$  introduced in section 2.6.

First of all, since this work strives to be indication-specific, in one partition of the data the treatment site is located in the cranial region, while the other partition deals with spinal cases. For a selection of patient cases in each treatment site a specific treatment setup such as shown in table 3.2, a pre-defined set of objects such as in table 3.1 and a list of plan quality criteria  $C_i$  are chosen. Geometrical features are extracted from the image-sets (CT, MR, ..) and the segmented organ structures, so that each patient case can be described as an  $m$ -dimensional vector of numerical values. Such features can be the volumes of certain structures, the spatial relationship represented in surface-to-surface distances, overlap-volume-histograms (OVH) and similar characteristics. Such features were applied by *Kazhdan et al.* [43] and *Shiraishi et al.* [72]. Altogether, in this work the following data aspects are of possible relevance:

- Values of the geometrical features describing the specific patient case
- Value of 'Weighting' parameter, e.g. Position of 'Weighting' slider
- Value of 'Normal Tissue Sparing' parameter, e.g. Position of 'Normal Tissue Sparing' slider
- Value of 'Modulation' parameter, e.g. Position of 'Modulation' slider

The detailed process for obtaining the training data is shown in figure 4.1. After definition of the treatment setup and the list of dosimetric criteria  $C_i$  the geometry vector is created for each patient. For all possible parameter configurations the optimization results are calculated

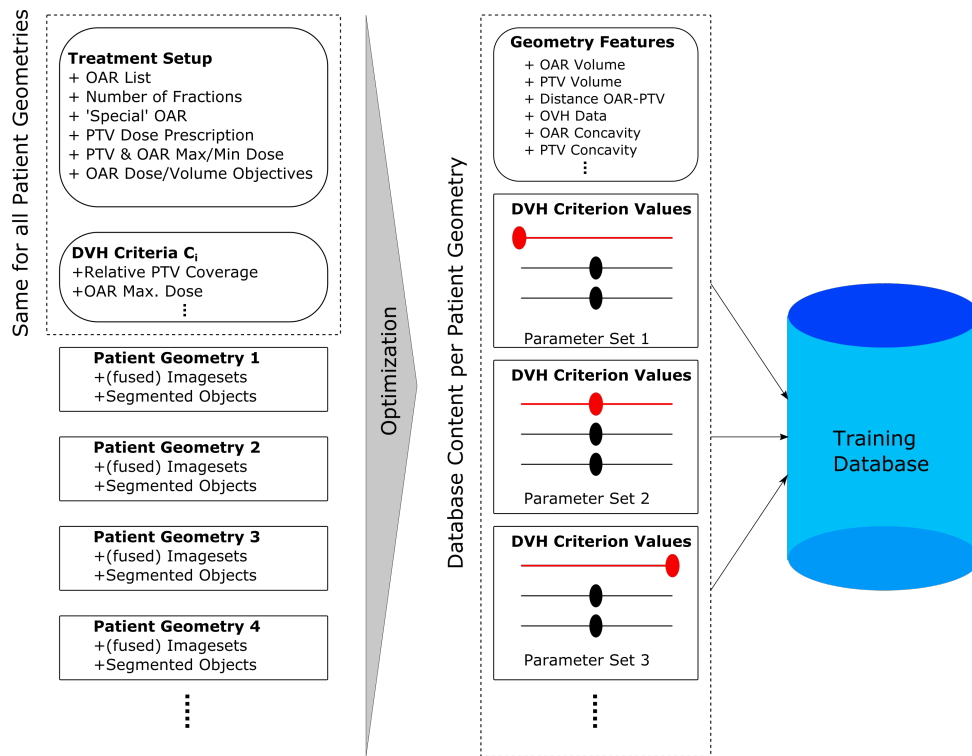


Figure 4.1: Generation of Training Data

and the values of the  $C_i$  are stored. In contrast to other approaches found in the literature [72], the complete 3D dose distributions are discarded after obtaining the  $C_i$ , since storage of thousands of entire 3D dose matrices, each taking up a couple of MB, would not be very memory efficient. The final training database entry for each patient with a certain treatment setup and list of characteristics  $C_i$  consists of the geometry vector and a list of all possible slider positions, together with the corresponding  $C_i$  values. All the various versions of the feature matrix  $X$  and target vector  $y$  used in this work are created from this data.

The idea behind the method presented in this thesis is the hypothesis that the behavior of the optimizer with respect to various patient geometries can be adequately represented with a high number of such data, and therefore, for a new case, the expected behavior can be predicted by inference from existing data. The structure of the AI-driven workflow developed in this work is shown in figure 4.2. As explained before, the ideal situation would be the availability of such a huge amount of data so that for a given patient geometry, treatment setup and list of criteria  $C_i$ , the achievable values of the  $C_i$ , together with the slider positions which are likely to produce those dose distributions, can be predicted. In such a scenario the user can immediately start exploring the most relevant part of the solution space. Due to the limited amount of data available ('limited' at least in terms of 'big data'), a more general prediction of the behavior of the optimizer is chosen here. It is a likely scenario that the treatment planner has chosen a specific set of slider positions, and is trying to change a specific feature of the resulting treatment plan (e.g. improving the PTV coverage). At this stage it would be very helpful to know whether trying out various additional positions of a specific parameter slider is likely to have the desired effect or not.

This consideration is the underlying concept of the presented algorithm. The left section of

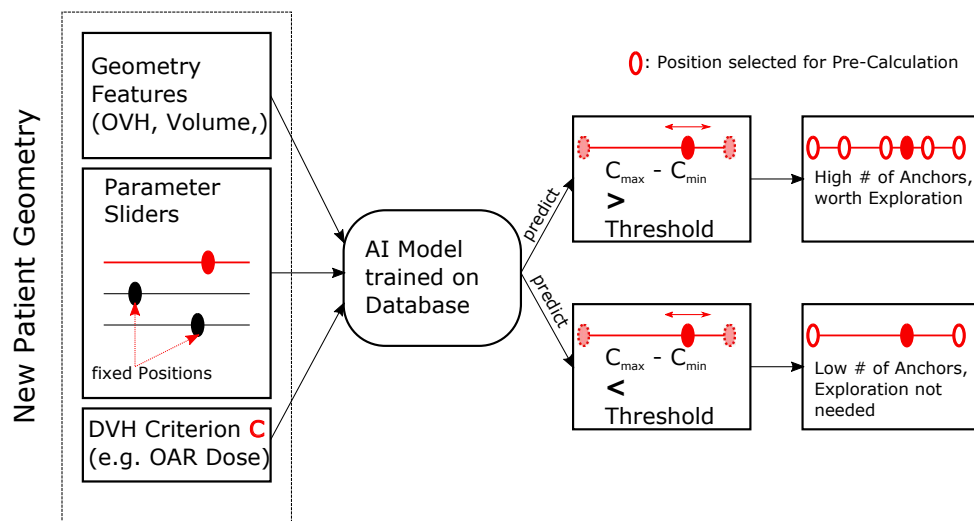


Figure 4.2: Schematic view of AI driven prediction algorithm

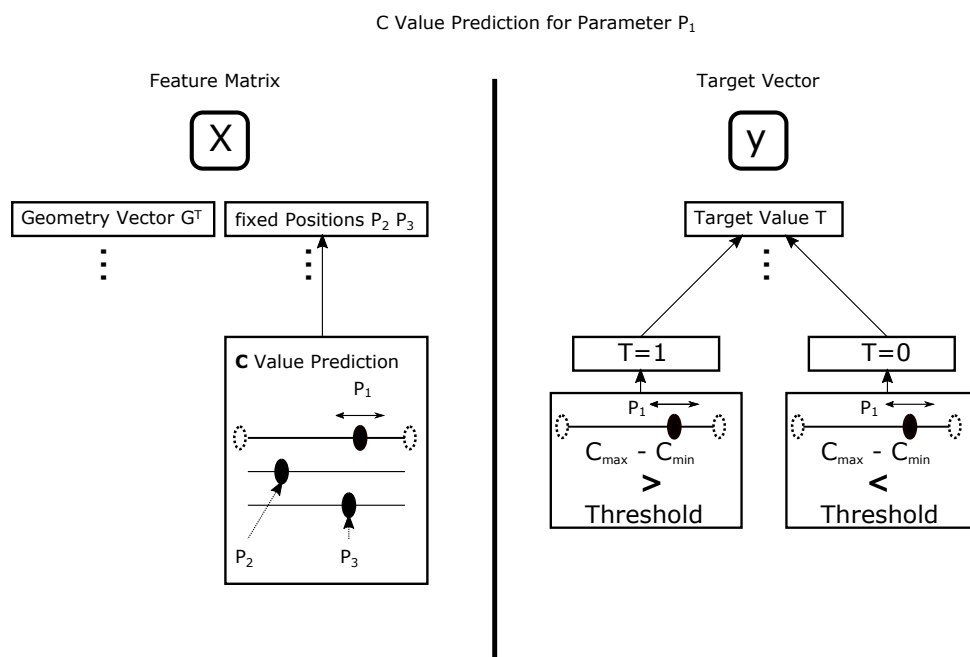
figure 4.2 illustrates the input values in this case and describes a certain moment during the planning process for a specific patient. The user wants to modify the value of a specific quality criterion  $C$  of the dose distribution resulting from optimization with the current slider settings via alteration of a specific slider control. By inference learned from the training data, the AI-model can give a prediction on the expected magnitude of the value range of  $C$  which might be achieved by trying out all remaining positions of this specific slider. The fact that the predicted range exceeds a certain threshold indicates that it is worth-while to perform more optimizations. If the range is likely to be under this threshold, the specific slider is unlikely to lead to any major changes of  $C$  in the current situation, and another strategy (e.g. focusing on a different slider first) has to be picked. Finding such a threshold, however, is not a straight-forward task. In the scope of this thesis, threshold values were picked by looking at the histograms of the available data.

In order to implement this prediction functionality, the following three types of AI models are evaluated in the remainder of this chapter:

- Support-vector-machine (SVM) with radial-basis-function (RBF) kernel
- Random-forest (RF) model
- Artificial neural network (ANN) with five hidden units using ReLu activation function and dropout[75] to avoid overfitting

This choice was made in order to be able to compare a relatively simple, fast and reliable model (SVM) as well as a very scalable state-of-the-art technology (ANN).

The RF has been chosen for its property to generalize well even for a limited amount of data and its unique capability of not only providing a classification result but also giving information on the relative importance of each feature in the decision process. The latter can be used for evaluation of the feature choice and feature design. The literature shows that sometimes also a combination of different classifiers leads to good results [89].



**Figure 4.3:** Schematic view of Feature Matrix  $X$  and Target Vector  $y$

### 4.2.1 Feature matrix and target vector generation

In order to train a specific AI algorithm for the task presented above, preparation of feature matrix  $X$  and target vector  $y$  is performed as shown in figure 4.3. Each row of  $X$  represents one sample optimization scenario, e.g. the geometric characteristics of the specific treatment case represented by the transpose of the geometry vector  $G$  and the parameter values  $P_1, P_2$  of the two sliders the user is not focused on at the current moment. In this work various versions of  $G$  have been studied, resulting in 3-dimensional, 7-dimensional and 11-dimensional feature vectors.


Since the prediction task is considered as a classification problem, the prediction result stored in the target vector  $y$  is either 1 or 0, depending on whether the achievable range of  $C_i$  values exceeds a certain threshold or not. The problem of finding a suitable threshold has to be addressed according to the chosen  $C_i$  and the treatment type (cranial vs spine) and treatment setup.

Once target vector and feature matrix are prepared, training and validation of the AI model can commence.

### 4.2.2 Generation of geometry vector $G$

As mentioned before, the geometry of a certain patient case can be described in varying detail. Focusing only on the PTV and the most relevant OAR, the following obvious values provide some important information about the spatial relationship and and geometric properties of these objects:

- Volume
- Concavity/convexity 2.5.3



**Figure 4.4:** Overlap-volume-histograms of the 'Brainstem'-objects in the cranial patient geometries analyzed in this work

- Closest surface-to-surface distance

More detailed geometry information can be extracted by the overlap-volume method (OVH) presented in section 2.5.1. While in theory the complete OVH data of each object could be added to the geometry vector  $G$ , this would result in a very high-dimensional feature space demanding a high calculation workload and resulting in poor fitting results. Figure 4.4 shows an example of real-world OVH data, and it is obvious that a more concise way of describing this data is needed.

In this case the *distance* between PTV and brainstem is defined by the smallest radius value where the first overlap occurs. The *penumbra* is the difference between the radius of the first overlap and the smallest radius value where complete overlap occurs. In addition the average slope, the number of inflection points and a triple of slope values (mean slope in the first 20%, middle 60% and final 20% section of overlap) were used as geometry features. An example illustrating these features can be seen in figure 4.5. In the spinal cases, where the spinal canal and the spinal cord represent the OAR object receiving special attention, the inverse OVH 2.5.2 is used instead of the original OVH in order to mitigate the effect of varying cropping lengths of these objects. In this work three different choices of geometry vectors  $G$  are considered. The three versions used for the cranial cases are shown in table 4.1 and the geometry vectors for characterization of the the spinal cases can be seen in table 4.2. Adding the two parameter values  $P2$ ,  $P3$  to each version of the geometry vector yields the complete representation of each sample in the 3-dimensional, 7-dimensional or 11-dimensional feature space. While a high-dimensional feature vector allows for a detailed description of each sample, this advantage comes with a couple of inevitable downsides. Computation effort increases with dimension of the feature space, resulting in increased training time for the model. While this might still be acceptable, since application of the final model is a very fast computation even with a lot of features, the features themselves become a problem. Finding a large number of meaningful features is a hard task, and simply e.g. adding the complete OVH data sampled in




Figure 4.5: Overlap-volume-histogram and the extracted features

1mm radius steps for each object will create a very high-dimensional feature vector containing a lot of redundant information. Additionally, due to the 'curse of dimensionality' phenomenon (described e.g. in [44]), with increasing dimension exponentially more data is needed to sample the feature space adequately and guarantee proper applicability of machine learning models. In tasks like image-recognition, where abundant labeled training data is readily available, this is not as big of a problem as in medical applications where obtaining training data is not as easy. Thus, a trade-off between number of features and the required amount of training data has to be found.

Table 4.1: Geometry vector in cranial cases

1D Geometry Vector	5D Geometry Vector	9D Geometry Vector
PTV Concavity	PTV Concavity	PTV Concavity
-	Volume PTV	Volume PTV
-	Volume Brainstem	Volume Brainstem
-	Distance Brainstem-PTV	Distance Brainstem-PTV
-	OVH Penumbra Brainstem	OVH Penumbra Brainstem
-	-	Slope OVH Brainstem sect. 1
-	-	Slope OVH Brainstem sect. 2
-	-	Slope OVH Brainstem sect. 3
-	-	# of inflection points OVH Brainstem

Table 4.2: Geometry vector in spine cases

1D Geometry Vector	5D Geometry Vector	9D Geometry Vector
PTV Concavity	PTV Concavity	PTV Concavity
-	Volume PTV	Volume PTV
-	Volume Spinal Canal	Volume Spinal Canal
-	Distance Spinal Canal-PTV	Distance Spinal Canal-PTV
-	Inv. OVH Penumbra Spinal Canal	Inv. OVH Penumbra Spinal Canal
-	-	Slope Inv. OVH Spinal Canal sect. 1
-	-	Slope Inv. OVH Spinal Canal sect. 2
-	-	Slope Inv. OVH Spinal Canal sect. 3
-	-	# of inflection points Inv. OVH Spinal Canal


### 4.2.3 Analysis of cranial cases with brainstem as prioritized OAR

#### OAR constraint point dose w.r.t. weighting slider


The first quality criterion  $C$  concerning the OAR object 'Brainstem' analyzed in this section is the  $D_{10\%}$  of the OAR object. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value. In other words, the task at hand is to predict in a certain situation whether trying out more positions of the 'Weighting' slider are likely to have a big impact on this criterion.

The left plot in figure 4.6 shows the histogram of the  $D_{10\%}$  values of the OAR object 'Brainstem' achieved in all optimization results for all available 44 cranial patient geometries (PTV prescription dose 15 Gy). The value range achievable with changes in the 'Weighting' slider only can be seen in the center, and a plausible threshold postulating that the achievable values should differ by at least 1 Gy is chosen, which leads to the two classes shown on the right. An important observation is the fact that the two classes are extremely unbalanced, a situation which has to be addressed by oversampling techniques (which were briefly introduced in section 2.6). Without oversampling there are not enough instances of the underrepresented class in each batch of training data and proper training is not possible. After random-oversampling, the classes are balanced (figure 4.7).

Figure 4.8 shows the fitting results for the SVM, ANN and RF classifiers for both the original



**Figure 4.6:** Histograms of OAR  $D_{10\%}$  values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 1 Gy OAR Constraint Dose range threshold (right)



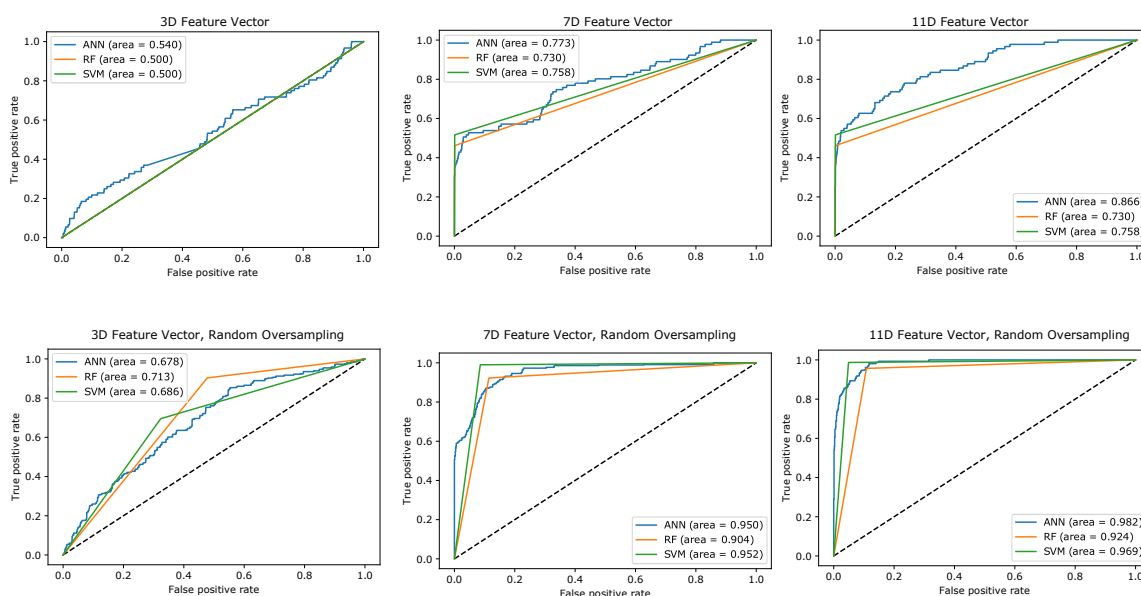
**Figure 4.7:** Histograms of the two sample classes resulting by introduction of a 1 Gy OAR Constraint point dose range threshold after random-oversampling

dataset and the balanced data after randomly duplicating samples from the underrepresented class using the random-oversampling technique. All models showed poor performance without

oversampling, in the case where the lowest number of features is used, only the ANN model received a prediction quality better than a random guess. With increasing number of features, the prediction quality improved. All three classifiers performed similarly. The increment from seven to eleven features yielded an increase in performance for the ANN model only.

Random-Oversampling had an enormous impact on model performances. With 7 features AUC values up to 0.95 could be achieved, and with 11 features the ANN model achieved an excellent AUC value of 0.98.

The comparison of the two available ANN models, one with 5 hidden units and one with 64



**Figure 4.8:** ROC curves for the OAR Constraint point dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

hidden units, confirmed a common rule-of-thumb in model construction, which postulates using roughly as many hidden units as the number of available features in the data. The 64 unit model achieved a *precision* of 0.96 and a *recall* of 1.0. The 5 unit model achieved a *precision* of 0.97 and a *recall* of 0.96. The training- and validation error and thus the quality of the model did not improve after already about 150 epochs.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that in this case the distance between PTV and OAR was the dominant feature, followed by the slope of the first section of the OAR OVH and the PTV volume.


### PTV rel. coverage w.r.t. weighting slider

The first quality criterion  $C$  concerning the PTV object analyzed in this section is the relative volume coverage of the PTV object with the prescribed dose of 15 Gy. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value and the prediction, whether trying out more positions of the 'Weighting' slider are likely to have a big impact on this value. The left plot in figure 4.9 shows the histogram of the relative PTV volume coverage values



achieved in all optimization results for all available 44 cranial patient geometries (PTV prescription dose 15 Gy). The value range achievable with changes in the 'Weighting' slider only can be seen in the center, and a threshold postulating that the achievable values should differ by at 0.05 (e.g. 5% of the PTV volume) is chosen, which leads to the two classes shown on the right. Class balancing is again achieved via random-oversampling.

Figure 4.10 shows the fitting results for the SVM, ANN and RF classifiers for both the original



**Figure 4.9:** Histograms of relative PTV volume dose coverage values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 5% PTV volume range threshold (right)

dataset and the balanced data after using the random-oversampling technique. As expected, without oversampling all models showed poor performance, in the case where the lowest number of features is used, prediction quality was not better than a coin-toss. With increasing number of features, prediction quality improved a lot and all three classifiers performed similarly. For predictions concerning this criterion, the increment from seven to eleven features did not yield an increase in performance.

Random-Oversampling had an enormous impact on model performances. With 7 features AUC values up to 0.99 could be achieved, and with 11 features the ANN model achieved a near-perfect AUC value of 0.998.

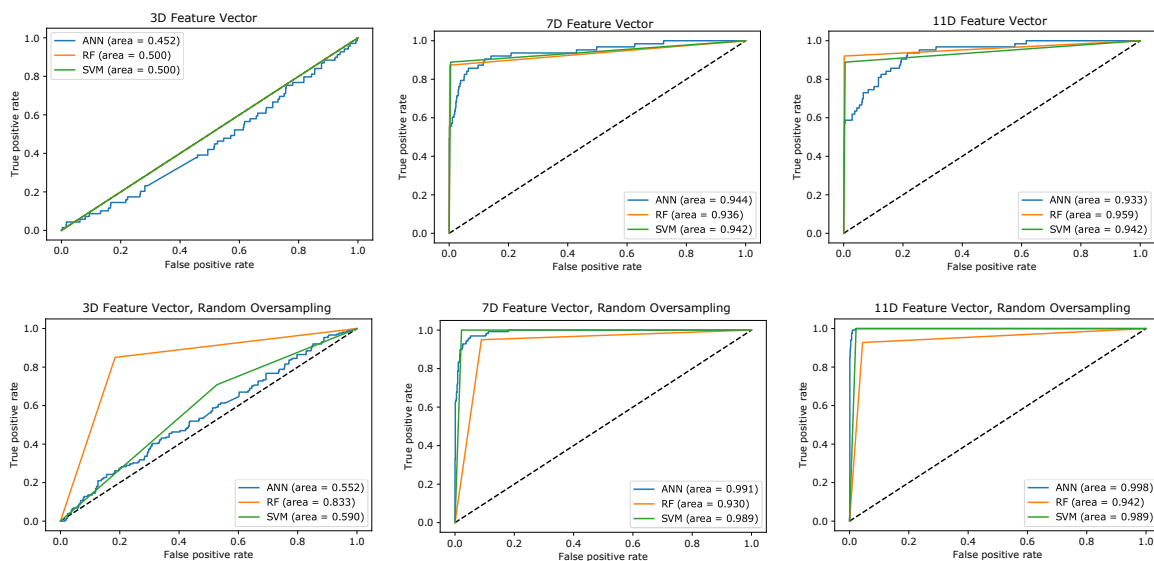
Prediction performance was hardly influenced by increasing the complexity of the ANN model. Both models achieved a *precision* of 0.98 and *recall* of 1.0. In both cases, training- and validation error and thus the quality of the model did not improve after already about 50 epochs.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that in this case the position of the 'Modulation' slider is by far the dominant feature. This is in accordance with the previously observed fact that in cranial cases the prescribed dose coverage could be always almost fulfilled, with deviations only in cases with very low 'Modulation'.

### PTV 1% dose w.r.t. weighting slider

The second quality criterion  $C$  concerning the PTV object analyzed in this section is the maximum dose received by the lowest volume-percent of the PTV. This is basically equivalent to a more robust version of the minimum dose which is not sensitive to a very low number of extremely cold outlier voxels. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value.


The left plot in figure 4.11 shows the histogram of the PTV-1%-Dose values achieved in all optimization results for all available 44 cranial patient geometries. The value range achievable



**Figure 4.10:** ROC curves for the Cranial PTV coverage classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

with changes in the 'Weighting' slider only can be seen in the center, and a threshold postulating that the achievable values should differ by at least 1 Gy is chosen, which leads to the two classes shown on the right. Again, the distribution of the classes suggests application of an oversampling technique.

Figure 4.12 shows the fitting results for the SVM, ANN and RF classifiers for both the original

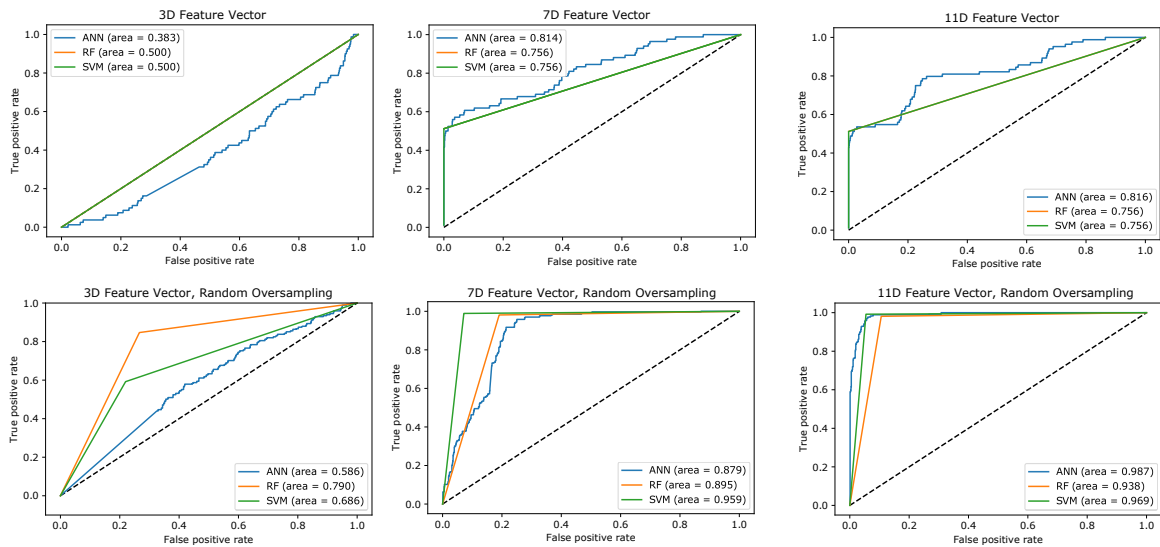


**Figure 4.11:** Histograms of PTV 1% Dose values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 1Gy range threshold (right)

dataset and the balanced data after randomly duplicating samples from the underrepresented class using the random-oversampling technique. Again, without oversampling, all models showed poor performance, especially in the case where the lowest number of features are used. By increasing the number of features the prediction quality could be improved, but still not to an acceptable degree. All three classifiers performed similarly and the increment from seven to eleven features did not yield an increase in performance.

Random-oversampling had an enormous impact on model performances. With 7 features AUC values up to 0.959 could be achieved, and with 11 features the ANN model achieved an excellent AUC value of 0.987.

A comparison of the two different ANN structures, one with 5 hidden units and one with 64



**Figure 4.12:** ROC curves for the Cranial 1% PTV Dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.


hidden units showed again that such a massive increase of model complexity had no impact on the prediction quality. Both models achieved a *precision* of 0.95 and *recall* of 0.99, differences showed up only in the 3rd decimal. The ANN training process also revealed that training- and validation error and thus the quality of the model did not improve after about 100 epochs. Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that the PTV volume was the predominant feature, followed by the section-slope values of the OAR-OVH and PTV OAR distance.

### Gradient index (GI) w.r.t. normal-tissue-sparing slider

As an example for a criterion  $C$  which is largely influenced by the 'Normal Tissue Sparing' parameter, the Gradient Index (GI) is analyzed in this section. Subject of the analysis is the degree of influence of the 'Modulation' parameter on this value, which gives information on how quickly the irradiation dose decreases outside the PTV.

The left plot in figure 4.13 shows the histogram of the GI values achieved in all optimization results for all available 44 cranial patient geometries (PTV prescription dose 15 Gy). The value range achievable with changes in the 'Normal Tissue Sparing' slider only can be seen in the center, and, in order to identify the instances with a nominal effect, a threshold postulating that the achievable values should differ by at least 1 is chosen, which leads to the two classes shown on the right. Again, the two classes are extremely unbalanced, and after random-oversampling, the classes are balanced.

Figure 4.14 shows the fitting results for the SVM, ANN and RF classifiers for both the original

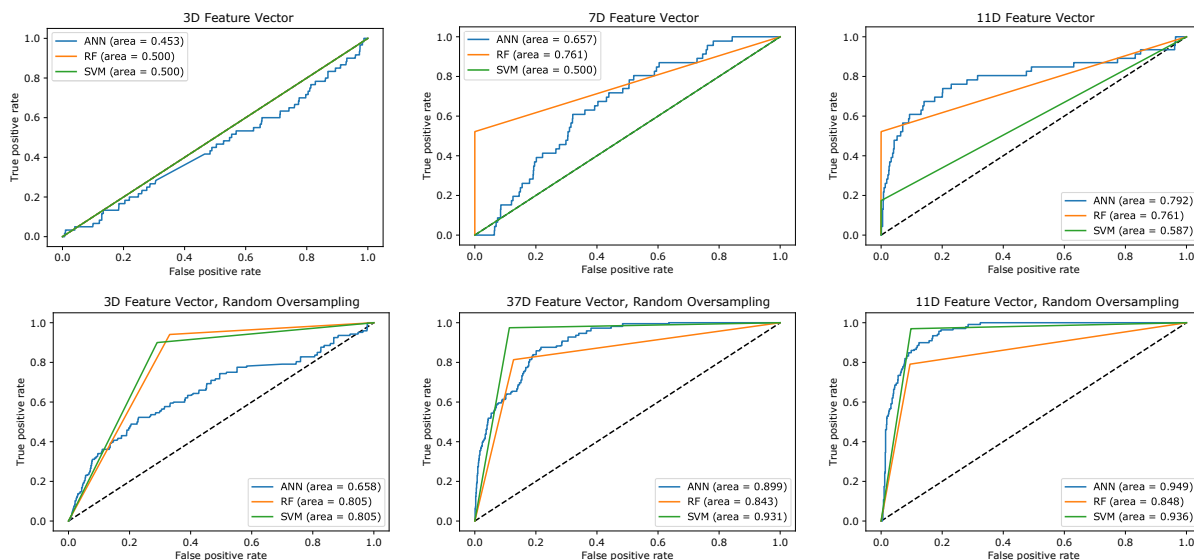


**Figure 4.13:** Histograms of GI values in the Cranial cases (left), the GI range achievable with the 'Normal Tissue Sparing' slider (center), and the two sample classes resulting by introduction of a GI range threshold of 1 (right)

dataset and the balanced data after randomly duplicating samples from the underrepresented class using the random-oversampling technique. As expected, without oversampling all models showed poor performance, in the case where the lowest number of features used prediction quality was not better than a coin-toss. With increasing number of features the prediction quality improved slightly. This time, the RF model and the ANN model had comparable AUC values, the SVM model could not handle the unbalanced data even with 11 features.

Random-Oversampling lead to a big improvement of model performance. With 7 features AUC values up to 0.93 could be achieved, and with 11 features the ANN model achieved an very good AUC value of 0.95.

A comparison of the two ANN structures showed that such a massive increase of model



**Figure 4.14:** ROC curves for the Cranial GI classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

complexity is able to increase the recall value from 0.96 to 1.0, however, both models had a precision of 0.94 and the 'perfect' recall of the more complex model probably is due to overfitting and does not necessarily mean better generalization of the model. The ANN training


process showed that training- and validation error and thus the quality of the model did improve massively in the first 100 epochs. After that, accuracy only improved very slowly with increasing number of epochs, but even after 200 epochs slight changes could be observed. Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that the PTV volume was the predominant feature.

**Monitor units (MU) w.r.t. modulation slider**

As an example for a criterion  $C$  which is largely influenced by the 'Modulation' parameter, the amount of dose monitor units (MU) needed to deliver each resulting treatment plan is analyzed in this section. Subject of the analysis is the degree of influence of the 'Modulation' parameter on this value.

The left plot in figure 4.15 shows the histogram of the MU values achieved in all optimization results for all available 44 cranial patient geometries (PTV prescription dose 15 Gy). The impressive value range achievable with changes in the 'Modulation' slider only can be seen in the center, and, in order to identify the most extreme occurrences, a threshold postulating that the achievable values should differ by at least 1000 is chosen, which leads to the two classes shown on the right. Again, the two classes are extremely unbalanced, and after random-oversampling, the classes are balanced.

Figure 4.16 shows the fitting results for the SVM, ANN and RF classifiers for both the original

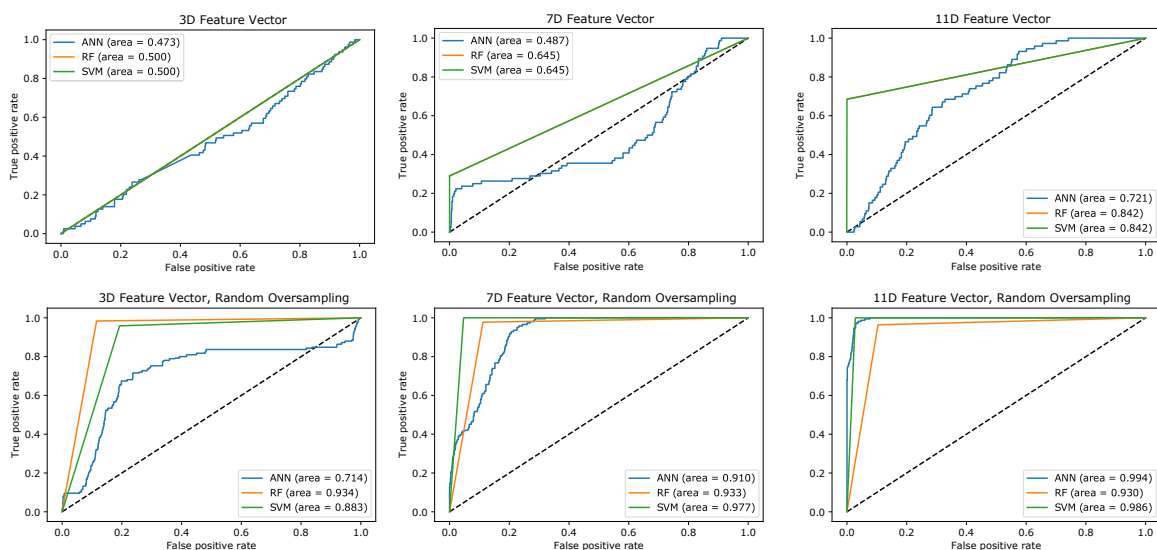


**Figure 4.15:** Histograms of MU values in the Cranial cases (left), the MU range achievable with the 'Modulation' slider (center), and the two sample classes resulting by introduction of a 1000 MU range threshold (right)

dataset and the balanced data application of the random-oversampling technique. As expected, without oversampling all models showed poor performance, in the case where the lowest number of features are used prediction quality was not better than a random guess. With increasing number of features the prediction quality improved, but still not to an acceptable degree. The simpler models (SVM, RF) performed better than the ANN, most likely the amount of data is not sufficient for proper ANN training without oversampling.

Random-oversampling had an enormous impact on model performances. With 7 features AUC values up to 0.97 could be achieved, and with 11 features the ANN model achieved an excellent AUC value of 0.99.

A comparison of the two available ANN structures, one with 5 hidden units and one with 64 hidden units showed that such a massive increase of model complexity can increase the recall value from 0.96 to 1.0, however, also in this case this might be an overfitting issue and does not necessarily mean better generalization of the model. The ANN training process showed that training- and validation error and thus the quality of the model did not improve after about



**Figure 4.16:** ROC curves for the Cranial MU classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

150 epochs.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that the OAR volume was the predominant feature, followed by the PTV concavity.

#### 4.2.4 Analysis of spine cases with spinal canal as prioritized OAR

In the second part of this machine-learning-based analysis the data for spinal indications will be analyzed. By choosing the same criteria  $C_i$  similarities and differences to the cranial data will be discernible.


##### OAR constraint point dose w.r.t. weighting slider

The first quality criterion  $C$  concerning the OAR object 'Spinal Canal' analyzed in this section is the  $D_{0.25cm^3}$  of the OAR object. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value. In terms of the problem definition presented above, the task at hand is to predict in a certain situation whether trying out more positions of the 'Weighting' slider are likely to have a big impact on this value.

The left plot in figure 4.17 shows the histogram of the  $D_{0.25cm^3}$  values of the OAR object 'Spinal Canal' achieved in all optimization results for all available 18 cranial patient geometries (optimization target for this criterion is not to exceed 10 Gy). The achievable value range shown in the central figure indicates that in most cases a range of 6 Gy is possible and the structure of the data does not directly suggest a certain threshold. In order to stay consistent with the methodology applied in this section, the threshold of 6 Gy is chosen, e.g. the model tries to find the samples where the achieved dose range is even bigger than the peak of the histogram (right figure)

## 4.2 Identification of relevant sections of the parameter space by prediction of achievable value ranges

Figure 4.18 shows the fitting results for the SVM, ANN and RF classifiers for both the original

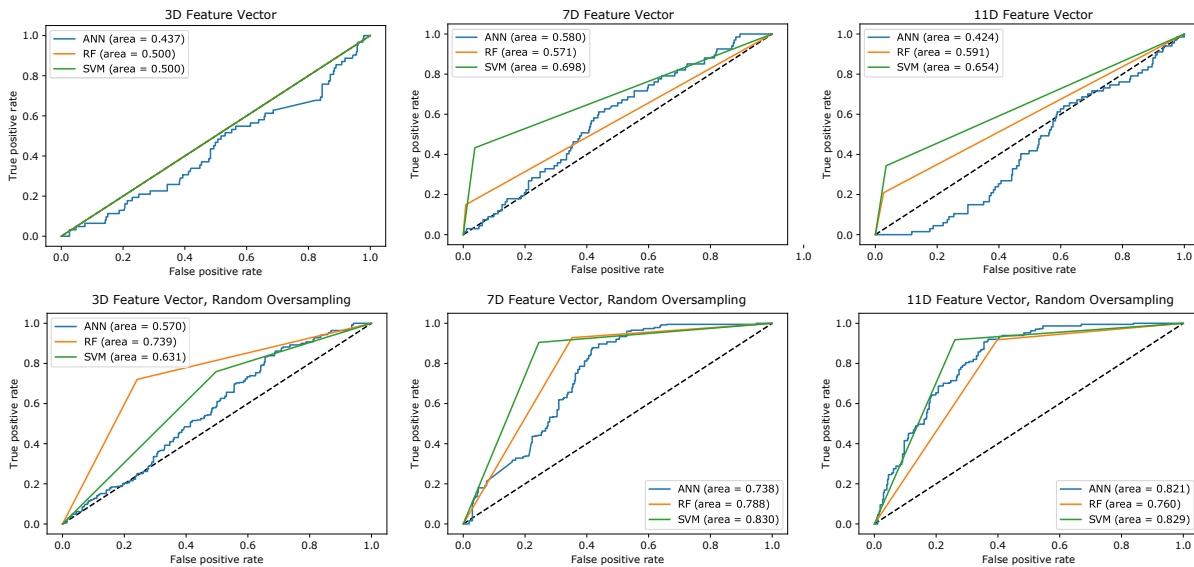


**Figure 4.17:** Histograms of OAR  $D_{0.25cm^3}$  values in the Spine cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 6 Gy OAR Constraint Dose range threshold (right)

dataset and the balanced data after randomly duplicating samples from the underrepresented class using the Random-Oversampling technique. Without oversampling all models showed poor performance, even with the maximum number of 11 features.

Random-Oversampling improved model performances. With 7 features AUC values up to 0.83 could be achieved, however, adding more features did not lead to noticeable improvements.

Comparison of different ANN structures showed that the 64 unit model achieved a *precision* of



**Figure 4.18:** ROC curves for the OAR Constraint point dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

0.82 and a *recall* of 0.98. The 5 unit model achieved a *precision* of 0.72 and a *recall* of 0.98. The ANN training process also showed that training- and validation error and thus the quality of the model still continued to improve slightly even after 200 training epochs.


Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that in this case the OVH penumbra of the OAR was the dominant feature.

### PTV rel. coverage w.r.t. weighting slider

The first quality criterion  $C$  concerning the spinal PTV object is the relative volume coverage of the PTV with the prescribed dose of 16 Gy for the spine cases. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value. In terms of the problem definition presented above, the task at hand is to predict in a certain situation whether trying out more positions of the 'Weighting' slider are likely to have a big impact on this value.

The left plot in figure 4.19 shows the histogram of the relative PTV volume coverage values achieved in all optimization results for all available 18 spinal patient geometries (PTV prescription dose 16 Gy for 95% of the PTV volume). The value range achievable with changes in the 'Weighting' slider only can be seen in the center, and a plausible threshold postulating that the achievable values should differ by at least 0.15 (e.g. 15% of the PTV volume) is chosen, which leads to the two classes shown on the right. Even though a slightly higher threshold might be used in this case, 15% was chosen in order to find out if easy-to-memorize rule-of-thumb threshold values might be established for future research. Compared to the cranial cases, a wider range of values is achieved. Still, the two classes are unbalanced, a situation which has to be addressed by oversampling techniques which were briefly introduced in section 2.6 if a reliable training of an AI model is desired. After random oversampling, the classes are balanced.

Figure 4.20 shows the fitting results for the SVM, ANN and RF classifiers for both the original



**Figure 4.19:** Histograms of relative PTV volume dose coverage values in the Spine cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 5% PTV volume range threshold (right)

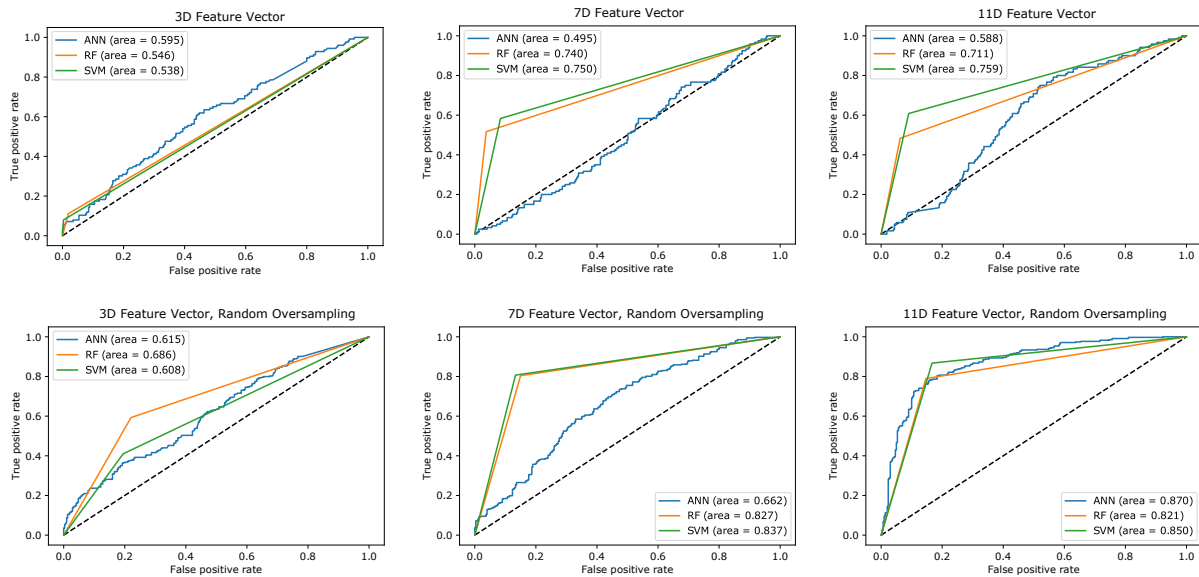
dataset and the balanced data after randomly duplicating samples from the underrepresented class using the Random-Oversampling technique. As expected, and as seen in the cranial data, without oversampling all models showed poor performance, and increasing the number of features did not help much, especially in the ANN model

Random-Oversampling had an impact on model performances, but only with the maximum number of 11 features an AUC of 0.87 could be achieved with the ANN model.

Again, ANN structure comparison indicated that such a massive increase of model complexity is not reflected by an increase in performance. The complex model achieved a *precision* of 0.87 and *recall* of 0.93 while the simple model the *precision* was 0.85 and the *recall* value was 0.83. The ANN training process also showed that training- and validation error and thus the quality of the model still improved after already about 200 epochs.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that in this case the object volumes and the PTV concavity had a high significance, followed by the position of the 'Modulation' slider.





**Figure 4.20:** ROC curves for the spine PTV coverage classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

### PTV 1% dose w.r.t. weighting slider


The second quality criterion  $C$  concerning the PTV object analyzed in this section is the maximum dose received by the lowest volume-percent of the PTV. This is basically equivalent to a more robust version of the minimum dose which is not sensitive to a very low number of extremely cold outlier voxels. Subject of the analysis is the degree of influence of the 'Weighting' parameter on this value. In other words, the task at hand is to predict in a certain situation whether trying out more positions of the 'Weighting' slider is likely to have a big impact on this value.

The left plot in figure 4.21 shows the histogram of the PTV-1%-Dose values achieved in all optimization results for all available 18 spinal patient geometries. The value range achievable with changes in the 'Weighting' slider only, and the huge effect it has on this criterion for this indication can be seen in the center. In this case the AI model can be applied in order to rule out the least effective positions. Therefore an threshold postulating that the achievable values should differ by less than 4 Gy is chosen, which leads to the two classes shown on the right.

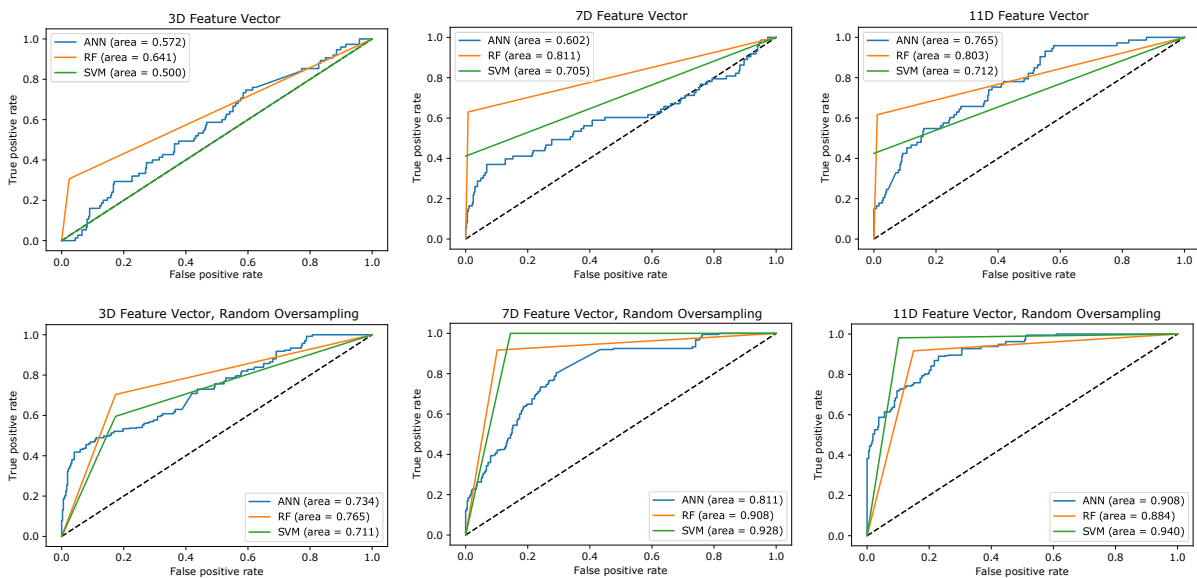
Figure 4.22 shows the fitting results for the SVM, ANN and RF classifiers for both the original dataset and the balanced data after using the random-oversampling technique. Without oversampling all models showed poor performance and increasing the number of features improved the prediction quality only slightly.

Random-Oversampling had noticeable impact on model performances. With 7 features AUC values up to 0.92 could be achieved, and with 11 features the SVM model achieved an AUC value of 0.94.

A comparison of different ANN structures, one with 5 hidden units and one with 64 hidden units showed the complex model achieved a *precision* of 0.95 and *recall* of 0.98 while the simple model the *precision* was 0.89 and the recall value was 0.94. The ANN training process also showed that training- and validation error and thus the quality of the model did not decrease



**Figure 4.21:** Histograms of PTV 1% Dose values in the Spinal cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 4 Gy range threshold (right)



**Figure 4.22:** ROC curves for the Spinal 1% PTV Dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

after about 150 epochs.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that the PTV volume was the predominant feature, followed by the PTV concavity.


### Gradient index (GI) w.r.t. normal-tissue-sparing slider

As an example for a criterion  $C$  which is largely influenced by the 'Normal Tissue Sparing' parameter, the gradient index (GI) is analyzed in this section. Subject of the analysis is the degree of influence of the 'Modulation' parameter on this value, which gives information on how quickly the irradiation dose decreases outside the PTV.

The left plot in figure 4.23 shows the histogram of the GI values achieved in all optimization results for all available 18 spinal patient geometries (PTV Prescription Dose 16 Gy). The value range achievable with changes in the 'Normal Tissue Sparing' slider only can be seen in the center, and, in order to identify the instances with a nominal effect, a threshold postulating

that the achievable values should differ by at least 1 is chosen, which leads to the two classes shown on the right. Again, the data for this problem turns out to be unbalanced.

Figure 4.24 shows the fitting results for the SVM, ANN and RF classifiers for both the original

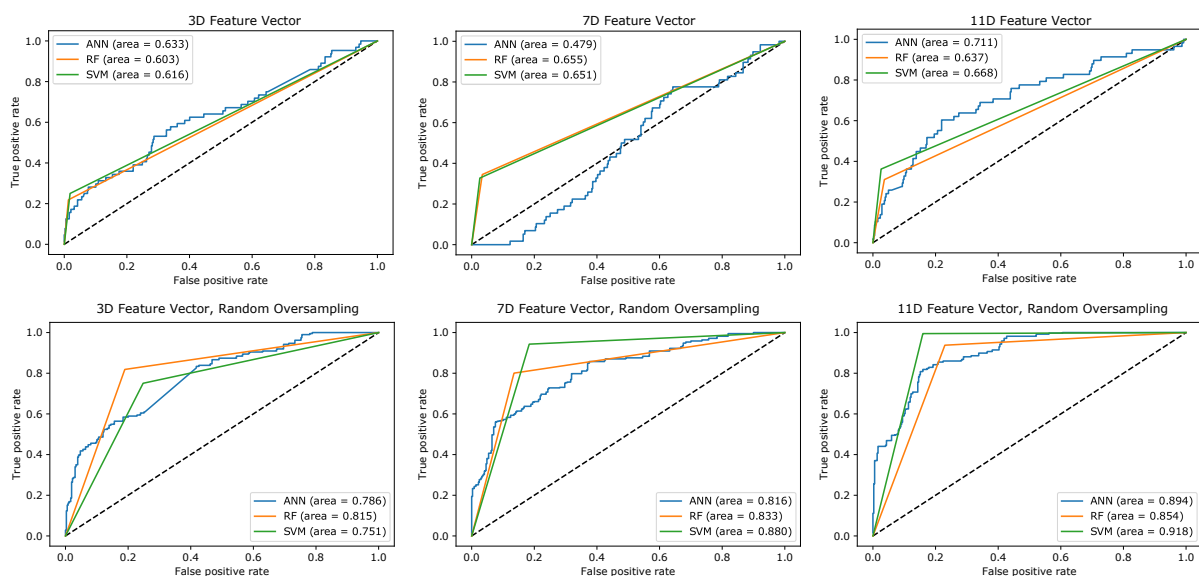


**Figure 4.23:** Histograms of GI values in the spine cases (left), the GI range achievable with the 'Normal Tissue Sparing' slider (center), and the two sample classes resulting by introduction of a GI range threshold of 1 (right)

dataset and the balanced data after randomly duplicating samples from the underrepresented class using the random-oversampling technique. For this criterion  $C$  and the spinal data, without oversampling all models showed poor performance, no matter which feature set was used.

As in the previous cases, random-oversampling lead to a big improvement of model performance while the effect of increasing the number of features was not as distinctive as in other analyses presented in this section.

The comparison of the 5-unit and the 64-unit ANN showed that such a massive increase



**Figure 4.24:** ROC curves for the Spine GI classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature vector. Upper row is without, lower row with random-oversampling to overcome class-imbalance.

of model complexity increases the *precision* value from 0.83 to 0.91. The 5-unit-model had a 'perfect' *recall* score but this is at the cost of a way worse *precision*. Even after 200 epochs slight increases in accuracy could be observed, however this is most likely a sign of overfitting.


Analysis of the information on the relative impact of specific features on the decision process

in the RF classifier showed that the slope of the first section of the OAR OVH was the most effective feature, followed by the PTV concavity and PTV volume, which makes sense since all of them are likely to influence feasible dose gradient outside the PTV.

### Monitor units (MU) w.r.t. modulation slider

As an example for a criterion  $C$  which is largely influenced by the 'Modulation' parameter, the amount of dose monitor units (MU) needed to deliver each resulting treatment plan is analyzed in this section. Again, the focus is placed on the degree of influence of the 'Modulation' parameter on this value.

The left plot in figure 4.25 shows the histogram of the MU values achieved in all optimization



**Figure 4.25:** Histograms of MU values in the Spine cases (left), the MU range achievable with the 'Modulation' slider (center), and the two sample classes resulting by introduction of a 4500 MU range threshold (right)

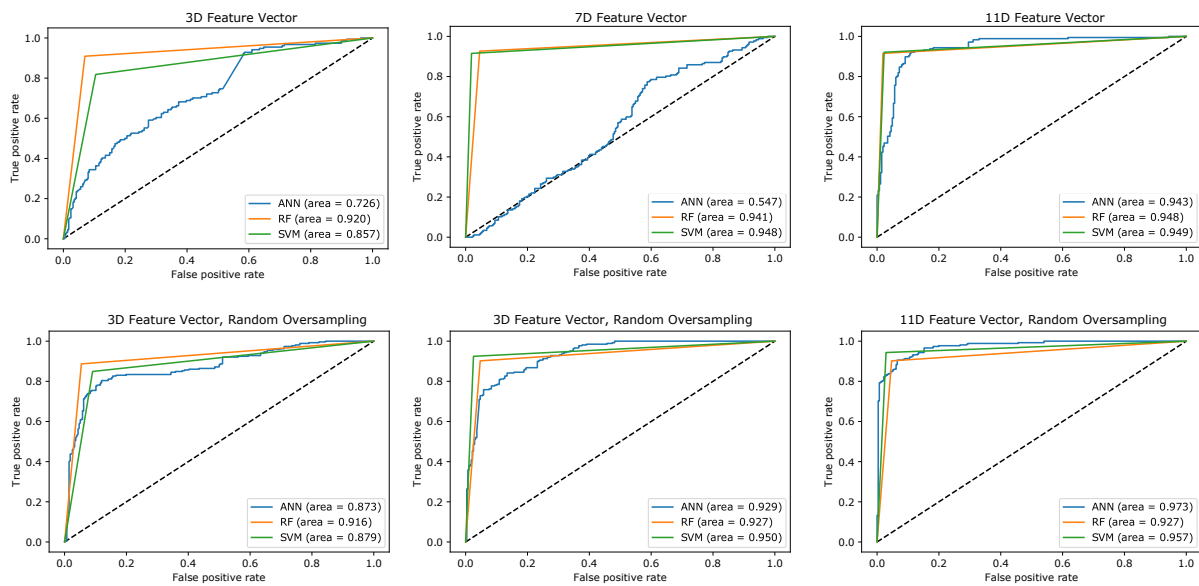
results for all available 18 spinal patient geometries (PTV prescription dose 16 Gy). The impressive value range achievable with changes in the 'Modulation' slider only can be seen in the center, and, in order to identify the most extreme occurrences, a threshold postulating that the achievable values should differ by at least 4500 MU is chosen, which leads to the two classes shown on the right. With such a threshold, the data is not as unbalanced as in other cases, making it an interesting case how the models deal with more equally distributed data. Figure 4.26 shows the fitting results for the SVM, ANN and RF classifiers for both the original dataset and the balanced data after randomly duplicating samples from the underrepresented class using the Random-Oversampling technique. Interestingly, performance of the SVM and RF classifier did not vary much with feature number or the use of random-oversampling or not. Only the ANN classifier achieved its highest AUC value of 0.973 with the highest number of features and oversampled data.

The comparison of ANN structures with 5 hidden units and one with 64 hidden units showed that in this case such a massive increase of model complexity does not increase the *recall* value of 0.96 and leads to a slightly increase from 0.95 to 0.97 in the *precision* value, only. The ANN training process showed that training- and validation error and thus the quality of the model did not improve after about 100 epochs already.

Analysis of the information on the relative impact of specific features on the decision process in the RF classifier showed that the PTV concavity was the predominant feature.

## 4.2.5 Results and discussion

In this chapter a machine-learning based method for evaluation of the slider-defined solution space of the cranial and spinal VMAT optimization problem was presented. The training and



**Figure 4.26:** ROC curves for the Spine MU classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.

test data for the machine learning process was taken from optimization results obtained from 44 cranial and 18 spinal patient geometries. For each geometry 343 VMAT optimized dosed distributions were calculated, resulting in a total of 21,266. For each dose distribution, the source geometry was described in a 1D, 5D and 9D vector of characterizing numerical values. The dose distributions and organ DVH data was not saved as a whole, but condensed down to a list of dosimetric criterion values  $C_i$ .

The combined data of slider-positions, patient geometry characteristics and dosimetric results was used to assemble the *feature matrix*  $X$  and *target vector*  $y$  needed for AI model creation. Despite the high overall number of optimization results, the relative low amount of 62 available patient geometries is not sufficient to train for direct prediction of dosimetric values based on slider position and geometrical data. Therefore a more general approach was chosen, not directly predicting numerical values, but rather giving a suggestion on whether, in a certain state of slider positions, a modification of the setting of a specific slider is likely to have a noticeable effect on a specific criterion  $C_i$ , or not.

For each treatment location (cranial and spine) three criteria were analyzed with respect to the 'Weighting' parameter, plus two additional criteria, one with respect to the 'Normal Tissue Sparing' parameter, the other one with respect to the 'Modulation' slider. Based on the histogram of all values achieved for each specific  $C_i$  in the data and the corresponding value range, a threshold discerning 'effective' and 'not effective' was chosen individually, resulting in a binary classification problem. While in some cases, e.g. the relative PTV dose coverage in the cranial geometries, the structure of the data naturally suggested a specific threshold value, in other cases the choice was not so straight-forward. Nevertheless, every time a threshold value suitable for showing the feasibility of the method could be found, and, since the final decision in the MCO treatment planning process is left to the user anyway, the choice of the threshold might be made according to the user's preference, too. The experiments in this section demonstrated that the threshold-based efficiency-prediction approach can be used for:

- Finding regions in the parameter space with unusually high impact on the  $C_i$  in terms of outlier-detection
- Ruling out regions in the parameter space, where a certain  $C_i$  will remain invariant
- Finding parameter regions deemed interesting because of the user's preference (e.g. the MU value varying by more than 1000 MU)

Adding the positions of the two slider parameters not under current analysis to the geometry vectors created feature matrices with 3, 7 and 11 columns, corresponding to a 3-, 7-, or 11D feature vector describing each sample. Depending on the  $C_i$  and choice of threshold, the number of samples for each class varied and was mostly highly imbalanced. In the cranial data typically only a few hundred of the 2156 samples were labeled as being interesting, in the spine cases, while the overall number was reduced to 882, the ratios were slightly more balanced. A standard regularization method was applied on the resulting data in order to eliminate effects of the individual features existing in different numeric scales and orders of magnitude.

Three types of machine learning model (SVM, RF and ANN) were trained and tested on the data with varying partitioning into training- and test set, mostly a 50-50 separation with a number of iterations for cross-validation was chosen. Two ANN models with varying number of hidden units (5 vs. 64) were evaluated. Prediction quality was measured with the area-under-curve (AUC) of the corresponding receiver-operating-characteristic curves, and precision- and recall values.

With the available amount of data, model training took typically less than a minute, with the ANN training with 200 epochs being the most time-consuming. Here training time could exceed 100s. Without further refinement of the training data, training in all models was greatly affected by the high grade of class-imbalance, and these problems generally could not be mitigated by choosing a higher number of features. The AUC values for all three types of classifier and the tree feature vector versions without oversampling can be seen in table 4.3 for the cranial cases and in table 4.4 for the spine cases.

**Table 4.3:** AUC values without Random Oversampling, Cranial

Criterion	AUC SVM 3D	AUC RF 3D	AUC ANN 3D	AUC SVM 7D	AUC RF 7D	AUC ANN 7D	AUC SVM 11D	AUC RF 11D	AUC ANN 11D
Brainstem $D_{10\%}$	0.500	0.500	0.540	0.758	0.730	0.773	0.758	0.730	0.866
PTV $V_{15Gy}$ rel.	0.500	0.500	0.452	0.942	0.936	0.944	0.942	0.959	0.933
PTV $D_{1\%}$	0.500	0.500	0.383	0.756	0.756	0.814	0.756	0.756	0.816
GI	0.500	0.500	0.453	0.500	0.761	0.657	0.587	0.761	0.792
MU	0.500	0.500	0.473	0.645	0.645	0.487	0.842	0.842	0.721

**Table 4.4:** AUC values without Random Oversampling, Spine

Criterion	AUC SVM 3D	AUC RF 3D	AUC ANN 3D	AUC SVM 7D	AUC RF 7D	AUC ANN 7D	AUC SVM 11D	AUC RF 11D	AUC ANN 11D
Sp. Canal $D_{0.25ccm}$	0.500	0.500	0.437	0.698	0.571	0.580	0.654	0.591	0.424
PTV $V_{16Gy}$ rel.	0.538	0.546	0.595	0.750	0.740	0.495	0.759	0.711	0.588
PTV $D_{1\%}$	0.500	0.641	0.572	0.705	0.811	0.602	0.712	0.803	0.765
GI	0.616	0.603	0.633	0.651	0.655	0.479	0.668	0.637	0.711
MU	0.857	0.920	0.726	0.948	0.941	0.547	0.949	0.948	0.943

Following a widely used practice in machine learning applications, the Random-Oversampling technique was applied to remove the class imbalances, which had significant impact on the

prediction quality. Now excellent *precision*- and *recall* values of 0.85 and higher could be achieved, even with the SVM model, which shows the effectiveness of the chosen approach. The improvement in prediction quality can also be seen in the AUC values for the cranial cases (table 4.5) and the spine cases (table 4.6).

**Table 4.5:** AUC values with Random Oversampling, Cranial

Criterion	AUC SVM 3D	AUC RF 3D	AUC ANN 3D	AUC SVM 7D	AUC RF 7D	AUC ANN 7D	AUC SVM 11D	AUC RF 11D	AUC ANN 11D
Brainstem $D_{10\%}$	0.686	0.713	0.678	0.952	0.904	0.950	0.969	0.924	0.982
PTV $V_{15Gy}$ rel.	0.590	0.833	0.552	0.989	0.930	0.991	0.989	0.942	0.998
PTV $D_{1\%}$	0.686	0.790	0.586	0.959	0.895	0.879	0.969	0.938	0.987
GI	0.805	0.805	0.658	0.931	0.843	0.899	0.936	0.848	0.949
MU	0.883	0.934	0.714	0.977	0.933	0.910	0.986	0.930	0.994

**Table 4.6:** AUC values with Random Oversampling, Spine

Criterion	AUC SVM 3D	AUC RF 3D	AUC ANN 3D	AUC SVM 7D	AUC RF 7D	AUC ANN 7D	AUC SVM 11D	AUC RF 11D	AUC ANN 11D
Sp. Canal $D_{0.25ccm}$	0.631	0.739	0.631	0.830	0.788	0.738	0.829	0.760	0.821
PTV $V_{16Gy}$ rel.	0.608	0.686	0.615	0.837	0.827	0.662	0.850	0.821	0.870
PTV $D_{1\%}$	0.711	0.765	0.734	0.928	0.908	0.811	0.940	0.884	0.908
GI	0.751	0.815	0.786	0.880	0.833	0.816	0.918	0.854	0.894
MU	0.879	0.916	0.873	0.950	0.927	0.929	0.957	0.927	0.973

Still, further analysis of the structure of the labeled data has to be done in order to evaluate the usefulness of the predictions, and generally more data is needed to refine this method. Some of the resulting *precision* and *recall* values are so good, that great care has to be taken in order to find out whether the models also generalize well on unknown data, or whether a big part of the seemingly good prediction quality has to be attributed to overfitting artifacts due to the limited amount and structure of the training data. An overview of the achieved precision and recall values can be seen in table 4.7 (cranial) and 4.8 (spine). Only the acquisition of vastly more data can clarify this issue, and, since the overall workflow does not contain any elements negatively affecting scalability towards larger amounts of data, except higher training time and the one-time overhead of data pre-processing, this is a feasible task. Theoretically, the models and software libraries used for this study can handle training data containing millions of samples and can be easily extended towards more powerful learning algorithms such as deep learning [66]. Additionally, with a high amount of data, the regression approach quickly discussed, and discarded because of lack of sufficient data, can be resumed, providing even more direct support to the user.

**Table 4.7:** Precision and recall values with 11D feature vector and Random Oversampling, cranial cases. Comparison of ANNs with 5 and 64 hidden units (HU), after 200 training epochs

Criterion	Precision	Precision	Recall	Recall
	5 HU	64 HU	5 HU	64 HU
Brainstem $D_{10\%}$	0.97	0.96	0.96	1.00
PTV $V_{15Gy}$ rel.	0.98	0.98	1.00	1.00
PTV $D_{1\%}$	0.95	0.95	0.99	0.99
GI	0.94	0.94	0.96	1.00
MU	0.99	0.99	0.96	1.00

**Table 4.8:** Precision and recall values with 11D feature vector and Random Oversampling, spine cases. Comparison of ANNs with 5 and 64 hidden units (HU), after 200 training epochs

Criterion	Precision	Precision	Recall	Recall
	5 HU	64 HU	5 HU	64 HU
Sp. Canal $D_{0.25ccm}$	0.72	0.82	0.98	0.98
PTV $V_{16Gy}$ rel.	0.85	0.87	0.83	0.93
PTV $D_{1\%}$	0.89	0.95	0.94	0.98
GI	0.83	0.86	0.91	1.00
MU	0.95	0.97	0.96	0.96



**Chapter 5**  
**Outlook & Conclusion**

## 5.1 Outlook: Combination of ML models for rapid exploration of the solution space

With the main parts of the research conducted during in this thesis being covered in the previous two chapters, this section shall provide an outlook on how all these results can be combined to create a streamlined treatment planning workflow with a high degree of automatization. The holy grail of radiotherapy treatment planning, of course, would be a completely automated planning process, where without any user interaction a deliverable optimal treatment sequence is generated individually for each patient and treatment site [83]. A fully automated MCO planning workflow for IMRT treatments has been implemented by *Voet et. al* [82]. However, due to the use of a fixed predefined lexicographic order imposed on the optimization objectives, it is only efficient within a narrow range of treatment indications and treatment sites. The nature of the multi-criteria optimization problem makes automatization of a general concept of optimality almost impossible, since many candidates for 'optimality' exist, and the final choice is depending on the preferences of a human operator.

In the literature many strategies aiming at finding a reasonable range of candidates for the optimal plan and facilitating this decision process for the user can be found. Pareto surface navigation has been done for IMRT[46] and VMAT[19], and commercially available solutions have been evaluated as being beneficial to the planning process, enabling even less experienced planners to produce excellent treatment plans [47] and achieve better sparing of healthy tissue [40].

With the recent rise of machine learning applications in virtually all areas of technology, radiotherapy treatment planning has been no exception, and several attempts have been made to leverage this powerful technology for saving time [2] [64], improving outcomes,[54] [73] [81] and facilitating quality assurance (QA) [78].

While promising progress has been made, the application of machine learning technology in radiotherapy is still in its infancy[27]. The techniques presented in this work have the advantage of not relying on any hard mathematical assumptions regarding the optimizer or planning constraints. In future development, they might be combined to implement the promising workflow presented below.

### 5.1.1 Outline and structure of algorithm

As a first enhancement step, the prediction method depicted in figure 4.2 based on training data assembled as shown in figure 4.3, can be repeated for every possible choice of fixed slider and prediction slider, for each  $C_i$  a specific model can be trained. The final goal is to be able to predict, at any given slider position combination and for any quality criterion  $C_i$ , which slider is the best choice for adjusting this specific  $C_i$ . Of course, this means that a total of three ANN classifiers have to be trained, but no additional data has to be acquired since all the necessary feature matrices and target vectors can be constructed from the database described in figure 4.1.

To achieve this goal, the most obvious classifier choice is the ANN predictor. One reason is its ease of scalability and flexibility for handling huge data sets, the other being the fact that in contrast to e.g. a SVM classifier the output is not a binary value, but merely the estimated probability, that a specific sample has a specific label. This property allows for a ranking of

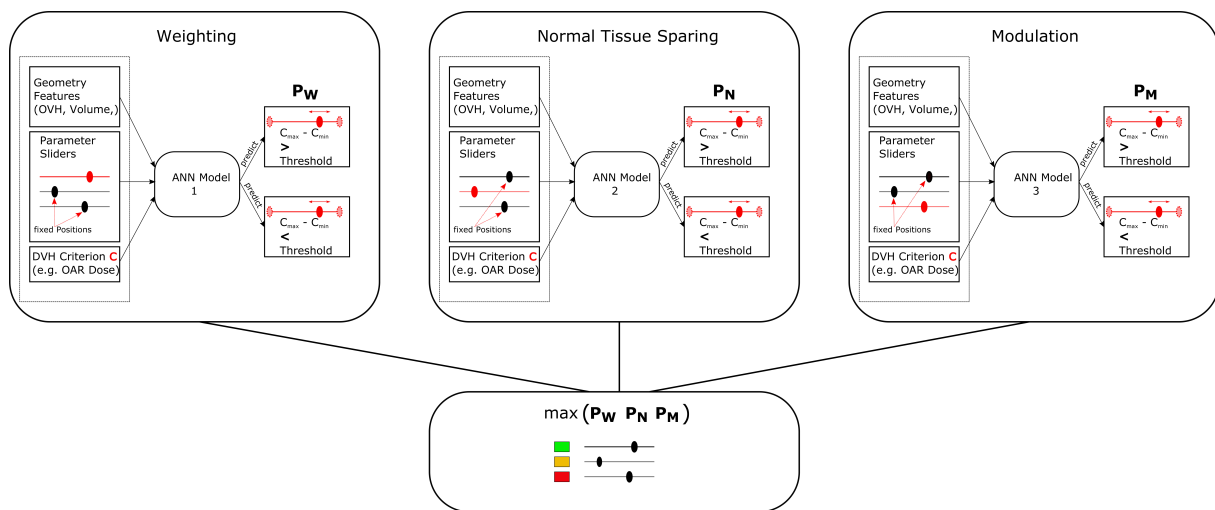


Figure 5.1: Complete prediction algorithm including all slider controls

predictions, in our case this means that, if more than one slider is a candidate for modification of a specific  $C_i$ , the slider with the highest probability can be picked. In terms of a graphical user interface in the treatment planning software, this might be achieved by a drop-down menu for selection of the  $C_i$ , the recommendation of the predicted slider control can be presented via a traffic-light-like color-coding scheme (a schematic overview of this algorithm is given in figure 5.1).

So far, only the relative value range of the  $C_i$  was considered for predicting slider efficiency. In the actual planning process however, absolute values might be more important. In this case, a different target vector could be used for training and prediction of which slider, when modified, given the current slider positions, is most likely to cause a specific constraint for a  $C_i$  to be fulfilled.

### 5.1.2 Additional enhancements

#### Use of big data/cloud-based data management

It is obvious that the database in figure 4.1 by no means has to end up being a static structure. Experience in the application of machine learning to various areas has shown that more training data is the biggest influence on training better models, and in most cases acquiring more data beats the application of more sophisticated algorithms in terms of achieving higher prediction quality [33] [90].

The models presented in this thesis could thus also benefit greatly from an ever increasing pool of training data. Every new case in a treatment center can be potentially immediately added to the training database, and an automated background process can perform the dose optimizations for all parameter combinations and re-train the models.

Such an incremental database and training process could even be extended across institution borders. In the US, pilot studies for national patient data registries for spine and cranial indications have been started [53], [70]. The benefit of data-based research on the outcome of cranial SRS procedures has already been shown [69], it is an obvious idea that such a vast amount of data might be used for machine-learning purposes, too. Of course, as pointed out by

*el Naqa et al.*[25], the application of machine learning techniques in radiation oncology is still in a very early developmental stage, and also a lot of bureaucratic and infrastructure-related problems still have to be dealt with.

### **Learning of user preferences**

It has been mentioned a couple of times in this thesis that the nature of an multi-criteria optimization problem is its abundance of potentially equally 'optimal' solutions. Sometimes some hierarchy of objectives, or some other means of defining a global optimum can be established. However, in many cases the final choice is still made by some human operator. In this context, the slider positions resulting in the finally chosen treatment plan can be saved for each patient case. A neural network can be trained using this position data and the geometry vector  $G$  supplemented with additional data concerning the selected treatment regimen such as prescribed dose, number of arc elements etc. for prediction of likely slider positions in a new patient case. Once a reasonable amount of data has been accrued, upload of a new patient case might trigger automatic pre-calculation of the most likely slider positions predicted by the ANN model. This procedure would provide a good starting point for the human decision maker, and with a growing number of training data the model is likely to learn the preferences of a specific user.

## 5.2 Summary & conclusion

So far, a number of topics has been covered in this thesis: After some general introduction into the theoretical aspects of radiotherapy treatment planning, the basic principles of multi-criteria optimization and machine learning, the Brainlab Elements VMAT Treatment Planning System was introduced. The concept of VMAT optimization with a low number of composite objectives grouping different aspects of the optimization process was briefly discussed and re-evaluated from a multi-criteria optimization perspective.

For a selection of cranial and spinal patient geometries, a high number of dose optimizations was performed. The influence of the three composite objective parameters for PTV vs OAR weighting, sparing of healthy tissue surrounding the PTV, and the amount of dose modulation in the resulting dose distributions (e.g. variation in dose rate, size and shape of MLC shapes) was analyzed by comparison of a selection of dosimetric criteria.

The influence of the position of a specific parameter slider on a specific criterion varied greatly with respect to the patient geometry, e.g. size, shape, and position of target volume and organ-at-risk structures, as well as the choice of criterion. In the cranial example cases where the goals for PTV dose coverage and the dose constraints could be easily fulfilled, the resulting dose distributions did not vary much. In artificially constructed, extremely challenging cases, those constraints often could not be fulfilled completely, but a much greater variation in the optimization results could be observed. For consistency reasons, all cranial treatment plans were optimized for the same symmetrical setup of three equidistant VMAT arcs, even if in some cases a more asymmetrical arrangement of the arcs would have been more favorable. In addition, the VMAT Elements treatment planning software offers more adjustment options than the three slider controls analyzed here. Therefore, in the scope of this study, the non-fulfillment of constraints is not of too much concern.

The spinal cases showed a higher variation in the optimization results and the dosimetric outcome was highly dependent on the shape of the PTV, the complete vertebral body with all pedicles being the most challenging one. Especially when the spinal canal was specified as OAR, the capabilities for modeling the various trade-offs were noticeable.

A methodology for approximation of the dose distributions resulting from not-yet-calculated parameter combinations based on trilinear dose interpolation was introduced and evaluated for two discretizations of the parameter space. Approximation using only eight pre-calculated optimization results for the extremal positions of the parameter sliders proved to be too coarse for all but the least challenging cranial cases for a number of reasons: On the one hand the value range of the slider parameters has been chosen in a way, that a set of intermediate slider positions is assumed to be the most likely region of the parameter space a potential user is going to explore. On the other hand, approximation using trilinear interpolation is best if the sought-after position is very close to one of the pre-calculated results (anchors). This means, that with only eight pre-calculated plans, the most relevant region of the parameter space coincides with the region where the deviations between approximated and actual result can be expected to be highest, which is not a desirable condition.

In order to mitigate this issue, optimization results for central slider positions were added to the pool of pre-calculated dose distributions, increasing the total number of anchors to 27. Even though this means some additional computational overhead, a general decrease of approximation errors, especially in the challenging geometries where the Pareto surface of

achievable solutions was of a more convoluted shape, could be observed. However, due to their fixed position in the parameter space, even this high number of pre-calculated results leads to problems, if a steep gradient in the solution space has to be approximated, warranting a more flexible distribution of the anchor dose distributions.

A machine-learning-based approach for finding regions in the parameter space with high variation in the resulting dose distributions, and thus steep gradients in the corresponding Pareto surfaces, was developed. The optimization data obtained during the experimental exploration of the solution space and the evaluation of the approximation method was leveraged to serve as basis for training and evaluation of machine learning models. Patient geometries were expressed by a vector of characterizing numerical values, such as PTV and OAR volume and -concavity, relative distance between objects, and properties of the overlap-volume-histograms. Combined with position information of two of the three composite parameter sliders, 3D-, 7D- and 13D- feature matrices were assembled. For a number of dosimetric criteria  $C_i$  a binary classification problem was generated as follows: Each sample feature vector was labeled according to a numerical threshold for  $C_i$ , stating whether movement of a specific slider is likely to noticeably influence the value of the  $C_i$ , or not. This information can then be used, e.g. for determining the position and number of anchors needed for approximation of the solution space. The classes turned out to be rather imbalanced in most cases, making it necessary to apply oversampling-techniques to restore class-balance, thus acquiring enough samples of each class to make a proper learning process possible.

Three kinds of machine-learning classifiers were trained and evaluated on the 3D-, 7D- and 13D-data: A support-vector machine (SVM), a random-forest classifier (RF) and artificial neural networks (ANN) with varying internal structure. Model performance was generally poor without oversampling, even with the highest number of features per sample. However, after random-oversampling, excellent precision-, recall-, and AUC values could be obtained in all cases, proving the effectiveness of the chosen approach.

In the final section, an algorithm for combining a number of such machine learning models was presented, which serves as a recommendation system for the treatment planner. Based on information from already planned similar treatments and the characteristics of the patient case currently at hand, the system automatically points out which slider control to modify, if the user wants to improve the value of a specific  $C_i$ . This has great potential for streamlining the existing treatment planning workflow, where the choice of slider positions for optimization is solely relying on user experience and random exploration.

Of course, there are also a number of limitations to this work to point out: The trilinear dose interpolation approach might seem naive at first glance, and except for the general statements presented in section 3.5.2 no exact mathematical error approximation is possible without demanding strong mathematical properties, such as convexity, from the optimizer. However, the data-driven approach chosen in this work is also its greatest advantage. Changes to the optimization algorithm, the exact nature and composition of the composite optimization parameters do not affect the general functionality of the approach. The possibility of adaptively distributing the the node points in the parameter space allows for higher resolution where needed, and the continuity and the choice of the parameter space in the optimizer guarantee that no massive irregularities in the resulting solution space have to be expected.

The effectiveness of the machine-learning techniques applied in section 4.1 is relying to a large extent on the amount and quality of the training data available. The author is completely aware

of the fact, that the amount of data available for developing the algorithms presented in this thesis only represents the bare minimum needed for showing feasibility and the potential to scale with much larger datasets. For widespread applications of machine learning like image recognition and natural language processing, labeled training sets comprising millions of samples are readily available for free. In medical applications handling extremely sensitive data like personal health records, labeled data on such scale is not yet available and, since topics like anonymization and cyber-security have to be dealt with, its acquisition is likely to be an expensive and time consuming process which can only be realized as a community effort. Nevertheless, aging populations and increasing cost for healthcare put increasing pressure on healthcare providers, and more and more resources are needed to provide effective and affordable treatment for a growing number of patients. Increasing automation of extremely labor-intensive tasks, such as radiotherapy treatment planning, with methods like the ones developed in this work, can play a big role in freeing up those resources.





**Chapter** **6**  
**Appendix**

## 6.1 Technical Background

### 6.1.1 Geometry of test patient data

**Table 6.1:** Geometry data of cranial test cases

Plan Name (Cranial)	Volume PTV [cm3]	Volume Brainstem [cm3]	Distance Brainstem PTV [mm]	OVH penumbra Brainstem[mm]	PTV concavity
Cranial Case #1	29.8	21.5	0	15	1.61
Cranial Case #2	9.8	19.6	3	35	1.26
Cranial Case #3	10.1	19.6	8	41	1.27
Cranial Case #4	12.2	19.6	1	43	1.25
Cranial Case #5	113.2	30.2	0	48	1.17
Cranial Case #6	3.9	30.2	5	55	1.15
Cranial Case #7	3.9	30.2	14	64	1.15
Cranial Case #8	3.9	30.2	30	67	1.15
Cranial Case #9	3.9	30.2	48	69	1.15
Cranial Case #10	3.9	30.2	15	50	1.15
Cranial Case #11	3.9	30.2	40	53	1.15
Cranial Case #12	3.7	30.2	1	66	1.18
Cranial Case #13	3.9	30.2	0	69	1.15
Cranial Case #14	3.9	30.2	1	64	1.15
Cranial Case #15	3.9	30.2	16	57	1.15
Cranial Case #16	3.9	30.2	38	54	1.15
Cranial Case #17	3.9	30.2	1	38	1.15
Cranial Case #18	6.7	19.6	0	35	1.27
Cranial Case #19	53.8	31.1	0	59	1.19
Cranial Case #20	7.8	25.3	6	44	1.5
Cranial Case #21	1.2	27.3	3	49	1.24
Cranial Case #22	16.1	27.2	0	38	1.84
Cranial Case #23	30.7	21.4	6	29	1.71
Cranial Case #24	4.1	26.2	5	42	1.26
Cranial Case #25	9.3	38	1	50	1.13
Cranial Case #26	1	13.3	9	38	1.18
Cranial Case #27	8.1	39.1	0	29	1.12
Cranial Case #28	1.5	29.2	6	39	2.03
Cranial Case #29	0.5	27.1	6	46	1.65
Cranial Case #30	3.9	30.2	17	32	1.15
Cranial Case #31	3.9	30.2	36	29	1.15
Cranial Case #32	4.7	21.5	0	32	1.36
Cranial Case #33	4.6	21.5	3	36	1.36
Cranial Case #34	4.7	21.5	1	35	1.36
Cranial Case #35	3.9	30.2	0	67	1.15
Cranial Case #36	3.9	30.2	9	57	1.15
Cranial Case #37	3.9	30.2	8	63	1.15
Cranial Case #38	17.4	30.2	1	60	1.9
Cranial Case #39	18.5	30.2	0	46	1.98
Cranial Case #40	22.1	30.2	0	48	1.66
Cranial Case #41	6.9	21.5	5	29	1.73
Cranial Case #42	8.2	21.5	1	42	1.65
Cranial Case #43	8.9	21.5	5	29	2.11
Cranial Case #44	15.5	21.5	5	26	1.89

**Table 6.2:** Geometry data of spinal test cases

Plan Name (Spine)	Volume PTV [cm3]	Volume Spinal Canal [cm3]	Distance Spinal Canal PTV [mm]	OVH penumbra Spinal Canal[mm] (inverse OVH)	PTV concavity
Spinal Case #1	45	22.6	0	38	2.24
Spinal Case #2	36.8	22.6	0	38	1.37
Spinal Case #3	57.3	22.6	0	38	2.92
Spinal Case #4	45.6	17.1	-1	40	1.85
Spinal Case #5	39	17.1	-1	39	1.39
Spinal Case #6	61.3	17.1	-1	40	2.81
Spinal Case #7	35.6	17.1	-1	35	1.82
Spinal Case #8	31.1	17.1	-1	35	1.34
Spinal Case #9	42.1	17.1	-1	35	2.07
Spinal Case #10	29.4	12	-1	35	2.03
Spinal Case #11	25.2	12	-1	35	1.94
Spinal Case #12	36.4	12	-1	35	2.26
Spinal Case #13	39.4	22.6	-1	39	1.84
Spinal Case #14	35.9	22.6	-1	39	1.7
Spinal Case #15	48.4	22.6	-1	39	2.24
Spinal Case #16	41	18.6	-1	41	1.97
Spinal Case #17	36.8	18.6	-1	41	1.55
Spinal Case #18	54.3	18.6	-1	41	2.75

## 6.1.2 Data generation and processing/software framework

**Table 6.3:** Overview of used software packages

<b>Purpose</b>	<b>Software package</b>	<b>Version</b>
Dose optimization	BL Elements Cranial Spine	var.
Data analysis (General framework)	Jupyter Notebook	Jupyter Notebook Server 4.3.1
Data analysis (Scientific computing)	Anaconda Python 2.7x	Python 2.7.13 / Anaconda 4.3.1
Data analysis (Data structures)	Pandas [55]	0.19.2
Machine learning (Framework RF, SVM)	Scikit learn	0.19.0
Machine learning (Framework for ANN)	Keras	2.2.2
Machine learning (ANN library)	Microsoft CNTK	2.6



# Bibliography

- [1] Brainlab AG. Integration with varian linacs, 2018. [[https://www.brainlab.com/wp-content/uploads/2018/11/ExacTrac\\_Sect6\\_No-Logo.jpg](https://www.brainlab.com/wp-content/uploads/2018/11/ExacTrac_Sect6_No-Logo.jpg), accessed March 5, 2019].
- [2] Aaron Babier, Justin J. Boutilier, Andrea L. McNiven, and Timothy C.Y. Chan. Knowledge-based automated planning for oropharyngeal cancer. *Medical Physics*, 45(7):2875–2883, 2018.
- [3] Rajamanickam Baskar, Kuo Ann Lee, Richard Ming-Chert Yeo, and Kheng-Wei Yeoh. Cancer and radiation therapy: Current advances and future directions. *International Journal of Medical Sciences*, 9(2):193–199, 2012.
- [4] Rasmus Bokrantz and Kaisa Miettinen. Projections onto the pareto surface in multicriteria radiation therapy optimization. *Medical Physics*, 42(10):5862–5870, 2015.
- [5] Rasumus Bokrantz. *Multicriteria optimization for managing tradeoffs in radiation therapy treatment planning*. PhD thesis, Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden, 2013.
- [6] Thomas Bortfeld. Imrt: a review and preview. *Physics in Medicine & Biology*, 51(13):R363, 2006.
- [7] Justin J. Boutilier, Tim Craig, Michael B. Sharpe, and Timothy C. Y. Chan. Sample size requirements for knowledge-based treatment planning. *Medical Physics*, 43(3):1212–1221, 2016.
- [8] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, Oct 2001.
- [9] M. Kara Bucci, Alison Bevan, and Mack Roach. Advances in radiation therapy: Conventional to 3d, to imrt, to 4d, and beyond. *CA: A Cancer Journal for Clinicians*, 55(2):117–134, 2005.
- [10] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16:321–357, 2002.
- [11] Huixiao Chen, David L. Craft, and David P. Gierga. Multicriteria optimization informed vmat planning. *Medical Dosimetry*, 39(1):64 – 73, 2014.
- [12] Wei Chen, Jan Unkelbach, Alexei Trofimov, Thomas Madden, Hanne Kooy, Thomas Bortfeld, and David L. Craft. Including robustness in multi-criteria optimization for intensity modulated proton therapy. *Physics in Medicine & Biology*, 57(3):591–608, 2012.
- [13] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

- [14] David L. Craft. Multi-criteria optimization methods in radiation therapy planning: a review of technologies and directions. *eprint arXiv:1305.1546*, 2013.
- [15] David L. Craft and Thomas Bortfeld. How many plans are needed in an imrt multi-objective plan database? *Physics in Medicine & Biology*, 53(11):2785–96, 2008.
- [16] David L. Craft, Tarek Halabi, Helen A Shih, and Thomas R Bortfeld. Approximating convex pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*, 33:3399–407, 10 2006.
- [17] David L. Craft, Tarek F. Halabi, Helen A. Shih, and Thomas R. Bortfeld. Approximating convex pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*, 33(9):3399–3407, 2006.
- [18] David L. Craft, Theodore S. Hong, Helen A. Shih, and Thomas R. Bortfeld. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *International Journal of Radiation Oncology • Biology • Physics*, 82(1):e83–e90, Jan 2012.
- [19] David L. Craft, Dualta McQuaid, Jeremiah Wala, Wei Chen, Ehsan Salari, and Thomas Bortfeld. Multicriteria vmat optimization. *Medical Physics*, 39(2):686–696, Feb 2012.
- [20] David L. Craft, Dávid Papp, and Jan Unkelbach. Plan averaging for multicriteria navigation of sliding window imrt and vmat. *Medical Physics*, 41(2):021709–1–021709–5, 2014.
- [21] Nello Cristianini and John Shawe-Taylor. *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press, 2000.
- [22] Yunfei Cui, Zhiqiang Geng, Qunxiong Zhu, and Yongming Han. Review: Multi-objective optimization methods and application in energy saving. *Energy*, 125:681 – 704, 2017.
- [23] Germund Dahlquist and Åke Björck. *Numerical Methods*. Dover Books on Mathematics. Dover Publications, 2003.
- [24] Matthias Ehrgott. *Multicriteria Optimization*. Springer, Berlin Heidelberg, 2005.
- [25] Issam El Naqa, Dan Ruan, Gilmer Valdes, Andre Dekker, Todd McNutt, Yaorong Ge, Q. Jackie Wu, Jung Hun Oh, Maria Thor, Wade Smith, Arvind Rao, Clifton Fuller, Ying Xiao, Frank Manion, Matthew Schipper, Charles Mayo, Jean M. Moran, and Randall Ten Haken. Machine learning and modeling: Data, validation, communication challenges. *Medical Physics*, 45(10):e834–e840, 2018.
- [26] Khaled Fawagreh, Mohamed Medhat Gaber, and Eyad Elyan. Random forests: from early developments to recent advancements. *Systems Science & Control Engineering*, 2(1):602–609, 2014.
- [27] Mary Feng, Gilmer Valdes, Nayha Dixit, and Timothy D Solberg. Machine learning in radiation oncology: Opportunities, requirements, and needs. *Frontiers in Oncology*, 8:110, 2018.

- [28] R.W. Freund and R.W. Hoppe. *Stoer/Bulirsch: Numerische Mathematik 1*. Springer-Lehrbuch. Springer Berlin Heidelberg, 2007.
- [29] Hastings Katherine G., Boothroyd Derek B., and Kappahn Kristopher et al. Socioeconomic differences in the epidemiologic transition from heart disease to cancer as the leading cause of death in the united states, 2003 to 2015: An observational study. *Annals of Internal Medicine*, 2018.
- [30] Serena Gianfaldoni, Roberto Gianfaldoni, Uwe Wollina, Jacopo Lotti, Georgi Tchernev, and Torello Lotti. An overview on radiotherapy: From its history to its current applications in dermatology. *Macedonian Journal of Medical Sciences*, 07 2017.
- [31] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [32] Tarek Halabi, David L. Craft, and Thomas Bortfeld. Dose–volume objectives in multi-criteria optimization. *Physics in Medicine & Biology*, 51(15):3809, 2006.
- [33] Alon Halevy, Peter Norvig, and Fernando Pereira. The unreasonable effectiveness of data. *IEEE Intelligent Systems*, 24(2):8–12, 2009.
- [34] Eric J. Hall. *Radiobiology for the Radiologist, 4th Ed.* Lippincott Williams & Wilkins, Philadelphia Baltimore New York London Buenos Aires Hong Kong Sydney Tokyo, 1994.
- [35] Horst.W. Hamacher and Karl.-H. Küfer. Inverse radiation therapy planning – a multiple objective optimization approach. *Discrete Applied Mathematics*, 118(1):145 – 161, 2002. Special Issue devoted to the ALIO-EURO Workshop on Applied Combinatorial Optimization.
- [36] Trevor Hastie, Robert Tibshirani, and Jerome H. Friedman. *The elements of statistical learning: data mining, inference, and prediction, 2nd Edition*. Springer series in statistics. Springer, 2009.
- [37] Haibo He, Yang Bai, Edwardo A. Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *Proceedings of the International Joint Conference on Neural Networks*, pages 1322 – 1328, 07 2008.
- [38] Theodore S. Hong, David L. Craft, Fredrik Carlsson, and Thomas R. Bortfeld. Multicriteria optimization in imrt treatment planning for locally advanced cancer of the pancreatic head. *International Journal of Radiation Oncology • Biology • Physics*, 72(4):1208–1214, Nov 2008.
- [39] Justin J. Boutilier, Taewoo Lee, Tim Craig, Michael B. Sharpe, and Timothy C. Y. Chan. Models for predicting objective function weights in prostate cancer imrt. *Medical Physics*, 42(04):1586–1595, 2015.
- [40] Sophia C. Kamran, Birgit S. Mueller, Peter Paetzold, Joseph Dunlap, Andrzej Niemierko, Thomas Bortfeld, Henning Willers, and David L. Craft. Multi-criteria optimization achieves superior normal tissue sparing in a planning study of intensity-modulated radiation therapy for rtog 1308-eligible non-small cell lung cancer patients. *Radiotherapy and Oncology*, 118(3):515 – 520, 2016.

- [41] Henry R. Kang. *Computational Color Technology (SPIE Press Monograph Vol. PM159)*. SPIE-International Society for Optical Engineering, 2006.
- [42] Tomas Kazda, Radim Jancalek, Petr Pospisil, Ondrej Sevela, Tomas Prochazka, Miroslav Vrzal, Petr Burkon, Marek Slavik, Ludmila Hynkova, Pavel Slampa, and Nadia N. Laack. Why and how to spare the hippocampus during brain radiotherapy: The developing role of hippocampal avoidance in cranial radiotherapy. *Radiation Oncology*, 9(1), 6 2014.
- [43] Michael Kazhdan, Patricio Simari, Todd Mcnutt, Binbin Wu, Robert Jacques, Ming Chuang, and Russell Taylor. A shape relationship descriptor for radiation therapy planning. In *Proceedings of the 12th International Conference on Medical Image Computing and Computer-Assisted Intervention: Part II, MICCAI '09*, pages 100–108, Berlin, Heidelberg, 2009. Springer-Verlag.
- [44] Eamonn Keogh and Abdullah Mueen. *Curse of Dimensionality*, pages 314–315. Springer US, Boston, MA, 2017.
- [45] Faiz M. Khan. *The Physics of Radiation Therapy, 2nd Edition*. Lippincott Williams & Wilkins, Philadelphia Baltimore New York London Buenos Aires Hong Kong Sydney Tokyo, 2005.
- [46] Fazal Khan and David L. Craft. 3d conformal planning using low segment multi-criteria imrt optimization. *Practical Radiation Oncology*, 5(2):e103–e111, Aug 2015.
- [47] Roel GJ Kierkels, Ruurd Visser, Hendrik P. Bijl, Johannes A. Langendijk, Aart A. van 't Veld, Roel JHM Steenbakkers, and Erik W. Korevaar. Multicriteria optimization enables less experienced planners to efficiently produce high quality treatment plans in head and neck cancer radiotherapy. *Radiation Oncology*, 10(1):87, Apr 2015.
- [48] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [49] Karl-H. Küfer, Michael Monz, Alexander Scherrer, Philipp Süß, Fernando Alonso, Ahmad S.A. Sultan, Thomas Bortfeld, David L. Craft, and Christian Thieke. Multicriteria optimization in intensity modulated radiotherapy planning. Technical Report 77, Fraunhofer (ITWM), 2005.
- [50] Michael Lahanas, Eduard Schreibmann, and Dimos Baltas. Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms. *Physics in Medicine & Biology*, 48(17):2843–71, 2003.
- [51] Manuel Lederman. The early history of radiotherapy: 1895-1939. *International Journal of Radiation Oncology • Biology • Physics*, 7(5):639–648, 1981.
- [52] Laura Masi, Raffaella Doro, Virginia Favuzza, Samantha Cipressi, and Lorenzo Livi. Impact of plan parameters on the dosimetric accuracy of volumetric modulated arc therapy. *Medical Physics*, 40(7):0717–18, 2013.



- [53] Matthew J. McGirt, Theodore Speroff, Robert S. Dittus, Frank E. Harrell, and Anthony L. Asher. The national neurosurgery quality and outcomes database (n2qod): general overview and pilot-year project description. *Neurosurgical Focus FOC*, 34(1):E6, 2013.
- [54] Chris McIntosh, Mattea Welch, Andrea McNiven, David A Jaffray, and Thomas G Purdie. Fully automated treatment planning for head and neck radiotherapy using a voxel-based dose prediction and dose mimicking method. *Physics in Medicine & Biology*, 62(15):5926–44, 2017.
- [55] Wes McKinney. Data structures for statistical computing in python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 51 – 56, 2010.
- [56] Andrea L. McNiven, Michael B. Sharpe, and Thomas G. Purdie. A new metric for assessing imrt modulation complexity and plan deliverability. *Medical Physics*, 37(2):505–515, 2010.
- [57] Michael Monz. *Pareto Navigation -interactive multiobjective optimisation and its application in radiotherapy planning*. PhD thesis, Department of Mathematics, Universität Kaiserslautern, Kaiserslautern, Germany, 2006.
- [58] Michael Monz, Karl-H. Küfer, Thomas R. Bortfeld, and Christian Thieke. Pareto navigation-algorithmic foundation of interactive multi-criteria imrt planning. *Physics in Medicine & Biology*, 53(4):985–98, 2008.
- [59] Michael A. Nielsen. *Neural Networks and Deep Learning*. Determination Press, 2015.
- [60] Andrzej Niemierko. Reporting and analyzing dose distributions: A concept of equivalent uniform dose. *Medical Physics*, 24(1):103–110, 1997.
- [61] Thais Mayumi Oshiro, Pedro Santoro Perez, and José Augusto Baranauskas. How many trees in a random forest? In Petra Perner, editor, *Machine Learning and Data Mining in Pattern Recognition*, pages 154–168, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [62] Ian Paddick. A simple scoring ratio to index the conformity of radiosurgical treatment plans. *Journal of Neurosurgery JNS*, 93(Technical note):219 – 222, 2000.
- [63] Ian Paddick and Bodo Lippitz. A simple dose gradient measurement tool to complement the conformity index. *Journal of Neurosurgery JNS*, 105(Supplement):194 – 201, 2006.
- [64] Richard Powis, Andrew Bird, Matthew Brennan, Susan Hinks, Hannah J. H. Newman, Katie Reed, John Philip Sage, and Gareth J. Webster. Clinical implementation of a knowledge based planning tool for prostate vmat. *Radiation Oncology*, 12(1):81, 2017.
- [65] Enzhuo Quan, Xiaoqiang Li, Yupeng Li, Xiaochun Wang, Rajat J Kudchadker, Jennifer L Johnson, Deborah Kuban, Andrew K Lee, and Xiaodong Zhang. A comprehensive comparison of imrt and vmat plan quality for prostate cancer treatment. *International Journal of Radiation Oncology • Biology • Physics*, 83:1169–78, 07 2012.

- [66] Berkman Sahiner, Aria Pezeshk, Lubomir M. Hadjiiski, Xiaosong Wang, Karen Drukker, Kenny H. Cha, Ronald M. Summers, and Maryellen L. Giger. Deep learning in medical imaging and radiation therapy. *Medical Physics*, 46(1):e1–e36, 2019.
- [67] Rolf Sauer. *Strahlentherapie und Onkologie*. Elsevier, Urban & Fischer, 2010.
- [68] Ji Serna, Michael Monz, Karl-H. Küfer, and Christian Thieke. Trade-off bounds for the pareto surface approximation in multi-criteria imrt planning. *Physics in Medicine & Biology*, 54(20):6299, 2009.
- [69] Jason P. Sheehan, Inga Grills, Veronica L. Chiang, Huamei Dong, Arthur Berg, Ronald E. Warnick, Douglas Kondziolka, and Brian Kavanagh. Quality of life outcomes for brain metastasis patients treated with stereotactic radiosurgery: pre-procedural predictive factors from a prospective national registry. *Journal of Neurosurgery JNS*, pages 1–7, 2018.
- [70] Jason P. Sheehan, Brian D. Kavanagh, Anthony Asher, and Robert E. Harbaugh. Inception of a national multidisciplinary registry for stereotactic radiosurgery. *Journal of Neurosurgery JNS*, 124(1):155 – 162, 2016.
- [71] Chenyang Shen, Yesenia Gonzalez, Peter Klages, Nan Qin, Hyunuk Jung, Liyuan Chen, Dan Nguyen, Steve B. Jiang, and Xun Jia. Intelligent inverse treatment planning via deep reinforcement learning, a proof-of-principle study in high dose-rate brachytherapy for cervical cancer. *arXiv:1811.10102 [physics.med-ph]*, 2018.
- [72] Satomi Shiraishi and Kevin L. Moore. Knowledge-based prediction of three-dimensional dose distributions for external beam radiotherapy. *Medical Physics*, 43(1):378–387, 1 2016.
- [73] Satomi Shiraishi, Jun Tan, Lindsey A. Olsen, and Kevin L. Moore. Knowledge-based prediction of plan quality metrics in intracranial stereotactic radiosurgery. *Medical Physics*, 42(2):908–917, 2015.
- [74] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016.
- [75] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- [76] Theodor Stewart, Oliver Bandte, Heinrich Braun, Nirupam Chakraborti, Matthias Ehr Gott, Mathias Göbelt, Yaochu Jin, Hirotaka Nakayama, Silvia Poles, and Danilo Di Stefano. *Real-World Applications of Multiobjective Optimization*, pages 285–327. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [77] Karen M. Teoh and Catharine et al Clark. Volumetric modulated arc therapy: A review of current literature and clinical use in practice. *The British Journal of Radiology*, 84(11):967–96, 2011.

- [78] Jim P. Tol, Max Dahele, Alexander R. Delaney, Ben J. Slotman, and Wilko F. A. R. Verbakel. Can knowledge-based dvh predictions be used for automated, individualized quality assurance of radiotherapy treatment plans? *Radiation Oncology*, 10(1):234, 2015.
- [79] Seiji Tomori, Noriyuki Kadoya, Yoshiki Takayama, Tomohiro Kajikawa, Katsumi Shima, Kakutarou Narazaki, and Keiichi Jingu. A deep learning-based prediction model for gamma evaluation in patient-specific quality assurance. *Medical Physics*, 45(9):4055–4065, 2018.
- [80] Ander Urruticoechea. Recent advances in cancer therapy: An overview. *Current Pharmaceutical Design*, 6(1):3–10, 2010.
- [81] Gilmer Valdes, Timothy D Solberg, Marina Heskell, Lyle Ungar, and Charles B Simone II. Using machine learning to predict radiation pneumonitis in patients with stage i non-small cell lung cancer treated with stereotactic body radiation therapy. *Physics in Medicine & Biology*, 61(16):6105, 2016.
- [82] Peter W.J. Voet, Maarten L.P. Dirksen, Sebastiaan Breedveld, Dennie Fransen, Peter C. Levendag, and Ben J.M. Heijmen. Toward fully automated multicriterial plan generation: A prospective clinical study. *International Journal of Radiation Oncology • Biology • Physics*, 85(3):866 – 872, 2013.
- [83] Huan Wang, Peng Dong, Hongcheng Liu, and Lei Xing. Development of an autonomous treatment planning strategy for radiation therapy with effective use of population-based prior data. *Medical Physics*, 44(2):389–396, 2017.
- [84] Steve Webb. *Intensity-Modulated Radiation Therapy*. CRC Press, 2001.
- [85] Steve Webb. The physical basis of imrt and inverse planning. *British Journal of Radiology*, 76(678):689, 2003.
- [86] Binbin Wu, Dalong Pang, Patricio Simari, Russell Taylor, Giuseppe Sanguineti, and Todd McNutt. Using overlap volume histogram and imrt plan data to guide and automate vmat planning: A head-and-neck case study. *Medical Physics*, 40(2):021714–1–7, 2013.
- [87] Binbin Wu, Francesco Ricchetti, Giuseppe Sanguineti, Misha Kazhdan, Patricio D. Simari, Ming Chuang, Russell Taylor, Robert Jacques, and Todd McNutt. Patient geometry-driven information retrieval for imrt treatment plan quality control. *Medical Physics*, 36(12):5497–505, 2009.
- [88] Kelly C. Younge, Martha M. Matuszak, Jean M. Moran, Daniel L. McShan, Benedick A. Fraass, and Donald A. Roberts. Penalization of aperture complexity in inversely planned volumetric modulated arc therapy. *Medical Physics*, 39(11):7160–7170, 2016.
- [89] Jiahan Zhang, Q. Jackie Wu, Tianyi Xie, Yang Sheng, Fang-Fang Yin, and Yaorong Ge. An ensemble approach to knowledge-based intensity-modulated radiation therapy planning. *Frontiers in Oncology*, 8:57, 2018.

- [90] Xiangxin Zhu, Carl Vondrick, Deva Ramanan, and Charless Fowlkes. Do we need more training data or better models for object detection? In *Proceedings of the British Machine Vision Conference*, pages 80.1–80.11. BMVA Press, 2012.
- [91] Xiaofeng Zhu, Yaorong Ge, Taoran Li, Danthai Thongphiew, Fang-Fang Yin, and Q Jackie Wu. A planning quality evaluation tool for prostate adaptive imrt based on machine learning. *Medical Physics*, 38(2):719–726, 2011.

# List of Figures

2.1	Modern clinical Linac with MLC, patient positioning system and treatment couch [1] . . . . .	7
2.2	Normalized relative depth dose (PDD) of a 6MV photon beam for 10x10mm(green), 60x60mm(red) and 100x100mm(blue) MLC square field size . . . . .	7
2.3	Schematic overview of target-outline (red), primary jaws (gray) and MLC leaves shaping the cross-section of the treatment beam (in reality, MLCs can have up to 120 leaf-pairs). . . . .	8
2.4	Aperture of an IMRT beam with beamlet discretization and fluence . . . . .	9
2.5	Dose-Volume-Histogram (DVH) graphs for two objects: PTV (red) and OAR (blue) . . . . .	10
2.6	2D example of Pareto optimality . . . . .	12
2.7	Illustration of 7D-approximation base . . . . .	15
2.8	Illustration of the approximation steps with a 7D-approximation base . . . . .	16
2.9	Illustration of the general trilinear interpolation problem . . . . .	17
2.10	First interpolation step . . . . .	18
2.11	Second interpolation step . . . . .	18
2.12	Final interpolation step . . . . .	19
2.13	Parameter space sub-division . . . . .	19
2.14	Adaptive resolution of the parameter space discretization . . . . .	20
2.15	Generation of Overlap-Volume-Histograms(OVH) . . . . .	21
2.16	OVH (blue) and inverse OVH (red, dashed) for OAR objects of different length with otherwise identical geometrical setup . . . . .	22
2.17	2-D object(blue) and its convex hull(red) . . . . .	22
2.18	ROC-curve (red) and AUC (green) . . . . .	26
2.19	2D-example of SVM with separable data . . . . .	26
2.20	structure of the Random-Forest(RF) method . . . . .	27
2.21	Input and output of as single neuron $n$ . . . . .	27
2.22	<i>Sigmoid</i> activation function $f(x)$ . . . . .	28
2.23	<i>ReLU</i> activation function $f(x)$ . . . . .	28
2.24	Three-layer neural network with 5 neurons in the hidden layer . . . . .	29
3.1	Screenshot of the main planning GUI of the Cranial Element . . . . .	32
3.2	Screenshot after optimization showing dose distributions and DVH . . . . .	33
3.3	The three slider controls for adjusting the composite meta-objectives . . . . .	34
3.4	Patient geometries of spine cases #7-#9. PTV (orange) and cropped spinal canal (OAR, purple). . . . .	37

3.5	Patient geometry images of cranial cases #5 and #36 including PTV (orange) and brainstem (OAR, purple). Other objects shown are the optical structures including the eyes. Case #5 is a much harder task for the optimizer due to PTV shape and size. In addition the PTV is directly engulfing large parts of the OAR.	37
3.6	Boxplots for relative PTV coverage, 343 optimizations per case	39
3.7	Boxplots for relative PTV coverage, only 'Modulation' values of 3 or higher allowed	39
3.8	Pareto surface of feasible PTV coverage in a selected cranial VMAT case (case #40), <i>Modulation</i> slider in minimum (left image) middle (center image) and maximum position (right image)	40
3.9	Boxplots of PTV Max. Dose, 343 optimizations per case	40
3.10	Boxplots of PTV min. Dose, 343 optimizations per case	41
3.11	Boxplots of PTV Mean Dose, 343 optimizations per case	41
3.12	Boxplots of Brainstem $D_{10\%}$ , 343 optimizations per case	42
3.13	Boxplots of Brainstem Mean Dose, 343 optimizations per case	42
3.14	Boxplots of Brainstem Min. Dose, 343 optimizations per case	43
3.15	Boxplots of Brainstem Max. Dose, 343 optimizations per case	43
3.16	Boxplots of Gradient Index (GI), 343 optimizations per case	44
3.17	Boxplots of Conformity Index (CI), 343 optimizations per case	44
3.18	Boxplots of monitor units (MU), 343 optimizations per case	45
3.19	Boxplots of monitor units (MU), only 'Modulation' values of 3 or higher allowed	45
3.20	Pareto surface of feasible $D_{10\%}$ values in the Brainstem of a cranial VMAT case, <i>Modulation</i> slider in middle position, same plan optimized with standard (7 equidistant positions, right) and double (14 equidistant positions, left) slider resolution	46
3.21	Boxplots for relative PTV coverage, 343 optimizations per case	47
3.22	Boxplots of PTV Max. Dose, 343 optimizations per case	47
3.23	Boxplots of PTV Min. Dose, 343 optimizations per case	48
3.24	Boxplots of PTV Mean Dose, 343 optimizations per case	48
3.25	Boxplots of Spinal Canal $D_{0.25cm^3}$ , 343 optimizations per case	49
3.26	Boxplots of Spinal Canal Mean Dose, 343 optimizations per case	49
3.27	Boxplots of Spinal Canal Min. Dose, 343 optimizations per case	50
3.28	Boxplots of Spinal Canal Max. Dose, 343 optimizations per case	51
3.29	Boxplots of Gradient Index (GI), 343 optimizations per case	51
3.30	Boxplots of Conformity Index (CI), 343 optimizations per case	52
3.31	Boxplots of monitor units (MU), 343 optimizations per case	52
3.32	Illustration of the trilinear interpolation problem mapped onto the 3D space spanned by the of VMAT optimization sliders	53
3.33	Linear interpolation along <i>Dose Modulation</i> axis	54
3.34	Linear interpolation along <i>Normal Tissue Sparing</i> axis	54
3.35	Linear interpolation along <i>PTV vs OAR Weighting</i> axis	55
3.36	Pattern of additionally pre-calculated slider positions (red) embedded in already available results (black). 2D slice of the parameter space	55
3.37	Optimization results taken into account for approximation at cross shown in green. 2D slice of the parameter space	55

3.38	Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases. . . . .	60
3.39	Pareto surfaces for the PTV dose coverage in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position . . . . .	60
3.40	Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases . . . . .	61
3.41	Pareto surfaces for the PTV dose coverage in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position. . . . .	61
3.42	Boxplots of difference between mean dose in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases. . . . .	62
3.43	Pareto surfaces for the Mean Dose of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position . . . . .	62
3.44	Boxplots of difference between mean dose in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases. . . . .	63
3.45	Pareto surfaces for the Mean Dose of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position. . . . .	63
3.46	Boxplots of difference between $D_{10\%}$ in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, cranial example cases. . . . .	64
3.47	Pareto surfaces for $D_{10\%}$ of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position. . . . .	64
3.48	Boxplots of difference between $D_{0.25cm^3}$ in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 8 anchor points, spinal example cases. . . . .	65
3.49	Pareto surfaces for $D_{0.25cm^3}$ of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 8 anchor points, 'Modulation' slider in central position. . . . .	65
3.50	Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases. . . . .	66
3.51	Pareto surfaces for the PTV dose coverage in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	67
3.52	Boxplots of difference between PTV dose coverage in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases. . . . .	67
3.53	Pareto surfaces for the PTV dose coverage in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	68
3.54	Boxplots of difference between mean dose in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases. . . . .	68

3.55	Pareto surfaces for the Mean Dose of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	69
3.56	Boxplots of difference between mean dose in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases. . . . .	69
3.57	Pareto surfaces for the Mean Dose of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	70
3.58	Boxplots of difference between $D_{10\%}$ in OAR 'Brainstem' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, cranial example cases. . . . .	70
3.59	Pareto surfaces for $D_{10\%}$ of OAR 'Brainstem' in cranial case #1. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	71
3.60	Boxplots of difference between $D_{0.25cm^3}$ in OAR 'Spinal Canal' in interpolated dose distribution and actual optimization results. Trilinear interpolation with 27 anchor points, spinal example cases. . . . .	71
3.61	Pareto surfaces for $D_{0.25cm^3}$ of OAR 'Spinal Canal' in spine case #4. Trilinear interpolation with 27 anchor points, 'Modulation' slider in central position. . . . .	72
4.1	Generation of Training Data . . . . .	78
4.2	Schematic view of AI driven prediction algorithm . . . . .	79
4.3	Schematic view of Feature Matrix X and Target Vector y . . . . .	80
4.4	Overlap-volume-histograms of the 'Brainstem'-objects in the cranial patient geometries analyzed in this work . . . . .	81
4.5	Overlap-volume-histogram and the extracted features . . . . .	82
4.6	Histograms of OAR $D_{10\%}$ values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 1 Gy OAR Constraint Dose range threshold (right) . . . . .	83
4.7	Histograms of the two sample classes resulting by introduction of a 1 Gy OAR Constraint point dose range threshold after random-oversampling . . . . .	83
4.8	ROC curves for the OAR Constraint point dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	84
4.9	Histograms of relative PTV volume dose coverage values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 5% PTV volume range threshold (right) . . . . .	85
4.10	ROC curves for the Cranial PTV coverage classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	86
4.11	Histograms of PTV 1% Dose values in the Cranial cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 1Gy range threshold (right) . . . . .	86



4.12	ROC curves for the Cranial 1% PTV Dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	87
4.13	Histograms of GI values in the Cranial cases (left), the GI range achievable with the 'Normal Tissue Sparing' slider (center), and the two sample classes resulting by introduction of a GI range threshold of 1 (right) . . . . .	88
4.14	ROC curves for the Cranial GI classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	88
4.15	Histograms of MU values in the Cranial cases (left), the MU range achievable with the 'Modulation' slider (center), and the two sample classes resulting by introduction of a 1000 MU range threshold (right) . . . . .	89
4.16	ROC curves for the Cranial MU classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	90
4.17	Histograms of OAR $D_{0.25cm^3}$ values in the Spine cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 6 Gy OAR Constraint Dose range threshold (right)	91
4.18	ROC curves for the OAR Constraint point dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance.	91
4.19	Histograms of relative PTV volume dose coverage values in the Spine cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 5% PTV volume range threshold (right) . . . . .	92
4.20	ROC curves for the spine PTV coverage classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	93
4.21	Histograms of PTV 1% Dose values in the Spinal cases (left), the dose range achievable with the 'Weighting' slider (center), and the two sample classes resulting by introduction of a 4 Gy range threshold (right) . . . . .	94
4.22	ROC curves for the Spinal 1% PTV Dose classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . .	94
4.23	Histograms of GI values in the spine cases (left), the GI range achievable with the 'Normal Tissue Sparing' slider (center), and the two sample classes resulting by introduction of a GI range threshold of 1 (right) . . . . .	95
4.24	ROC curves for the Spine GI classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature vector. Upper row is without, lower row with random-oversampling to overcome class-imbalance. . . . .	95
4.25	Histograms of MU values in the Spine cases (left), the MU range achievable with the 'Modulation' slider (center), and the two sample classes resulting by introduction of a 4500 MU range threshold (right) . . . . .	96

---

4.26 ROC curves for the Spine MU classification problem with ANN, RF and SVM classifier and 3D-,7D- and 11D-Feature Vector. Upper row is without, lower row with Random-Oversampling to overcome class-imbalance. . . . . 97

5.1 Complete prediction algorithm including all slider controls . . . . . 103

# List of Tables

3.1	Geometric organ structures in the data . . . . .	37
3.2	Planning constraints . . . . .	38
4.1	Geometry vector in cranial cases . . . . .	82
4.2	Geometry vector in spine cases . . . . .	82
4.3	AUC values without Random Oversampling, Cranial . . . . .	98
4.4	AUC values without Random Oversampling, Spine . . . . .	98
4.5	AUC values with Random Oversampling, Cranial . . . . .	99
4.6	AUC values with Random Oversampling, Spine . . . . .	99
4.7	Precision and recall values with 11D feature vector and Random Oversampling, cranial cases. Comparison of ANNs with 5 and 64 hidden units (HU), after 200 training epochs . . . . .	99
4.8	Precision and recall values with 11D feature vector and Random Oversampling, spine cases. Comparison of ANNs with 5 and 64 hidden units (HU), after 200 training epochs . . . . .	100
6.1	Geometry data of cranial test cases . . . . .	110
6.2	Geometry data of spinal test cases . . . . .	110
6.3	Overview of used software packages . . . . .	111



# List of Publications

## **Publications that concern subjects covered by this thesis:**

1. Christian Harrer, Stefan Schell, and Jan J. Wilkens. Multicriteria-optimized treatment planning for volumetric modulated arc therapy (VMAT) using high-level composite optimization parameters and trilinear dose interpolation  
*AAPM meeting abstract, Medical Physics* 45 (6), 2018
2. Christian Harrer, Stefan Schell, and Jan J. Wilkens. Technical Note: approximation of dose quality criterion values in multi-criteria optimized (MCO) volumetric modulated arc therapy (VMAT) treatment planning using trilinear dose interpolation.  
*manuscript in preparation*
3. Christian Harrer, Wolfgang Ullrich, and Jan J. Wilkens. Data driven prediction of multi-criteria optimization (MCO) parameter efficiency in volumetric modulated arc therapy (VMAT) treatment planning using machine learning (ML)  
*manuscript in preparation*
4. Christian Harrer and Wolfgang Ullrich. Intelligent optimization setting adjustment for radiotherapy treatment planning using patient geometry information and artificial intelligence  
*International Patent Application* PCT No.: PCT/EP2019/055201, 2019

## **Publications that do not concern subjects covered by this thesis:**

1. Christian Harrer et al. Semi-automatic detection of Gd-DTPA-saline filled capsules for colonic transit time assessment in MRI  
*Medical Imaging 2008: Image Processing, SPIE Proceedings* 6914, 2008
2. Christian Harrer et al. Detection of the heartbeat in cranial accelerometer data using independent component analysis  
*United States Patent Application* No. US 2018/0296107 A1, 2018  
PCT No.: PCT/EP2015/059279, 2015



# Acknowledgements

Working towards a PhD is a big project, and successful completion is virtually impossible without people willing to provide support, feedback and encouragement in times of need.

First of all I would like to thank Prof. Jan J. Wilkens from the TUM for accepting me as his external PhD student and for always finding time in his busy schedule to discuss the progress of my work, give feedback and show ways for improvement. Your calm, positive and always encouraging way was an invaluable support during the past years! I would also like to thank everybody else from the TUM, especially Prof. Franz Pfeiffer and Prof. Martin Zacharias for taking the remaining roles in the final examination committee.

Secondly, I owe a big thank you to everybody involved from Brainlab, especially Marc Göpfert and Wolfgang Ullrich, who initially gave me the the chance to embark on this journey, and Stefan Schell, who, together with Wolfgang, always found time for supervision and fruitful discussions. I am also extremely grateful to Cornelis Kamerling for sharing his experience in academia and providing unbiased yet always constructive feedback in the last year, and to all the remaining great people from the R&D Radiotherapy department at Brainlab.

Finally, I would like to thank everybody else who provided support in so many ways during during these challenging years, especially my parents. Without you I would have never come this far!