

DIPLOMARBEIT

Advancing Deep Learning for Automated Treatment Planning in Adaptive Radiotherapy

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Technische Physik

eingereicht von

Simon Glatzer

Matrikelnummer: 12019265

ausgeführt am
Atominstitut der Fakultät für Physik
Technische Universität Wien

in Zusammenarbeit mit der
Universitätsklinik für Radioonkologie
Medizinische Universität Wien

Betreuung
Priv.-Doz. Mag. Peter Kueß, PhD
Priv.-Doz. Gerd Heilemann, MSc PhD
Lukas Zimmermann, MSc PhD

Wien, 23.06.2025

Unterschrift Verfasser

Unterschrift Betreuer

Abstract

Radiotherapy is an important treatment modality for a variety of cancers and especially prostate cancer, which is one of the most commonly diagnosed malignancies in men globally. The increasing incidence, coupled with the demand for individualized treatment strategies, underscores the need for more efficient and adaptable planning workflows. Currently treatment planning requires the calculation of a set of parameters, which is typically reliant on expert-driven, iterative optimization processes in commercial Treatment Planning Systems (TPSs). These are however time-consuming and inflexible, and as such ill-suited for modern online adaptive radiotherapy.

In this thesis, two new strategies were implemented and tested with the goal of improving a previously developed Deep Learning (DL)-based model for generating machine-ready Volumetric Modulated Arc Therapy (VMAT) plans for prostate cancer. The model can predict the necessary machine parameters for a Linear Accelerator (LINAC), i.e. the Multi-Leaf Collimator (MLC)-sequences and the Monitor Units (MUs) directly from Computed Tomography (CT) images, structure information and a predefined dose distribution. Building upon the fundamental work and suggestions provided in previous work [1], [2], this thesis introduces several modifications aimed at improving the model's generalizability and the produced dose distributions. Firstly, an object detection based model was developed and the network adapted to this end. This modification replaced the previously used loss function, which was based solely on the deviation between predicted and reference machine parameters with a more spatially informed loss function. Secondly, the network was modified to work with a 3D slice of the full 3D representation of the patient obtained from medical imaging. Furthermore, a dose calculation algorithm was implemented to validate the predicted dose distributions. The initial test and validation set consisted of 117 patients and a final set of 310 training and 104 validation data sets was obtained through data augmentation. The performance of the modified models was then evaluated by analyzing the training performance and performing dosimetric tests on the generated plans for 34 patients.

The proposed modifications to the model ultimately failed to improve the performance. The Gamma Passing Rate (GPR) of the modified models dropped by about 10 % percentage points compared to the base model. While all models achieved the same mean dose as in the ground truth in the Planning Target Volume (PTV), the volume receiving 95 % of the dose decreased by 5 % and 23 % for the two new models respectively, indicating less coverage. However, especially the object detection based method shows great promise if paired with more advanced methods and a specifically tailored model. Furthermore, it was observed that the model stopped producing better dose distribution early on in the training and instead focused primarily on mimicking the original machine parameters better. This shows that the current geometrical approach of training the model directly on the prediction of the machine parameters has limits to its accuracy. Furthermore, it highlights the need for a model which incorporates direct dose information in the training process.

This work contributes to ongoing efforts toward fully autonomous, AI-driven planning pipelines capable of supporting modern, adaptive oncology workflows.

Zusammenfassung

Strahlentherapie ist eine wichtige Behandlungsmethode für verschiedene Krebsarten, insbesondere für Prostatakrebs, eine der weltweit am häufigsten diagnostizierten Krebserkrankungen bei Männern. Die steigende Inzidenz sowie der wachsende Bedarf an individualisierten Therapieansätzen unterstreichen die Notwendigkeit effizienterer und flexiblerer Planungsprozesse. Derzeit basiert die Therapieplanung auf der Berechnung eines Parametersatzes, der üblicherweise durch expertenbasierte, iterative Optimierungsverfahren in kommerziellen Planungssystemen (TPS) erfolgt. Diese Verfahren sind jedoch zeitaufwändig und unflexibel, und somit nur bedingt für moderne, online-adaptive Radiotherapie geeignet.

In dieser Arbeit wurden zwei neue Ansätze implementiert und evaluiert, um ein bereits bestehendes DL-basiertes Modell zur Generierung fertiger Pläne für die Behandlung von Prostatakrebs zu verbessern. Das Modell sagt hierbei die erforderlichen Maschinenparameter für einen Linearbeschleuniger (LINAC), d.h. die Positionen der Kollimatoren (MLC) und die Monitor-Einheiten (MU), direkt auf Basis von Computertomografie (CT)-Bildern, Strukturdaten und einer vorgegebenen Dosisverteilung voraus. Aufbauend auf Erkenntnissen aus früheren Arbeiten [1], [2] wurden in dieser Arbeit mehrere Erweiterungen vorgenommen, um die Generalisierbarkeit des Modells sowie die Qualität der vorhergesagten Dosisverteilungen zu verbessern. Erstens wurde ein object detection-basiertes Modell implementiert und das Netzwerk entsprechend angepasst. Diese Modifikation ersetzte die zuvor verwendete Verlustfunktion, die lediglich auf der Abweichung zwischen vorhergesagten und Referenzmaschinenparametern basierte, durch eine räumlich informierte Verlustfunktion. Zweitens wurde das Netzwerk dahingehend erweitert, dass es mit einem 3D-Ausschnitt der 3D-Repräsentation der Patientenanatomie arbeiten kann. Darüber hinaus wurde ein Dosisberechnungsalgorithmus implementiert, um die Qualität der vorhergesagten Dosisverteilungen zu validieren. Der ursprüngliche Datensatz bestand aus 117 Patienten, welcher durch Datenaugmentierung dieser auf 310 Trainings- und 104 Validierungsdatensätze erweitert werden konnte. Die Leistung der modifizierten Modelle wurde anschließend durch Analyse der Trainingsverläufe und dosimetrische Tests an Plänen, welche für 34 Testpatienten generiert wurden, bewertet.

Die implementierten Modifikationen führten letztlich nicht zu einer Verbesserung der Modellleistung. Die Gamma-Pass-Rate (GPR) der modifizierten Modelle sank im Vergleich zum Basismodell um etwa 10 Prozentpunkte. Während alle Modelle dieselbe mittlere Dosis wie die Originalpläne im Zielvolumen der Prostata (PTV) erreichten, nahm das Volumen, das 95 % der verschriebenen Dosis erhält, um 5 % bzw. 23 % bei den beiden neuen Modellen ab – ein Hinweis auf eine reduzierte Zielvolumenabdeckung. Besonders der object detection-basierte Ansatz zeigt jedoch Potenzial, insbesondere in einer möglichen Kombination mit weiterentwickelten Methoden und einer spezifisch angepassten Modellarchitektur. Darüber hinaus wurde beobachtet, dass das Modell bereits früh im Training keine Verbesserung der Dosisverteilung mehr erzielte und stattdessen

hauptsächlich darin trainiert wurde, die ursprünglichen Maschinenparameter besser zu imitieren. Dies verdeutlicht die Limitierungen des aktuell verwendeten geometrischen Trainingsansatzes, bei dem direkt auf Maschinenparameter optimiert wird, und unterstreicht die Notwendigkeit, die Dosisverteilung selbst in den Trainingsprozess zu integrieren.

Diese Arbeit leistet einen Beitrag zu laufenden Forschungsbemühungen, vollautomatisierte, KI-gestützte Planungssysteme zu entwickeln, die moderne adaptive Arbeitsabläufe in der Radioonkologie nachhaltig unterstützen können.

Danksagung

Zunächst möchte ich mich bei meinen Betreuern, Peter Kueß und Gerd Heilemann, für ihre wertvolle Unterstützung während dieser Arbeit bedanken. Sie hatten immer ein offenes Ohr für meine Fragen, ob fachlich oder organisatorisch. Besonders dankbar bin ich auch für ihre Flexibilität und dafür, dass sie sich kurz vor der Abgabefrist trotz voller Terminkalender meine Entwürfe immer sofort angesehen haben. Weiters möchte ich mich noch bei Lukas Zimmermann bedanken, der mir besonders bei der Arbeit mit der Objekterkennung viele hilfreiche Tipps und Ratschläge gegeben hat.

Ich bedanke mich auch beim Prof. Joachim Widder und Prof. Dietmar Georg für die Möglichkeit meine Diplomarbeit an der Universitätsklinik für Radioonkologie schreiben zu dürfen.

Ganz besonders möchte ich mich auch bei meiner Freundin Carola und meiner Familie bedanken. Ohne sie wäre diese Arbeit nicht möglich gewesen. Besonders ihre Fähigkeit, mich auch in belastenden Phasen, wenn ich nur die Arbeit im Kopf hatte, abzulenken, hat mir die nötige Motivation um durchzuhalten gegeben.

Contents

Abstract	i
Zusammenfassung	ii
1 Introduction	1
2 Physical, Technological and Medical Background	4
2.1 Radiotherapy	4
2.1.1 Modalities	4
2.1.2 Dose Deposition	5
2.1.3 Dose Calculation	7
2.2 Medical Background	10
2.2.1 Standard RT Workflow for prostate cancer patients	10
2.2.2 Treatment Planning	11
2.2.3 Delivery	12
2.3 Deep Learning and Neural Networks	15
2.3.1 Convolutional neural networks	15
2.3.2 Encoder-Decoder networks	17
2.3.3 Residual networks	17
2.3.4 U-Net	18
2.3.5 Object detection	18
2.3.6 Transformers and attention	20
2.4 AI in Treatment Planning	21
2.5 Purpose of this thesis	22
3 Material and Methods	24
3.1 Patient Data	24
3.2 Preprocessing	26
3.2.1 Augmentation Strategies	27
3.2.2 Limitations	28
3.3 Network Architecture	29
3.3.1 Dataloading	30
3.3.2 Encoder	31
3.3.3 Latent space manipulations	32
3.3.4 Decoder	33
3.3.5 Loss function	33
3.4 Experimental Setup	34
3.4.1 Network training and validation	35
3.4.2 Dosimetric evaluation	36
3.4.3 Statistical Tests	38
3.5 Soft- and Hardware	39

4	Results	40
4.1	Training and validation results	40
4.2	Geometric evaluation	43
4.3	Dosimetric evaluation	48
4.3.1	Gamma Pass Rate	48
4.3.2	Dose-Volume Histograms	49
5	Discussion	62
5.1	Interpretation of results	62
5.2	Comparison with state of the art	63
5.3	Future challenges and improvements	65
6	Conclusion and outlook	67
	References	74

1 Introduction

Radiation therapy has been one of the cornerstones of modern oncology and the effective treatment of cancer for decades. More particularly, radiotherapy is essential for the treatment of prostate cancer, which is one of the most common forms of cancer among men globally. According to recent estimates, over 1.4 million new cases of prostate cancer were diagnosed worldwide in 2022, a number projected to rise to nearly 2.9 million by 2040 [3], [4]. Especially in many Western nations, the demographic shift towards an aging population leads to an increase in patient numbers and the associated health-care demands. This makes prostate cancer a serious public health challenge and further stresses the need for more efficient therapy workflows. The increasing focus on individualized treatment planning necessitates the development of new techniques to ensure the radiotherapy sector remains capable of addressing future clinical challenges.

The modern treatment planning process is based on an optimal dose distribution, created by a team of experts from different medical backgrounds, such as radiation oncologists and medical physicists, that takes into account the patient's specific anatomy as well as any tumor characteristics. This dose distribution is usually created with the help of a Treatment Planning System (TPS), a specialized software solution that allows for iterative optimization, dose comparison and treatment simulation. Treatment planning further includes extensive Quality Assurance (QA) to ensure safety and effectiveness. Despite its obvious merits in terms of accuracy and clinical reliability, the conventional approach is too slow and inflexible. As a result, it is unsuitable for online adaptive radiotherapy as it requires specialized expertise and binds resources and personnel that could be allocated more efficiently. Furthermore, once the initial plan has been created for a patient, it is usually kept unchanged for the entirety of the treatment process, notwithstanding the fact that the patients anatomy might slightly change in between sessions. For example in the case of prostate cancer, a full bladder behaves differently than an empty one. This also does not take tumor regression over time into account.

To address these limitations, recent research has focused on utilizing the immense potential of Artificial Intelligence (AI), mostly in the form of special types of Deep Learning (DL) models, and integrating these into the specific part of the treatment planning workflow. Especially Deep Neural Networks (DNNs) have shown promise in terms of automating and accelerating different tasks [5], such as rapid generation of viable dose distributions as well as automatic segmentation. These improvements not only allow for massive reductions in time and labor spent on certain tasks, but also allow clinic personnel to focus more on the specific challenges each patient poses and to adjust the treatment strategy dynamically to anatomical and biological changes during therapy.

This thesis investigates the possibility of using DL to automate treatment planning, with a focus on prostate cancer. The work builds on a previously established proof-of-concept developed in-house at the Department of Radiation Oncology (AKH Vienna, MedUni

Vienna), which demonstrated the feasibility of using specialized Neural Network (NN) to accurately generate Volumetric Modulated Arc Therapy (VMAT) plans [1]. Unlike other proposed models [6], [7], [8], the one presented here focuses on directly predicting machine parameters which can be used for delivery with minimal additional work and massively speed up the actual treatment plan generation. While the initial models showed great promise, there still remained several critical limitations, particularly in terms of data diversity and generalizability. Furthermore, there was no actual physically informed feedback in the model as its loss function was based on the difference between the original and predicted machine parameters, which is only a surrogate for the actual dosimetric impact. Additionally, more advanced DL techniques like object detection have since become available and offer new possibilities [9].

The basis of every AI model is high-quality data. This is also incidentally one of the most serious challenges for the effective usage of DL in radiotherapy. Although, as mentioned, prostate cancer is highly prevalent, due to data protection laws, patient data cannot be shared without dedicated approvals, and data on treatments is usually kept in individual institutions and not shared on a large scale. Each clinic might have different clinical protocols and goals, for example dose thresholds in specific regions, for specific treatment challenges and use different TPS to actually create the plans. Furthermore, differences in terms of hardware vendors, i.e. the manufacturer of the Linear Accelerator (LINAC) used, as well as planning workflows, for example what kind of imaging and how it is used, lead to vast, heterogeneous datasets. This hinders the efficient compilation of large, standardized datasets, which are nonetheless essential to build and train generalized and effective Machine Learning (ML) models.

The original model on which this thesis is based also suffered from these problems. Currently, it is limited to a specific combination of software (TPS) and hardware (LINAC) relies on training data derived from a narrowly defined clinical cohort. Thus, it can only be applied for creating plans in a select set of scenarios. This lack of generalizability impedes its performance on unfamiliar datasets and limits its usefulness for a wide range of clinical applications.

To address these issues, a number of new avenues for improving the current model and to increase physical accuracy and adaptability are explored in this work. The first major innovation is the integration of object detection to explicitly identify and localize relevant anatomical structures within the input data. This enhanced spatial awareness is expected to improve the model's ability to generalize to patients with varying anatomies and tumor characteristics. The second modification is the inclusion of a pseudo-dose engine to create a physically informed model that can better quantify the impact of the machine parameters on the actual physical dose by allowing for a quick approximation of the dose defined by the currently predicted machine parameters.

Furthermore, this work also investigates the possibility of generating synthetic training

data by creating variations of existing clinical plans through changes in the anatomy and by applying different optimization targets to create new plans from the original ones, thus emulating data from different treatment centers. These changes to data preprocessing not only aim to increase the available data pool and target the issue of data scarcity, but also offer the possibility to enhance the model performance by increasing generalization and reducing the risk of overfitting.

Ultimately, this thesis contributes to the broader goal of making radiotherapy planning more accessible, efficient, and adaptable. A generalized, soft- and hardware vendor-independent DL model has the potential to be deployed across a wide range of clinical environments and situations, including those with limited finances or expertise in advanced treatment planning. This is especially relevant in low- and middle-income countries where the population and thus number of patients is growing fast, but there is no longstanding tradition of radiotherapy and little financial resources. By significantly reducing the time and complexity of plan generation as well as the reliance on specialized know-how, such a system could enable faster patient throughput and allow access to high quality treatment planning using models trained on global expert institutes. Furthermore, it could facilitate personalized treatment adaptation and improve equity in cancer care delivery.

The remainder of this thesis is structured as follows:

- In Section 2 clinical and technical background on prostate cancer treatment and radiotherapy in general is provided. Furthermore, key concepts of deep learning are introduced and a short summary of the relevant physics behind dose calculation is given. Lastly, the current state of AI in radiotherapy is presented.
- In Section 3 an overview of data used and the state of the original model is given. Alongside this, the changes that were made to it as well as how they were evaluated is explained.
- Following this, in Section 4 the results of the numerical experiments are presented.
- In Section 5 the aforementioned results are interpreted and analyzed with respect to the research objectives.
- Finally, Section 6 gives a final summary and outlook on future goals.

The aim for this work is to bring together the clinical practice in radiation oncology and modern AI research to get closer to truly personalized and online radiotherapy.

2 Physical, Technological and Medical Background

The following section will give an overview about the physical and technological background behind this thesis. In Section 2.1 information about the modalities of and physical theory behind radiotherapy with a focus on prostate cancer treatment is presented. Following this, a short introduction to the medical practice for prostate cancer treatment is given. In Section 2.3 all relevant theory for the software part behind the project is presented, while Section 2.4 contains a short overview of the current use of artificial intelligence and machine learning in radiotherapy treatment planning. Lastly, in Section 2.5 the scientific goal of this thesis is laid out in more detail.

2.1 Radiotherapy

Radiotherapy is one of the primary methods for treating various forms of cancer. Among the different types of radiation therapy, External Beam Radiation Therapy (EBRT) is the most commonly employed approach and also the focus of this work [10]. Other common variants include Brachytherapy, which is a form of internal radiation therapy whereby radioactive materials or temporary catheters are placed near the affected tissue to treat the cancer. These seeds emit radiation over a longer period of time and remain harmless after the material degrades, while the catheters are removed after treatment [11].

2.1.1 Modalities

The fundamental principle utilized by EBRT is based on the targeted delivery of high-energy photons or particles to malignant tissues. This process in turn induces molecular and atomic interactions in the tissue, which ultimately damage cancer cells and inhibit tumor development [12], [13].

More specifically, the therapeutic effect of radiotherapy using photons is primarily caused by ionization and excitation processes of the atoms and molecules in the tumor, when energy is deposited by the photons. These processes either directly damage the cells in the case of ionization or indirectly through the emission of secondary radiation from excitation. The most critical biological target for radiation-induced damage is the DNA as the photons can cause double-strand breaks and other damage within the DNA helix. This impairs cellular replication and triggers apoptosis (programmed cell death) within the cancerous region [13].

Typically, photon beams in the energy range of 6 to 18 MeV are used for EBRT to achieve a balance between optimal tissue penetration and dose deposition, and sparing the healthy cells around the tumor region. The depth-dose distribution for MV-photons, among other commonly employed beam types, can be seen in Figure 2.1. The biggest energy-dependent variation in percentage depth dose (PDD) curves is observed in the buildup region. This buildup is due to secondary electron scattering. Higher photon

energies exhibit a more extended and pronounced buildup due to increased secondary electron range. To achieve the balance between optimal tissue penetration and dose deposition, different imaging modalities, including Computed Tomography (CT), Positron Emission Tomography (PET) and Magnetic Resonance Imaging (MRI), are employed beforehand to identify regions of interest, such as critical organs and the tumor region. Furthermore, there are different techniques to actually deliver the dose to the patient. The main ones are VMAT and Intensity Modulated Radiation Therapy (IMRT) [13], which will be explained in more detail in Section 2.2.3.




Figure 2.1: Percentage depth-dose curves for different beams: ^{60}Co -beam, 6 MeV, 10 MeV and 18 MeV x-ray beams [14].

2.1.2 Dose Deposition

In radiotherapy physics, knowledge about the absorbed dose is crucial for effectively creating treatment plans. In general, the dose is used as a quantitative connection between the purely physical radiation and the actual biological effects in the irradiated tissue. Principally, dose is defined as the amount of energy deposited by ionizing radiation per unit mass in a given material, i.e. biological tissue in the radiotherapy. Mathematically, it can be defined as:

$$D = \frac{dE}{dm} = \frac{1}{\rho} \frac{dE}{dV} \quad (2.1)$$

The most commonly used unit is Gray (Gy), where 1 Gy is defined as the absorption of one Joule of energy per kilogram of matter.

Another relevant concept is Kinetic Energy Released per unit MASS (KERMA), which is defined as the sum of the kinetic energies of the charged particles created by the uncharged ionizing radiation per unit mass at a given point. Mathematically, this leads to:

$$K = \frac{dE_{tr}}{dm} \quad (2.2)$$

where dE_{tr} is the sum of the energies and dm is the mass. KERMA can in general not be equated with the absorbed dose. While both quantities are approximately equal in regions of Charged Particle Equilibrium (CPE), they differ in the initial buildup region, where one may find a high KERMA but a lower dose, due to more electrons leaving this region than entering. CPE is said to exist, if in the observed volume the charged particles leaving the volume, are exactly replaced by identical particles of the same energy. Both quantities are usually measured in Gy.

As mentioned above, photons traveling through matter deposit energy. This happens via photons indirectly ionizing matter via the release of charged particles, usually electrons. These then interact with and transfer energy to orbital electrons of atoms through Coulomb interactions. This deposited energy can in turn be related to a specific dose.

The diminishing intensity of a photon beam as it travels through a material can be modeled by the exponential attenuation law:

$$I(x) = I_0 \exp^{-\mu(E)x}, \quad (2.3)$$

where I_0 is the initial intensity at the entry point, μ is the linear attenuation coefficient, which is both energy dependent and different depending on the material, and x is the thickness of the material. There are different interactions between the incident photon beam and the traversed matter, at the relevant energy scales the photo-electric effect, Compton scattering and pair-production are the most important ones [15].

The photo-electric effect is a phenomenon wherein a bound orbital electron interacts with an incident photon, absorbing it and then being emitted with a kinetic energy equaling the difference between the incoming photon energy and the electron binding energy. The probability of the photo-electric effect occurring is inversely proportional to its energy cubed [15].

Compton (incoherent) scattering occurs when a photon interacts with a loosely bound electron inhabiting the outer shell of an atom. In this case, the photons transfers a portion of its energy, dependent on the scattering angle, to the electron and changes its

direction [15].

Electron-positron pair production is an effect in which a photon near an atomic nucleus converts its energy to particle mass and creates an electron and a positron with kinetic energy equal to the difference between the photons energy and twice the rest energy of the electron, i.e. $E = 2 \cdot 511$ keV. The electron and positron then give off their energy via ionization, excitation and Bremsstrahlung. The positron will then annihilate with an electron emitting two photons [15].

At lower energies in the keV-range the photo-electric effect dominates, in the intermediate range the Compton effects dominates, as it depends only slightly on the photon energy, and starting in the MeV-range pair production becomes most dominant, as only in this range the photons have the required energy to produce an electron-positron pair. The effects further depend on the atomic number Z , with the dependency for the photo-electric effect being about $\propto Z^4$ and for pair production $\propto Z^2$, while Compton scattering is only weakly dependent on Z [15].

2.1.3 Dose Calculation

Accurate dose calculations are a fundamental requirement in treatment planning, to ensure adequate coverage of the tumor region while minimizing exposure of healthy tissues nearby. However, the calculation of the dose distribution within a patient is a non-trivial problem, as there are different contributions to be considered. The most important components of the dose are highlighted in Figure 2.2.

For the actual computation, a number of different algorithms have been proposed over the years, such as:

- Empirical models, which rely on measured dose distributions and empirical correction factors, e.g. Clarkson integration
- Convolution-superposition models, which explicitly account for both primary and scattered radiation using kernels, e.g. the Pencil Beam (PB) model and the Point Beam kernel/Collapsed Cone Convolution (CCC)
- Monte Carlo (MC) methods, which use stochastic simulations to track individual photon and electron interactions

The actual categorization might differ as the distinction for some models is not always clear cut [17].

While MC simulations offer the highest accuracy, they are also much more time consuming and expensive as simpler, kernel-based models. The latter are based on utilizing a kernel, which in effect is a function describing the spatial distribution of the energy deposition within the patient. The dose at a position in the patient from the point kernel

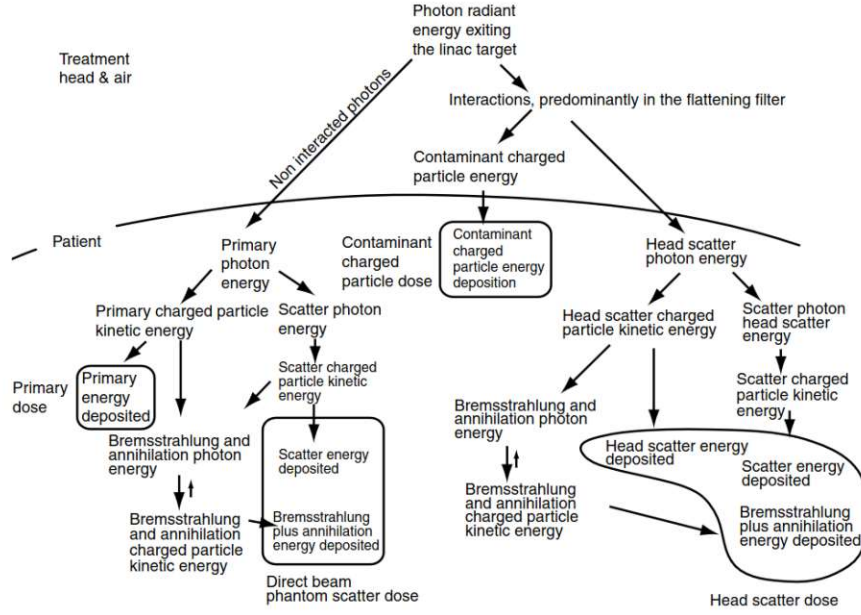


Figure 2.2: A schematic overview of the most important interactions and dose components for photon beams, the four most important are highlighted [16].

model can be calculated by integrating contributions from all discrete interaction sites within the irradiated volume. This is particularly useful for handling inhomogeneities, but requires significant computational effort compared to simpler methods like the pencil-beam model [18].

Faster and only slightly less accurate, the CCC algorithm takes advantage of the fact that the dose contribution from distant voxel decreases rapidly with distance. The dose is calculated by accounting only for voxel in certain directions instead of all voxels, with neighboring voxel being more likely to lie in one of those directions. To ensure energy conservation, i.e. account for the energy of unsampled voxels, the energy from all voxels is collapsed in to the average direction of each sampling direction. This method can thus also account for inhomogeneities in the material depending on direction [18].

A reasonable compromise between speed and accuracy is the PB kernel for which the dose at a point $P(x, y, z)$ in the patient can be given mathematically by the following convolution integral:

$$D(x, y, z) = \iint_{\text{field}} \Psi_E(x', y') K_{PB}(x - x', y - y', z) dx' dy', \quad (2.4)$$

where Ψ_E is the incident energy fluence at the beam entry and K_{PB} is the PB kernel, which includes contributions from the primary photon beam, secondary electrons and scattered photons. This corresponds in essence to a convolution of the PB kernel and the

field intensity distribution on the incident layer. The full photon beam in this scenario is represented as being made up of many smaller beamlets (pencil beams), which when added together can be used to compute the dose from the full beam. Using this model one can explicitly account for beam intensity variations, for example due to wedges and collimators and dose due to the lateral spread of secondary electrons and scattered photons. The kernels and the fluence can be obtained from MC simulations and polyenergetic beams can be represented by adding monoenergetic beams [18].

A graphic representation of the principle behind the PB model is given in Figure 2.3.




Figure 2.3: A graphic representation of the principle behind the PB model. The dose at a point P is calculated from the contributions of the individual pencil beams [18].

In practice, the convolution is only performed at five different depths and intermediate values are interpolated. Furthermore, there are a number of corrections which might be employed to increase the accuracy of the PB model. For example, factors for the obliqueness of the skin at the beam entry and for different inhomogeneities can be introduced [18].

A reasonable analytic representation of the PB kernel in cylindrical coordinates is given by:

$$K_{PB}(r, z) = \frac{A_z e^{-a_z r} + B_z e^{-b_z r}}{r}, \quad (2.5)$$

where r is the radius from the main pencil-beam axis and the fitting coefficients, which depend on the depth z , can be obtained for example through MC simulations [18].

Advantages of the PB model are for example the superior speed compared to more expensive methods such as MC simulations, while maintaining reasonable accuracy for specific scenarios. Furthermore, in homogeneous tissue this model works quite accurately and the calculation of the convolution can be sped up by using Fourier transforms. There are however also a number of limitations, such as decreased accuracy in inhomogeneous media, such as the lungs or bones and near interfaces, e.g. lung boundaries. Also, higher order scattering is not fully accounted for with the PB model [18].

2.2 Medical Background

Prostate cancer ranks as the second most frequently diagnosed cancer among men worldwide. Furthermore, it is also the fifth leading cause of cancer-related mortality for men [19]. Epidemiological data from recent years shows substantial geographical variance. Incidence rates per 100,000 men varied ranged from 7.6 to 168. Mortality rates also exhibited large differences, with values ranging from 2.9 per 100,000 men and to more than 24.0 in some countries [19].

Details on how a radiotherapy treatment plan is created and ultimately delivered are given in Section 2.2.1, Section 2.2.2 and Section 2.2.3.

The treatment of prostate cancer is uniquely suited for the development of a deep learning model to assist in treatment planning. The high occurrence rate of prostate carcinomas provides more training and testing data than most other cancer types. Furthermore, the relatively simple geometry of the prostate compared to the lungs for example, should allow the models to more efficiently learn the relations between the applied dose and the treatment plan parameters.

2.2.1 Standard RT Workflow for prostate cancer patients

Before a patient can be treated via radiotherapy, thorough diagnostic and imaging measures have to be taken. In the following, a short summary of the standard workflow for radiotherapy at the Department of Radiation Oncology (AKH Vienna, MedUni Vienna) is given [20].

These include for example taking into account the patients previous medical history and histological examinations. From these the Gleason score and WHO grading can be es-

established, both markers to assess the tumor [20]. Furthermore, the Prostate-Specific Antigen (PSA) level is measured both at the initial diagnosis as well as right before the start of the treatment. The main imaging measures taken are CT, Transrectal Ultrasound (TRUS) and MRI. The CT used for planning is performed with a full bladder and using a so called rectal balloon, which helps to keep the prostate position stationary and reproducible and is also used during treatment [20]. The TRUS helps to further establish the target volume, while the advantages of the MRI are an assessment of the definition of the prostate apex and extent of the tumor. Depending on indication, a cystoscopy (examination of the bladder), a rectoscopy (examination of the rectum) or a bone scintigraphy might be performed as well [20].

For the actual radiotherapy treatment several different approaches might be taken, such as EBRT, brachytherapy, a combination of both or hormonal therapy. The specific method depends among other indicators on the stadium of the cancer, the Gleason score and PSA levels [20]. Furthermore, the age and health of the patient are taken into account. Lastly, the so-called Roach formula, which depends on the Gleason score and PSA level, is used to establish a risk of the cancer spreading to nearby lymph nodes. This for example means for low risk patients EBRT or seeds, without additional hormonal therapy, are used. In contrast to this, for high risk patients, a combination of EBRT and brachytherapy or of EBRT and hormonal therapy. Additionally, for EBRT patients gold fiducial markers are placed inside the prostate to help guide and optimize tumor visualization and dosimetry [20].

Depending on the risk, different volumina might be chosen as the target. For lower risk patients only the prostate and maybe the basis of the seminal vesicles might be irradiated, while for higher risk patients the prostate and seminal vesicles as well as the surrounding lymphnodes are chosen. Organs that might be at risk during the irradiation include for example the bladder, the rectum and the femoral heads [20].

The treatment planning happens in specially designed software and is discussed in more detail in Section 2.2.2.

2.2.2 Treatment Planning

A TPS is a special type of software used in radiotherapy. They play an important role in modern radiation oncology by standardizing and simplifying the treatment planning workflow, and allowing for a more personalized treatment of patients. Usually, these software platforms utilize different kinds of imaging data, such as CT, MRI, and PET scans, to generate a 3D representation of the patient's anatomy. This representation can then be used to identify critical structures and delineate the tumor regions, and also to define clinical goals for plans, such as dose thresholds in sensitive organs and in other surrounding healthy tissue. As mentioned in Section 2.1.3, TPS use different kinds of dose calculation algorithms to compute the dose distribution and optimize it based on the predefined clinical goals. This dose distribution is then translated into machine usable

parameters to be applied with the LINAC.

TPS also allow the simulation of different delivery techniques, such as IMRT or VMAT, mentioned above. Furthermore, TPS also enable a variety of QA procedures and other treatment evaluations by providing a Dose Volume Histogram (DVH), direct comparisons of plans and other checks. Modern TPS might even include applications utilizing artificial intelligence, which can further speed up and improve the treatment planning process. Advanced applications of TPS include adaptive radiotherapy, which adjusts treatment plans based on changes in tumor size, position, or patient anatomy during the overall course of treatment and also online monitoring of the patient to adapt to the real-time changes in anatomy and position while the patient is being irradiated.

However, TPS are usually rather complex and expensive, and still require a lot of human expertise to create the best possible plans. A sketch of the treatment planning workflow can be seen in Figure 2.4.

Examples of TPS include the RaySearch software by RaySearch Laboratories [21] and the Monaco software from ELEKTA [22]. At the Department of Radiation Oncology (AKH Vienna, MedUni Vienna) the former software is used and was thus also the basis for the statistical analysis of the model.

2.2.3 Delivery

The previous sections focused on the steps preceding the actual therapy of the patient and the theory behind it. In the following, a short summary of how the dose is actually delivered to the tumor region is given.

IMRT is a beam delivery technique based on modulating the intensity of the radiation beams, which offers high precision and conformity. The dose is usually delivered from a small selection of angles. This makes it especially effective for tumors which are close to critical organs, such as breast, lung, brain, gynecologic, and gastrointestinal cancers. However, it is also comparatively cost intensive, both in terms of time and computing power, and requires extensive dosimetry and treatment QA. Furthermore, it is prone to geometric errors, which necessitate a strong emphasis on the patient position, especially regarding the day-to-day changes in some region volumes, e.g. the bladder might be more or less filled. Advanced imaging techniques such as MRI and PET, and single-photon emission computed tomography (SPECT) have further increased the effectiveness of IMRT [13].

Compared to this, VMAT, which is a variant of IMRT, allows for a higher treatment efficiency by delivering the dose through a continuous 360° rotation of the beam around the patient. During this rotation, the beam intensity and the shape can be modulated by Multi-Leaf Collimators (MLCs). This leads to an improved sparing of Organ At Risk

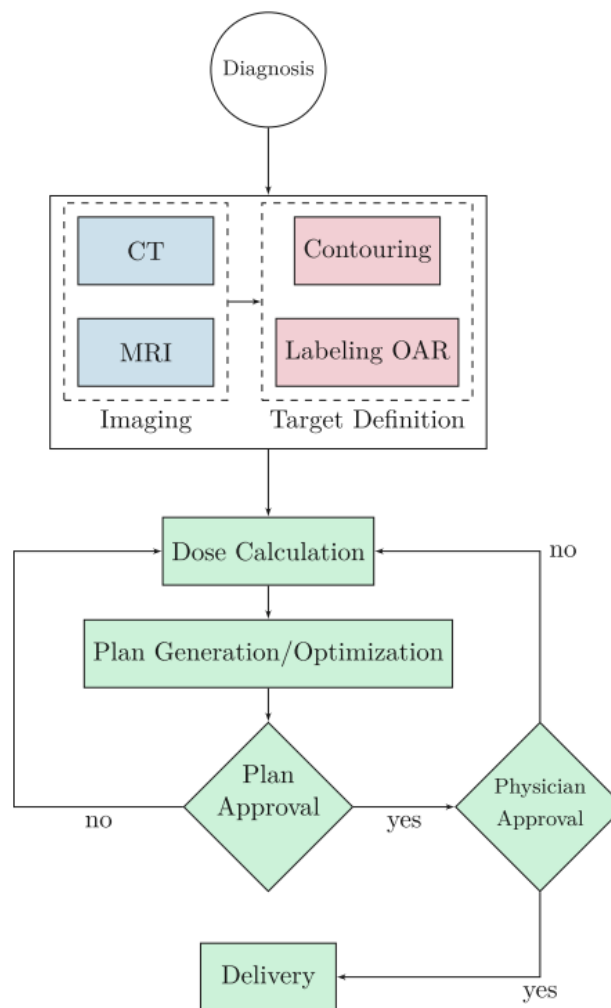


Figure 2.4: Treatment planning workflow: the patient is diagnosed and after creating a 3D representation of the patient anatomy using Imaging, the relevant regions are identified. The dose is then calculated and plans are created iteratively until it satisfies predefined clinical goals. After QA by a physician the patient can then be treated using this treatment plan [23].

(OAR) and a more homogeneous dose compared to conventional IMRT. However, the higher low-dose radiation exposure in surrounding tissue leads to an increased chance of secondary malignancies and thus necessitates a long-term follow-up of patients. Similarly to IMRT, VMAT has been used for treatment planning of prostate, thoracic, gastrointestinal, gynecological, head and neck, brain, and breast tumors [13]. At the Department of Radiation Oncology (AKH Vienna, MedUni Vienna) VMAT is primarily employed for treatment and all plans used for this project were of this type.

The beam emitted by the LINAC is shaped by collimators. While the exact make up and positioning of these collimators might vary between machines, for all plans used in this thesis the machine specifications included 80 vertically stacked leaf collimator pairs, which modulate the beam by moving horizontally, and two jaw collimators which move along the vertical axis. Furthermore, the intensity of the beam is measured using so called Monitor Unit (MU). How this looks in a typical TPS can be seen in Figure 2.5.




Figure 2.5: The "Beams-Eye-View" in RayStation, showing how the vertically moving jaws and the horizontally moving leaves are modulated to irradiate the target area. The area left open by the jaws is given in orange and the affected leaves are shown in orange with their opening being the outlet for the beam during treatment. The differently colored outlines of relevant regions are used to define clinical goals, for example how much average dose should be in the prostate region in light pink [Screenshot from RayStation].

Independent of what specific technique is used to modulate the beam, there are also further concepts that improve the therapeutic outcome of radiotherapy.

An essential aspect regarding the delivery of the actual dose is the concept of fractionation. Rather than delivering the entire radiation dose in a single session, the treatment is separated into multiple smaller doses administered over a series of sessions, typically with a day between each treatment. Fractionation exploits the differential repair capacities of normal and cancerous tissues. Normal cells generally possess a superior ability to repair sublethal DNA damage compared to cancer cells. By allowing time between treatments, healthy tissues can recover, while cancer cells, which repair more slowly, accumulate irreparable damage [13].

This approach not only enhances the therapeutic ratio but also reduces the risk of side effects and the strain on patients. Standard fractionation schemes used in the Department of Radiation Oncology at the Medical University of Vienna deliver doses in the range of 1.8 to 2 Gy per fraction, but hypofractionated and hyperfractionated protocols, i.e. more dose over a shorter time frame and less dose over a longer time frame respectively, are also used depending on the clinical scenario. For example, hypofractionated treatments might include single doses of up to 3 Gy per fraction or even more [20], [24].

As mentioned before, the full target dose for a treatment is not applied during a single treatment session, but rather split over several days or even weeks to exploit the difference in radiation recovery between healthy and tumor cells. However, for more complex cases, such as cancer which has already spread to surrounding lymph nodes, it might be required to split the dose of a single session into several arcs. This method involves rotating the gantry, from which the beam is emitted, for more than one full rotation of 360° . This effectively doubles the degrees of freedom one has to achieve the clinical goals, as the actual radiation delivered by each beam is in principle independent, thus allowing for more angles to avoid excessive irradiating of OAR, while increasing treatment time. For localized prostate tumors however, a single arc is typically sufficient to achieve an effective dose distribution.

2.3 Deep Learning and Neural Networks

DL is a subset of machine learning, using neural networks with multiple processing layers to learn hierarchical structures in data, both supervised and unsupervised [25]. While deep learning models have been used in widely differing domains, such as speech recognition, object detection or natural language processing, in recent years they have also featured more prominently in the radiotherapy sector. Deep learning uses backpropagation to iteratively adjust the internal parameters of the model to refine feature representations across successive layers. There have been many different kinds of NN such as Convolutional Neural Network (CNN), Recurrent Neural Network (RNN) and other variants proposed, of which the most relevant for this thesis will be presented in the following. The ability of deep learning to find underlying information in large datasets has positioned it as useful tool in radiation oncology, where it is increasingly used to enhance treatment planning, adaptive radiotherapy, and outcome prediction [26].

2.3.1 Convolutional neural networks

CNNs are a special class of deep learning models, which are designed to work on structured data, such as images or dose distributions in the radiotherapy case. These networks use convolutional layers to detect structures and patterns in the images, such as edges, textures or gradients. They also involve other operations, such as pooling layers to reduce spatial dimension while preserving essential information or fully connected layers for final classifications. CNNs further employ parameter sharing across the full images,

which reduces the model size [27], [28].

There is a variety of applications for CNNs with the most important being [28]:

- image classification/recognition, i.e. identifying benign or malignant skin lesions
- object detection, i.e. detecting whether or not a special object is present or not in a given image
- segmentation, i.e. dividing images into different regions or segments, for example healthy and tumor tissue

An example for a pooling convolution operation can be seen in Figure 2.6.




Figure 2.6: Example for the computation of a pooling operation on a 5×5 input matrix, reducing the output to a 3×3 matrix [29].

Since the introduction of the CNN, many variants and modifications have been proposed. Some of these strategies, like deep architectures, residual connections, and batch normalization are also used in the project of this thesis and will be explained in more detail in later sections. However, there are still issues present, chief among them the need for large labeled datasets to allow for efficient training as otherwise overfitting might become a problem on smaller datasets [28], [27]. This requires extensive data augmentation and transfer learning to overcome training issues, especially in fields in which there is a lack of large uniform data such as radiotherapy.

2.3.2 Encoder-Decoder networks

Encoder-decoder networks are a special class of deep learning architectures designed for tasks that require transforming a given input data into a feature representation and then again reconstructing an output from that representation. This is achieved in a two-step process. First, the encoder extracts relevant feature information from the input data, reducing its spatial dimensions and transforming it in to the so called latent space. For this purpose, depending on the specific application, layers using convolutions or transformers might be employed. The data representation at the fully encoded bottleneck, where only the essential aspects are present, is then used by the decoder. This makes such networks useful for tasks such as image segmentation, machine translation, and anomaly detection [30], [27].

The decoder then takes this compressed representation and reconstructs it into the desired output format. It performs an inverse operation compared to the encoder, gradually transforming the latent space representation back into a structured output using upsampling, deconvolution, or attention mechanisms. The encoder-decoder framework is widely used in applications such as autoencoders for unsupervised learning, sequence-to-sequence models for language translation, and U-Net architectures for medical image segmentation [30].

A representation of such a framework can be seen in Figure 2.8, however in this specific case of a U-Net there are skip-connections between the down- and upsampling path, which would not be present in a conventional encoder-decoder network.

2.3.3 Residual networks

The Residual Network (ResNet) architecture, introduced in 2015, is a special NN architecture to allow for deeper neural networks, which before that often ran into the so-called degradation problem, i.e. additional layers leading to higher errors instead of improving performance. This is avoided in ResNets by using skip-connections between layers, which additionally to the adapted input also pass on the unchanged input to the next layer, which allows the network to learn residuals instead of the full mapping directly. This approach makes it easier to optimize deeper networks, as the model effectively learns modifications to an identity function rather than fitting a complex function from scratch [31], [28].

Advantages of the ResNet architecture include the possibility of effectively training deep neural networks, without having to worry about the vanishing gradient problem as the residual connections allow gradients to propagate more efficiently through the layers. This has the additional advantage of improving the generalization capabilities of the network, as deep networks tend to generalize better on unseen data compared to shallower ones. However, one has to consider the additional computational complexity. Deep versions of ResNet are computationally expensive, both regarding processing power and memory.

Furthermore, overfitting becomes an increasing problem when smaller datasets are used to train deep networks, which in turn requires data augmentation and regularization [31].

An example of such a residual learning block can be seen in Figure 2.7.




Figure 2.7: Sketch of the principle behind residual learning. Image taken from [31].

2.3.4 U-Net

Another class of deep learning networks, more specifically encoder-decoder networks, is the U-Net framework, so called because of its symmetric "U"-shaped design, which uses skip connections, similarly to the ResNet architecture mentioned above, to directly forward input by bypassing deeper layers from the encoder to the decoder. A schematic representation of the U-Net design structure can be seen in Figure 2.8. As before, the encoder extracts feature information from the input while downsampling, but now the shortcut connections help retain spatial information which would otherwise be lost during the reduction of spatial dimensions. This makes it especially effective for image segmentation [32], [28].

Even though the U-Net architecture has proven very useful for image segmentation, it also has drawbacks, such as being computationally expensive when working with high-resolution images. Furthermore, this framework necessitates extensive data augmentation to allow for good generalizations [32]. Nevertheless, it is a widely used architecture and has spawned several variants, such as 3D U-nets for volumetric data [33].

2.3.5 Object detection

Object detection is a specific computer vision task designed for localizing objects in imaging data and classifying it according to a given set of classes. In contrast to conventional image classification, in object detection the aim is to predict bounding boxes for objects within the image, which has found widespread use in areas from autonomous vehicles to medical imaging. First, CNNs are used to extract visual feature information, which is then used to identify potential objects in the image by surrounding them with a bounding box. The accuracy of these bounding boxes is determined by applying the

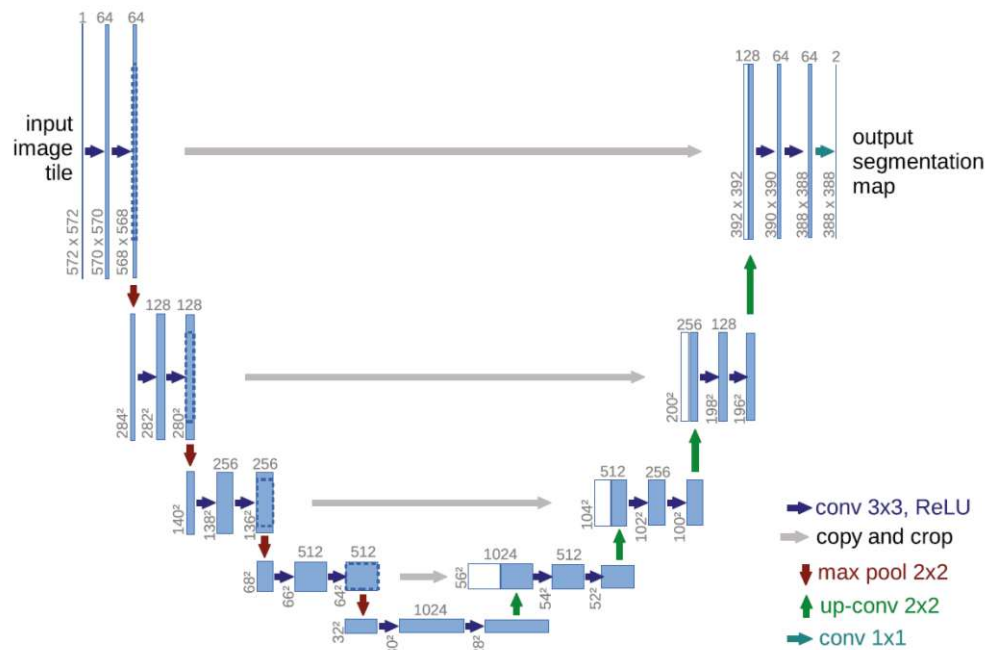


Figure 2.8: Image of the U-Net architecture: the encoder can be seen as the path on the left side which reduces the spatial dimensions down towards the bottleneck at the bottom and then upsamples this representation back to be used. Image taken from [32].

intersection-over-union metric. This is done by taking the intersection area of both the prediction and the ground truth and then dividing the area of the intersection by the area of the union of both boxes. These bounding boxes are then assigned a label for a class [28]. A variant of this, is the so called You Only Look Once (YOLO) framework, designed for very fast assessment of images.

The original YOLO network was put forward in 2015 and improved the speed at which images could be detected by employing a single neural network for both prediction of bounding boxes and class probabilities in one go [34]. This made it much faster than comparable models at the cost of reducing localization accuracy while maintaining a high generalization performance. Since then, numerous improved versions of the YOLO framework have been developed.

An image explaining how the bounding boxes and classification works can be seen in Figure 2.9.

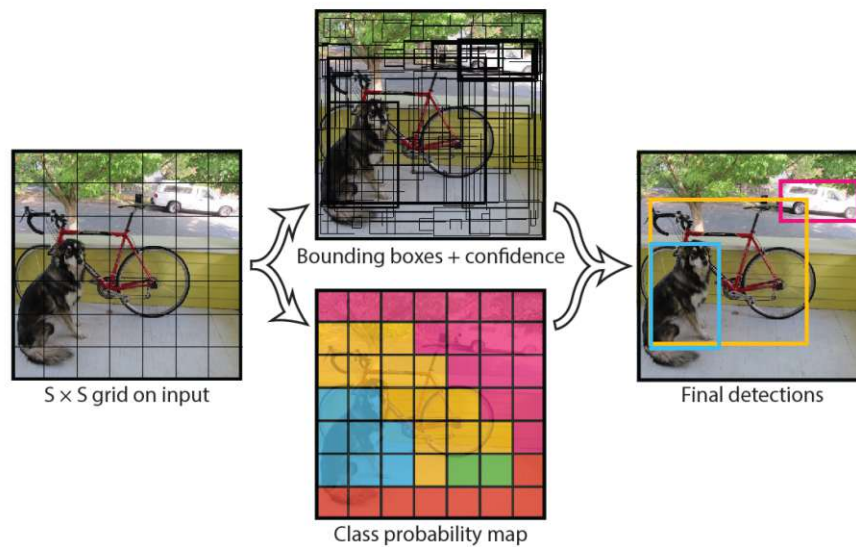


Figure 2.9: The image is first split into a grid on which bounding boxes and their respective confidence for containing an object of a given class are superimposed. These are then combined to detect the objects on the actual image [34].

2.3.6 Transformers and attention

Transformers are a special deep learning architecture, which were initially designed for Natural Language Processing (NLP), but have since been adapted for tasks in computer vision and speech recognition. The original model was introduced in 2017 [35] and developed the self-attention mechanism, which enables the network to attribute different weights to different inputs and having these weighting coefficients themselves be dependent on other input values, allowing for a better capturing of sequential data. It also introduced a much improved capability for parallelization, which proves immensely useful for large datasets. Instead of using only single attention, transformers employ multi-head attention, where each head independently learns different relationships among the data, after which these representations are connected together [28].

Initially, the input data has to be tokenized, e.g. subwords in language processing or image patches in computer vision, and then embedded in a higher dimensional vector space. Further, the tokens are given additional information about their ordering by positional encoding. This representation is then fed through an encoder-decoder structure and finally generates an output [28].

Advantages of Transformers include the much faster training compared to earlier RNNs, as here all tokens can be processed simultaneously. Furthermore, the self-attention mechanism allows the model to capture long-range relationship between inputs more effec-

tively. Transformers also display a high potential for transfer learning, i.e. transformers pre-trained on large, general datasets can be fine-tuned for special applications by then training on smaller, specialized training data. However, it has to be noted, that transformers and especially large-scale models require massive memory and computational resources. Furthermore, due to the inherent black box nature of neural networks, the way the networks comes up with a given output becomes very hard to interpret [28]. The success of the transformer architecture has led to them also being implemented in well-known models such as BERT or ChatGPT [36].

2.4 AI in Treatment Planning

As already mentioned before, there are several use cases in the field of radiation oncology which could benefit from AI, and more specifically machine learning in the form of deep learning. A number of frameworks have been proposed for plan optimization, however treatment planning remains a time-consuming process, which is very reliant on human experience and input. In the following, a short summary of various DL-based techniques for Automated Treatment Planning (ATP) will be given.

One possibility for the implementation of ML is in automated beam orientation selection, wherein a DL model predicts the optimal direction for the beam, reducing planning time [37]. There is also an approach using reinforcement learning [38]. Automated dose map prediction is another field in which CNN-based models are used to accurately predict 3D dose distributions. These approaches outperform traditional DVH-based predictions by incorporating spatial dose information, leading to improved clinical usability.

Another possible application for DL is automated fluence map generation and optimal dose prediction, where deep learning techniques streamline the conversion of predicted dose maps [39], [40]. These methods speedup the overall treatment workflow and further reduce the need for manual adjustments, decreasing the possibility of human error and dependence on human experience. ML models have also been used for segmentation purposes, where the model directly predicts the delineation of the relevant organ and tumor regions [41].

There are however still many challenges to be overcome in the use of AI in radiotherapy, such as the limited dataset availability, overfitting risks, and the need for better generalization across diverse clinical scenarios. There is also an issue with different tumor types and treatment techniques requiring different deep learning models, making generalized ATP solutions difficult to implement [26].

The model predictions also require rigorous quality assessment to ensure that the ML-generated plans are at least as effective and adherent to clinical standards as manually created ones. Legal and ethical considerations also have to be considered, especially the challenges posed by the "black-box" nature of DL, which makes interpretability and ac-

countability difficult. All of these issues and challenges further highlight the importance of multi-institutional collaborations and the development of standardized evaluation metrics [26].

The specific project this work builds on, starts by taking a pre-optimized dose distribution, creating a latent-space representation of it and then producing a DICOM Radiotherapy (RT) plan file by predicting the MLC positions and the beam intensity. The thus generated plan could then immediately be recalculated in a conventional TPS and then delivered by a LINAC. These RT plan files can be produced in a matter of seconds, which allows for rapid online creation and modification of treatment plans. The full project pipeline also includes a dose prediction model which could be used together and implemented into a larger workflow, further speeding up the creation of RT plans. A more thorough explanation of the model architecture and the build-up of the whole pipeline is given in Section 3.

2.5 Purpose of this thesis

This thesis is based on the project developed in [23], [42] and [1] which uses DL to enhance the speed, efficiency, and quality of radiotherapy treatment planning. The specific primary objectives are:

- Advancing the fully automated treatment planning pipeline established at the Department of Radiation Oncology, Medical University of Vienna.
- Developing DL-driven methodologies to facilitate real-time adaptive radiotherapy planning.

This work builds upon an existing encoder-decoder network and seeks to optimize its performance through the integration of advanced techniques, such as object detection, embedding mechanisms and pseudo-dose computations. The research will focus on refining the network to achieve more efficient and higher-quality treatment plan generation and evaluating its performance against established clinical standards.

A fundamental challenge lies in training the model to learn the underlying physical principles of treatment planning, particularly machine parameters, delivered dose, and resulting dose distribution. The dose distribution is of utmost importance, as it directly determines the quality of a treatment plan. To address this, this study will investigate the integration of a pseudo-dose calculation model into the network's loss function, allowing the model to incorporate these physical relationships more effectively. Subsequent statistical analyses by creating full treatment plans from the model output using conventional TPS will be performed to assess whether this approach enhances treatment planning performance.

Expected outcomes include improved planning efficiency, more consistent achieving of clinical goals, and broader accessibility to advanced oncological care. By supporting

the development of adaptive radiotherapy, this research contributes to the overarching objective of precision medicine in oncology.

3 Material and Methods

This thesis is based on the automatic treatment planning model, which was developed in [1]. The initial proof of concept was given by [23] and further refined by [42]. The study was approved by the institutional review board of the Medical University Vienna under ethics approval number 1373/2023. In theory, for the final clinical workflow a model for predicting a dose distribution purely from patient imaging data is part of the larger pipeline and could be used before the plan generator model. However this work focused exclusively on the generation of the treatment plans based on pre-optimized dose distributions.

In Section 3.1 an overview of the kind of data, that the project is based on, is given. This data has to be preprocessed before it can be used for the ML model, the process of which is described in Section 3.2. The specific architecture of the network and the introduced changes are detailed in Section 3.3. How these changes were tested and evaluated can be seen in Section 3.4. Finally, a quick overview of the used software, packages and hardware is given in Section 3.5.

3.1 Patient Data

For the training and testing of the DL model, treatment plans from 169 patients treated at the Department of Radiation Oncology (AKH Vienna, MedUni Vienna) were used. Earlier models used more heterogeneous data regarding the treatment fractions, total dose and modalities. The patient data for this thesis however, was exclusively post 2021, after a TPS was introduced. The new TPS, RayStation (RaySearch AB, Stockholm, Sweden) facilitated more standardization with scripted templates for plan generation and optimization. Furthermore, all patients were treated using 20 fractions with a total dose of 60 Gy using the VMAT technique and treated using an ELEKTA LINAC. All personal information was anonymized before exporting from the clinical database.

These 169 patients were then split into a training and validation set of 129 patients and a test set of 40 patients. To increase the actual number of training data, augmentation techniques were used. How these augmentations were applied is specified in Section 3.2.1.

All patient data relevant to this project is stored for use in treatment planning using the so called Digital Imaging and Communications in Medicine (DICOM) standard. This is both a standardized file format and a network communication protocol in worldwide use, for the handling, storing, and sharing of data related to medical imaging. It further contains related information, including for example personal information of patients, such as age and gender, but also the specifics of treatment plans, such as the total dose prescription. This format ensures that the medical information necessary for a patients treatment is consistently formatted and can be used and transmitted across institutions using different devices and systems, which is especially useful today, as patients might be treated for the same or different reasons at different clinics.

DICOM data typically includes a meta information header and the actual data. The former stores general information, such as how to decode the data, and is not directly relevant for the purpose of this thesis. The data itself is stored as a series of elements, identified by a tag and a certain data type, as well as the length and actual value representation of the data. The DICOM standard also includes radiotherapy specific objects, as defined in [43], to help facilitate the transferring of data between departments. These objects include those for imaging data, such as CT, MRI or X-ray images, which can be stored both compressed or uncompressed. While only images can be stored, an attribute might include a set of these, which allows for multiframe imaging or the stacking of images to create 3D shapes. Furthermore, it also includes dose objects to store dose distributions computed by TPS, for example as 2D- or 3D-dose grids, groups of dose points or isodose curves. There are also specific information objects to work with files storing the relevant regions for radiotherapy, as identified by clinicians using imaging technology. Lastly, the DICOM standard for radiotherapy allows the efficient handling of data containing the dosimetric and geometric data relevant for planning treatments using EBRT or brachytherapy.

While the DICOM standard has made the storing and sharing of data across different institutions easier by offering a uniform layout, it also has drawbacks. For example inconsistency in what optional data fields are actually filled poses issues for applications which need a wide array of different information. This case was also encountered during this work on this thesis, as some sets of patient data did not include a field directly containing the total prescribed dose for a treatment plan, which had to be inferred from other attributes.

The specific data relevant for the model training and testing, can be extracted from the patient DICOM files and includes among other data for each patient a set of CT image slices, which are stacked along a vertical axis to allow for a 3D representation of the pelvis region. Furthermore, the data set for a single patient also consists of a so called structure set file, which delineates the regions of interest within this 3D view. One patient data set also includes a dose distribution file, which defines how much dose is deposited within each voxel. Depending on the number of arcs there might also be more dose files. In this case, a dose distribution computed by the TPS was used, however in the full automatic planning pipeline, a second deep learning model can be used before this one, to predict the dose distribution. This dose prediction model was however not relevant for the purpose of this thesis. Lastly, the patient data also contains a plan DICOM file, which contains information about the actual beam specifications of the EBRT treatment as well as information about the dose delivery.

3.2 Preprocessing

After the patient data was exported from the clinical data base and anonymized, it was sent through the preprocessor to format it in a way that could directly be used for training.

To this end, the CT files were stacked vertically to a resolution of $(x, y, z) = (128, 128, 320)$, i.e. a higher resolution along the vertical z -axis, as this was also the axis parallel to the MLC and thus held crucial information about their positions. CT images had a resolution of 0.81 mm and a slice thickness of 2 mm. The coordinate system chosen was such that the center correspond to the beams isocenter, i.e. its center of rotation. The dose distribution was resampled to also fit the above mentioned resolution, which lead to the voxel size changing from $(4, 4, 4)$ to $(3.125, 3.125, 1.250)$ (all in mm). Lastly, the structure set data was also positioned into this coordinate system, but contrary to the single channel CT data, a number of channel corresponding to the regions of interest was used. These regions were the prostate target region itself, now denoted as Planning Target Volume (PTV), the left and right femoral head, the bladder and the rectum. Unlike CT and dose data, these structures were also saved as binary data, as they do not represent intensity, but only whether a voxel is part of a given region.

Generally speaking, any VMAT plan file can be subdivided into so called control points. These are spread around the gantry rotation path at specific angles to represent the collimator positions, storing both leaf and jaw positions, as well as the MU, i.e. intensity. However, not all plans have the same placement of these control points or even the same number. Thus, for the purpose of making the data more uniform and thus potentially more effective for the training process, the control points of all plans were linearly interpolated to a set of 180 uniformly spread, a resolution of 2° , control points. The initial interpolation process would sometimes lead to negative MU values and unphysical movement of leaves, thus the interpolation process was overhauled to fix these issues. This was applied on the preprocessing for all models.

All data was stored using NumPy arrays, the structure set data for example being stored in an array of shape $(c, 128, 128, 320)$ depending on the number of channels c , while CT and dose data had only one channel, but the same spatial dimensions. The leaf positions were stored in an array of shape $(2, 180, 80, 1)$, where the first dimension represents the two sides of the collimators, i.e. left and right leaves, the second corresponding to the 180 control points, the third to the 80 leaves on each side and the last dimension being for easier handling. The jaw positions were stored similarly, replacing the 80 leaves with 2 jaws, and the MU given by just a single number per control point.

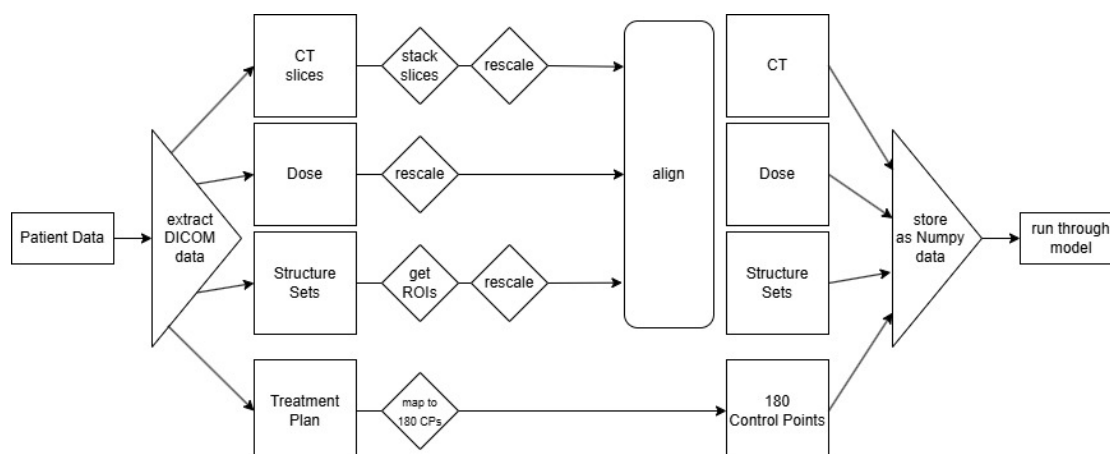


Figure 3.1: Preprocessing pipeline as used for the patient data for the model training and testing.

3.2.1 Augmentation Strategies

In total, three different developed augmentation strategies, previously developed in-house in another work, were employed.

First, within the planning of a treatment plan in a TPS, one can define certain objective functions, which the system tries to achieve during optimization. This could be for example a dose threshold in a certain region. To this end, all plans from the Department of Radiation Oncology at the University Hospital Vienna were recalculated once again, but this time using a new template for these objective functions, effectively doubling the number of available patients.

Furthermore, both the bladder and rectum positions were slightly altered for each patient and the plan, using the original objective functions, was then recalculated to incorporate these changes to the anatomy.

However, it has to be noted that not all augmentations could successfully be applied to each of the patients, some of them because of software reasons and others because of geometrical reasons, for example when a movement or change in size of an organ was not possible as it would have overlapped with another one. For these patients just the original plan was included in the training data.

In total, of the 129 patients that were selected for the training and validation set, 117 were found to be fit for the actual training process. The remaining twelve patient treatment plans had to be excluded. For example, some had different prescribed dose values or differences in the plan geometry, i.e. two arcs instead of one. From the 40 patients of the test set, six patients had to be excluded. The excluded treatment plans consisted of

more than one arc and the model can only predict machine parameters for plans with one arc, which would have made it very hard to effectively compare the original and predicted plans.

The 117 patients in the training and validation set were then used for the augmentation algorithms. From the first augmentation, i.e. using the new template for the objective functions, 94 new data sets could be generated. The alterations of the patient anatomy yielded 101 and 102 for the bladder and rectum modifications respectively. The 414 resulting data sets were partitioned into training (75%) and validation (25%) subsets, containing 310 and 104 sets respectively.

The final number of patients and data sets can be seen in Table 3.1.

Table 3.1: Overview of the number of patient data sets. Note that the patients from the test set were not subject to augmentations.

usage	data sets	data sets after augmentations
training	117	310
validation		104
testing	34	34

3.2.2 Limitations

The input data for DL models is of utmost importance and as such a correct handling and preprocessing of data is crucial to achieve good performances. However, there are limitations to how the data can be preprocessed. For example, the interpolation of the control points loses some information. Furthermore, while there is a lot of data available for prostate cancer plans, it is not always uniform, for example in terms of arcs or number of collimator leaves.

During preliminary testing of the influence of the preprocessing on the training and model performance, the preprocessing pipeline was adapted such that both plans with more than one arc/beam and different number of collimator leaves could be included. For example, the MLC head from Varian machines has 60 leaves compared to the 80 from Elekta machines.

In the first case, the data specific to each arc, i.e. the control points for the beam and its dose, were split to be each treated as “new” patients, both using the same CT and structure set data. For the latter case, plans created for the Varian machines mentioned above, i.e. using 60 leaves, were modified in such a way that the topmost and bottom-most ten leaves were split in half to result in a full leaf set of 80. The positions of each

pair split leaf could then be averaged to form one and lead back to the original layout.

This meant that also plans with simultaneously irradiated lymph nodes, as the available plans of this type were all using two arcs, and furthermore plans created for Varian [44] machines could in theory be included in the training.

However, when using training and validation sets including plans with lymph node involvement or plans made for Varian machines, the performance of the model suffered. Thus only plans without lymph node involvement, but more than one arc, were actually added to the training and validation set for the final models. It has to be noted however, that there is still much to be tested in terms of generalization of the input data, which would have been beyond the scope of this work.

3.3 Network Architecture

The DL network for the plan generation is set up with an initial encoder path based on the ResNet architecture, followed by latent space operations and ending with three separate decoders for leaf positions, jaw positions and MU. In Figure 3.2 the model architecture is shown.

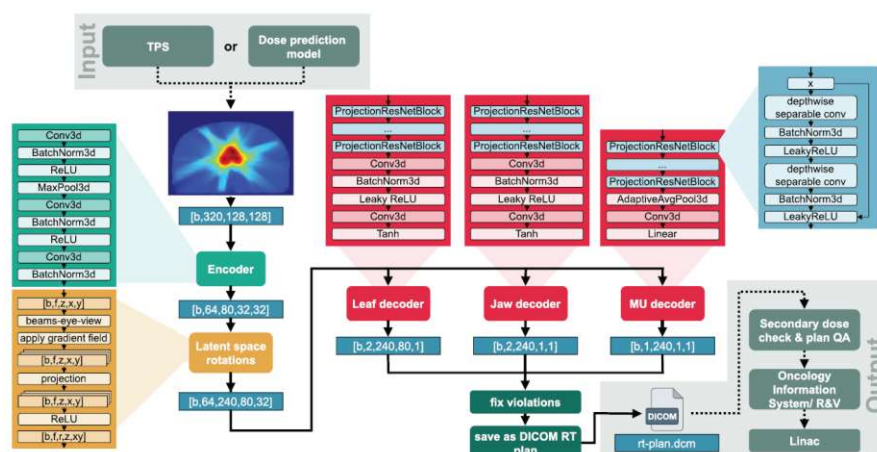


Figure 3.2: Network architecture depicting the path of the initial input data through the encoder, latent space rotations, the three separate decoder heads and finally postprocessing to be assembled into a machine usable DICOM plan. Some of the specific dimensions in the figure have since changed, but the overall structure of the model remains the same. The image was taken from [1].

In the following, the architecture introduced above is explained in more detail, split into the most important parts of the network. Alongside the original state of the model, also the respective changes, i.e. the new models, from this thesis are mentioned. A short

summary of the changes is given at the beginning of Section 3.4 detailing the setup of the numerical experiments.

3.3.1 Dataloading

The base model used the same data format that was created during the preprocessing pipeline. Thus, for each training element first the preprocessed machine parameters were loaded to be used as ground truth in the loss function. Further, the preprocessed 3D patient data was stacked along a feature dimension to produce arrays (now called tensors, to fit the nomenclature of the used software packages) of the shape $(c', 320, 128, 128)$, with the spatial dimensions transposed for easier handling in later steps. The first dimension of size c' included the structure set, dose and CT data stacked on top of each other. Several of these tensors could also be sent through the model at once as batches of patient data of shape $(b, c', 320, 128, 128)$, with the fixed value b being the number of patients in the batch. When using batches of several patients, the respective machine parameters were stacked accordingly as well.

Here, the first changes in the model path were implemented. To allow for deeper models, i.e. more layers, the amount of data that was sent through the model had to be reduced. Initially, the model input would be the whole 3D view of the patient at once, i.e. data of shape $(320, 128, 128)$. The input data was reduced to a set of l leaves out of each patient and the corresponding voxels of the 3D representation, as interference effects were judged to be negligible between slices. The fixed parameter l had to be specified before beginning the training of the model.

This also meant that when finally using this model for predicting machine parameters, overlapping slices of leaves could be predicted. For example, when predicting the leaf positions for a slice of five leaves with the central leaf at index 40 out of 80, the whole slice would encompass leaves 38 to 42. The next index could then be set as 43, which would mean the slice of leaves from 41 to 45. Thus, when finally creating a plan for this model, it would be possible to average the predicted positions for overlapping leaves, potentially gaining accuracy.

A graphical representation of the slicing can be seen in Figure 3.3.

To accommodate the changes for the YOLO object detection, the data had to be further reshaped. All parts concerning the monitor units remained unchanged, for the collimators, both leaves and jaws, the initial format of storing the position of each collimator was changed. Rather than storing the actual coordinate values of each pair, the data was now reshaped to be stored as the center of the bounding box of the opening between the two collimators alongside the width of this opening. This can be interpreted as the bounding box of the object, i.e. the opening, that should be detected. As the model would work with binary classification, detecting whether the collimators are closed or open, with the center and width pair only being relevant in the latter case, another di-




Figure 3.3: The reduction of the full set of 80 leaves to small subset of, in this example, five leaves between the red lines [Screenshot from RayStation].

mension for the class and one for the confidence was added. Thus, the beam data for the leaves changed shape from $(2, 180, 80, 1)$ to $(4, 180, l, 1)$, with the first dimension being the class, confidence, center and width for the YOLO implementation and the third dimension being the number of leaves, which might be the full 80 but can be reduced using the aforementioned slicing.

For further analysis during the training, functions to map the data from one format to the other were implemented. These also incorporated the possibility to include more than one supercell for the YOLO format. This can be pictured as splitting the full range of the possible leaf positions into S equally large parts, S being the number of supercells. The model would then also learn in which supercell the center of the opening of the leaves could be found, while the width might also overlap into another supercell.

3.3.2 Encoder

For the encoder, the employed ResNet extracts feature information from the 3D data and reduces spatial information. This is done by using convolutional layers which reduce the initial data dimension by exactly halving it. In between convolutions the encoder blocks also used batch normalizations and Rectified Linear Units (ReLUs) as activation functions. The data was sent through several of these blocks until the data had been sampled down to the desired dimensions.

The encoder had to be adapted to fit the specific dimensions set by the size of the slices mentioned before. This meant that some of the convolutions had to be changed to yield the correct spatial dimensions of the data.

Furthermore, for the YOLO based model, more changes were implemented. Firstly, the pooling operation was changed from one retaining the maximum activation to an average pooling operation to lose less context in the data. Additionally, the encoder was deepened to include more convolutional layers. However, striding was removed from most of these layers which meant that the resolution after encoder was no longer (80, 32, 32) but rather (160, 64, 64) from the initial (320, 128, 128). Furthermore, all normalization layers in the encoder were changed from batchwise normalization to instance normalization. This was done because the object detection required a reduced batch size for memory reasons as well as high spatial accuracy.

After preliminary testing showed that the structure set data decreased the encoders performance for the YOLO based model, this data was removed from the previous encoder path entirely. This means only dose and CT data were sent through the encoder, after it had been adapted as mentioned. Instead, the information from the structure set was embedded in 64 new dimensions which would then fused together with the dose and CT data.

The embedding was done by transforming the binary structure set data with several channels, one for each relevant region, to a single channel label map, with each voxel having exactly one index, corresponding to the respective region, or the background, i.e. “normal” tissue. This data was then embedded into a feature space of the same dimension as the features from the original encoder.

After the embedding, both the embedded structure set and the encoded dose/CT data were fused together in a fusion block consisting of a single convolution, followed by an instance normalization layer and ReLU activation.

3.3.3 Latent space manipulations

After the data had been encoded, the latent space representation of the initial 3D data input was first rotated around the z -axis along the positions specified by the 180 control points. These 180 volumes were then stacked together. Before that however, for each of these volumes, a projection along the beam path through the patient, the so called beams-eye-view, was created by resampling volumetric data at the beam’s isocenter, i.e. eliminating the depth of the representation. However, since these projections reduce the information to two dimensions, the view from two directly opposing control points would look the same. To account for the physical difference between the beam entering from one or the other side, a gradient field was introduced. This gradient field had a decreasing value depending on the depth within the 3D representation. Thus, after multiplying the data tensor with this gradient field, the voxels on the beam’s entry side had heavier weights than to those on the exit side.

During the course of this work it was discovered that some plans incorporated an additional tilt of the collimators. This can be visualized by rotating the leaves and jaws in Figure 3.3 but keeping the colored organs in position. To account for this, an additional

rotational matrix was implemented in the latent space manipulations.

3.3.4 Decoder

For the basic model, both the leaf and jaw decoder heads are designed similarly: a number of convolution blocks, with residual connections and two depthwise separable convolutions each, to progressively reduce the data to match the expected output shape. The MU decoder has six convolution blocks followed by average pooling and a final output convolutional layer. The tanh activation function is used for the output of the leaf and jaw decoders, while the MU decoder directly uses the output from its final convolution layer. The model then predicts the MLC sequence, the jaw collimator positions and the MUs.

In contrast to this, for the YOLO based model, the two depthwise separable convolutions in each block were replaced with a single 3D convolution to preserve spatial information, which was previously lost due to the separate convolution of the z and (x, y) -dimensions.

Furthermore, the activation layer was also changed. Specifically, for the leaf and jaw collimators, the activation for the classification and confidence was changed to the sigmoid activation function, as this function is well-suited for classification tasks. No activation was used for the decoder output for the center of the bounding boxes, while ReLU as the activation function for the width as this value would be non-negative. The activation for the MU output remained the same.

Lastly, for all models using the slicing algorithm to reduce the size of the input data, all functionality regarding the jaw decoder was removed. Because many slices used for training did not actually incorporate the part of the 3D representation in which the jaw positions were, the models did not learn to predict these coordinates. Furthermore, it was also discovered that the YOLO based models struggled with the jaw positions and thus this decoder was removed there. During prediction, the values for the jaw positions were set by taking the extremal vertical extent of the PTV and fixing the jaws slightly above and below that. This could be done as the jaws rarely moved much beyond the above mentioned extremal positions in most plans anyway.

3.3.5 Loss function

For the training of the base model, the $L1$ -loss was calculated at each control point for the positions of each of the jaw and leaf collimators and the MUs. These were then summed up with the leaf position errors being weighted 100 times stronger than the other factors, to emphasize the learning of this property.

For the YOLO object detection, the loss was calculated depending on whether or not an object was detected as being present, i.e. whether or not the collimators were open or closed. In the former case, the mean squared error for the center and the mean squared

error of the root of the width prediction was calculated. Furthermore, the confidence loss was computed by using the aforementioned intersection over union. For the latter case where the leaves were predicted to be closed, only the confidence loss was computed. The error in the classification was calculated in both cases. Finally, all the components were added up to produce the full loss.

Furthermore, a loss function based on the calculation of a pseudo dose was implemented. This pseudo dose engine was originally created by researchers from the University of Umeå in Sweden and adapted for the present case. It was based on a pencil beam kernel and was created based on the algorithm by [45].

This engine was built such that it can be used with only MLC positions and MUs as input together with the 3D representation of the patient in form of the stacked CTs, meaning without the dose and structure set data. It then calculated the dose distribution for these machine parameters using a pencil beam kernel. The loss was then calculated by computing the mean absolute errors between the two normalized dose distributions.

3.4 Experimental Setup

To assess the effectiveness of the above mentioned changes, three different models were trained on the same set of preprocessed patient data. The created models were then evaluated on the test set of 34 patients.

The first model that was trained was the baseline model M1, against which the performance of the changes to the different parts of the network were tested. This model was in essence the same as the one from [42], that did not use the pretrained encoder weights or online training data augmentation, but included both structure set and CT data. Slight adjustments were made to the preprocessing, which did however affect all of the models compared here, and the selection criteria of the patient data. The model used a ResNet18 architecture as the encoder.

The second model M2 featured the adaption of the base model M1 to the slicing of the input data. For this model, slices of size five, i.e. one central leaf and the two above and below it, were chosen. The model. Thus, the model could be trained with a deepened decoder, using a ResNet50 architecture. Also, the en- and decoder layers had to be adapted to fit the new dimensions of the reduced input data, however their overall structure remained the same.

The third model was the object detection based model M3 using the full input data without slicing, similar to the base model M1. Compared to the base M1, this model featured changes to the representation of the machine parameters when loading the data as well as the splitting the structure set data from the encoder and embedding their information. Furthermore, the decoder was slightly adapted and a new loss function based

on the YOLO algorithm was implemented.

The fourth model M4 compared to model M3 featured changes similar to those from M1 to M2. The YOLO object detection loss was employed and all other changes were applied as well. The key difference to model M3 was that the input data for each patient was split into four slices of 20 leaves each, which were then sent through the model at once as one batch.

The fifth model 5 was the pseudo-dose engine based model. This model differed from the baseline model M1 only in the vastly more complex loss function, which did not rely on simple regression from the L1 loss of the machine parameters, but rather on the difference between the original dose distribution and the approximated dose distribution from the dose engine based on the currently predicted collimator positions and MUs. However, this engine was not actually used as the trainable loss function as it would have drastically increased the time for each iteration from a few seconds to several minutes. This was both because of the speed of the calculations itself and because of the need to facilitate a transferring of gradients between the TensorFlow based dose engine and the rest of the model, which was based on PyTorch.

In Table 3.2 an overview of the trained model and their training data sets is given.

Table 3.2: Overview of trained models and their respective characteristics. Note that the models sometimes also differed in more ways than these characteristics, but these chosen here for clarity's sake to distinguish the models. The full extent of the differences between the models can be found in the respective subsections of Section 3.3.

model	characteristics
M1	base model
M2	en- and decoder adapted to work with slices of input data
M3	YOLO object detection loss
M4	YOLO object detection loss en- and decoder adapted to work with slices of input data
M5	dose engine for validation in loss function

3.4.1 Network training and validation

All models were trained using the Lion optimizer [46] running for 20k iterations using a batch size of 5 for the base model M1 and the model M3, while for memory and training speed reasons a batch size of 1 was used for the other models. The model training was done on the cluster of the Medical University Vienna using an Nvidia A100 GPU with 40GB VRAM.

While for the other models in each training or validation step a randomly chosen patient was passed through the model, for model M2 a slice of 5 leaves was further extracted from this patient and then this slice was used for training. The slice of 5 leaves was centered around a randomly chosen index. For the training process, the selection of this slice was weighted with a normal distribution to ensure that the model would more efficiently learn the positions of the central leaves. This was done as the top- and bottommost leaves are usually closed and hidden behind the jaw collimators anyway and thus contribute less to the final dose distribution.

The YOLO based models M3 and M4 were trained using three supercells $S = 3$.

After each epoch, the L1 loss for each of the leaf and jaw coordinates as well as the MU values was computed for comparison, as the loss values themselves could not be compared on their own. To this end, the output of the YOLO-based models M3 and M4 had to be reconverted from the center and width format to the original coordinate format.

These errors were used as the first metric to evaluate the performance of the models. However, it has to be noted that these values are only surrogate indications for the quality of a predicted treatment plan, as different collimator positions and MU values can lead to very similar dose distributions. Thus, further dosimetric analysis was performed to analyze the performances in more detail.

3.4.2 Dosimetric evaluation

After the training was completed, all models except model M5 were used to predict machine parameters for the treatment plans of the test patients. Model M5 was excluded as the architecture of the model was the same as for the base model M1. The predicted machine parameters were loaded into a predefined DICOM RT plan template.

This was done by first copying the CT and DICOM structure set file from the original patient data set and matching their IDs to the new plan file. This had to be done as each patient treatment plan has a set of associated files with specific IDs and the new plan would have to mimic a real clinical plan in order to be loaded into RayStation.

Then the predicted parameters for each control point had to be interpolated again from the standardized 180 control points onto the set of 151 points defined by the template as most plans used at the Department of Radiation Oncology (AKH Vienna, MedUni Vienna) only incorporate the central 151 control points, cutting off those at the edges. However, since the specific positions of all control points agreed for both the predicted set and the predefined template, only the information about the first and last few control points was lost.

Currently, some postprocessing is still necessary to ensure compatibility of the predicted plans with the specifications of the TPS. This includes for example ensuring that the MU value at each control point does not fall below the minimal threshold or that the leaf and jaw movements do not violate the minimal and maximal positions allowed by the TPS.

Once this has been done, the plans can be loaded into the TPS for further dosimetric analysis. The new set of machine parameters can be used to calculate the thus specified dose distribution with the TPS using the CCC algorithm. To assess the quality of the produced dose distributions, the Gamma Passing Rate (GPR) was computed and analysis of plan quality using DVHs was performed.

Gamma Pass Rate

The GPR is an extensively used quantitative metric for the assessment of dose distribution agreement between a reference dose and a calculated dose, for example a dose distribution from a TPS in the context of treatment plan verification or QA by clinical personnel. Thus, in the clinical practice, the GPR is used to compare an experimentally measured dose distribution and one from the TPS, for example from a Monte Carlo simulation. However, here it is used to evaluate the agreement of the original dose and the predicted dose. It takes both difference in the dose values and spatial discrepancies between two dose distributions into account.

The gamma evaluation involves calculating a gamma index γ for each point in the reference dose distribution. Mathematically, this gamma index can be defined as:

$$\gamma(\vec{r}) = \min_{\vec{r}' \in R} \left\{ \sqrt{\left(\frac{\Delta d(\vec{r}, \vec{r}')}{\Delta d_{\text{crit}}} \right)^2 + \left(\frac{\Delta D(\vec{r}, \vec{r}')}{\Delta D_{\text{crit}}} \right)^2} \right\} \quad (3.1)$$

where $\Delta d(\vec{r}, \vec{r}')$ is the spatial distance between the reference point \vec{r} and the evaluation point \vec{r}' , $\Delta D(\vec{r}, \vec{r}')$ is the dose difference between these same two points, Δd_{crit} is the Distance-To-Agreement (DTA) criterion (maximal distance between points, typically 2 to 3 mm) and ΔD_{crit} is the Dose Difference (DD) criterion (maximum dose difference between points, typically 2 to 3%). For a global analysis, the dose difference $\Delta d(\vec{r}, \vec{r}')$ is calculated depending on a normed dose, which was taken to be the prescribed dose of the original plans. The GPR can be calculated using a global or a local approach, where the latter normalizes dose differences to the maximum dose in the reference distribution, while the former normalizes to the dose at each comparison point, making it more sensitive in low-dose regions. For this analysis the global approach was used.

The γ -criterion is then defined as follows:

$$\begin{aligned}\gamma(\vec{r}) &\leq 1 : \text{test passes} \\ \gamma(\vec{r}) &> 1 : \text{test fails}\end{aligned}$$

The GPR is then defined as the percentage of points within the respective region for which $\gamma \leq 1$, i.e. the percentage of points that “pass”. For instance, a GPR of 90% using 3%/3 mm criteria means that 90% of the evaluated points met the combined dose and spatial accuracy thresholds. This criteria was also used in this work.

Dose-Volume Histograms

DVHs, specifically cumulative DVHs, are generated by computing the dose distribution across a structure, e.g. a tumor or OAR, and plotting the percentage or absolute volume of the structure receiving at least a given dose. This results in a curve that shows how much of the structure receives a minimum dose level, which helps to analyze the dose uniformity and coverage. Thus, while the GPR is a measure of how similar two dose distributions are, the DVH allows for the treatment plans to be compared with respect to the dose coverage and the sparing of healthy tissue.

For the analysis of a treatment plan, different types of DVHs might be used. In this study we defined the following metrics:

- D_{mean} : mean dose in the volume of the structure,
- D_{98} : dose received by 98% of the volume of the structure,
- D_2 : dose received by 2% of the volume of the structure,
- V_{95} : volume of the structure receiving at least 95% of the dose

The first three metrics were calculated for all the five input structures: the PTV, the left and right femoral head, the bladder and the rectum. The V_{95} was only calculated for the PTV, as we were especially interested in how well the prostate is covered in the dose distribution. For irradiation targets, the prostate in this case, the goal is generally to achieve high dose conformity, i.e. a steeply descending DVH at the prescribed dose. For OARs, they can be used to ensure that dose thresholds are not exceeded.

3.4.3 Statistical Tests

The results for both the errors in the machine parameters as well as the calculated DVH metrics on the test set were subjected to statistical tests to detect significant differences.

First, the data was evaluated with the Friedman test. The Friedman test is a non-parametric statistical test and can be utilized to detect differences in data sets across multiple test attempts. This test does not require that the data is normal distributed, making it particularly suitable for the present use case. It operates by ranking the data for each patient and assessing whether the ranks differ significantly across the predictions. This makes it useful for comparing the performance of the different models.

Following a significant Friedman test result, post-hoc pairwise comparisons were used to identify which specific groups differ from each other. The Wilcoxon signed-rank test is a non-parametric statistical test and used for comparisons between two related samples. It evaluates whether the median difference between paired observations is zero by ranking the absolute differences and considering their signs. Similar to the Friedman test, this test does not assume that the differences between paired individuals is normal distributed but rather assumes only that these differences are distributed symmetrically around a central value.

3.5 Soft- and Hardware

The code for the data preprocessing and the model itself was written entirely in the Python programming language, primarily using the PyTorch Lightning [47] library, which is based on the PyTorch software (San Francisco, California) [48]. The dose engine in the loss function of one the models used the TensorFlow (Mountain View, California) [49] library. The preprocessor is written using mostly the NumPy [50] and pydicom [51] package.

All models were trained on an NVIDIA A100 GPU with 40 GB of VRAM on the HPC cluster of the Medical University of Vienna. The NVIDIA NGC Pytorch container 22.07-py3 was used as a base for the interpreter and the tracking software Weights & Biases was used for monitoring the network training and the hyperparameter optimization.

4 Results

In the following section the results of the numerical experiments detailed in Section 3.4 are presented. In Section 4.1 the performances of the models during the training process are shown. In Section 4.2 a direct comparison of the machine parameters predicted for the patients of the set is given. Lastly, in Section 4.3 a dosimetric evaluation of the plans predicted by the trained models is presented.

After the training was completed, all models except model M5 were used to create machine ready plans for all patients of the test set. The predicted jaw and leaf sequences as well as MUs were also used to perform further statistical analysis. For all analysis further than the training performance evaluation, model M5 was omitted as it was basically the same model as model M1 in terms of the network and thus also produced very similar results. Furthermore, during the creation of the plans it became apparent that the YOLO based model using the slicing algorithm, i.e. model M4, did not consistently produce valid plans. While all models required some postprocessing to ensure that the predicted machine parameters did not violate any constraints from the LINAC, the leaf sequences for this model would have required massive adjustments. Thus, no predictions from this model were used for the in-depth model performance comparison.

4.1 Training and validation results

The training process for all models was logged using the Weights & Biases software (San Francisco, California) [52]. Using this software the loss values at fixed intervals were logged. Furthermore, the mean absolute errors in the machine parameters, averaged over all control points and leaves/jaws, were also logged during training.

The loss function values over the iterations are not presented here, as they would not be comparable across models. The loss for models M1 and M2 was the basic regression loss, made up of the mean absolute errors in the predicted and ground truth machine parameters with the leaves weighted by 100. In contrast the loss for the YOLO based models M3 and M4 was calculated from classification and regression components. The former applying always and the latter only being calculated, when the model detected an object, i.e. the leaves being opened. The regression component was then made up of the mean squared errors in the leaf and jaw positions as well as the confidence of the prediction. The MU loss was left as a simple mean absolute error regression. Lastly, the dose engine loss for model M5 calculated the differences between the original and predicted normalized dose distributions.

However, all model outputs during the training were reconverted to the original coordinate format for the collimators to enable comparison. Graphical representations of the mean absolute errors for the different machine parameters are given in Figure 4.1, Figure 4.2 and Figure 4.3.

For the MLC the two models working with the original network, i.e. models M1 and M5, performed the best, both starting with relatively low errors of about 3 mm and ending with the lowest errors at slightly below 1 mm. Compared to that, the two YOLO based models M3 and M4 and model M2 showed higher errors from the start. The best errors that these models achieved were slightly more than 2 mm and slightly less than 3 mm for model M2 and M3 respectively. Model M4 consistently had the highest errors finally ending the training with errors at about 8 mm.

Since the jaw path was removed from the decoder for most models, only models M1 and M5 predicted these parameters with very similar errors ending at about 1 mm.

The MU were again predicted by all models. Again, the base model M1 and model M5 performed best achieving validation errors of about 1 MU. Similar to the leaf positions, model M2 had the next highest final errors, finishing just below 1.5, while models M3 and M4 had errors just slightly above 1.5. However the errors from all models remained below 2 MU.




Figure 4.1: Comparison of the mean absolute error in the leaf positions from the validation for all models.

Model M5 was trained with the same settings and the same architecture as model M1, apart from the additional pseudo dose loss function in the validation. The loss value decreased sharply in the first 2000 iterations. After this, the loss varied between iterations and did not properly converge during the course of the further training. In Figure 4.4 the pseudo dose from the validation of model M5 is shown. It can be seen that the

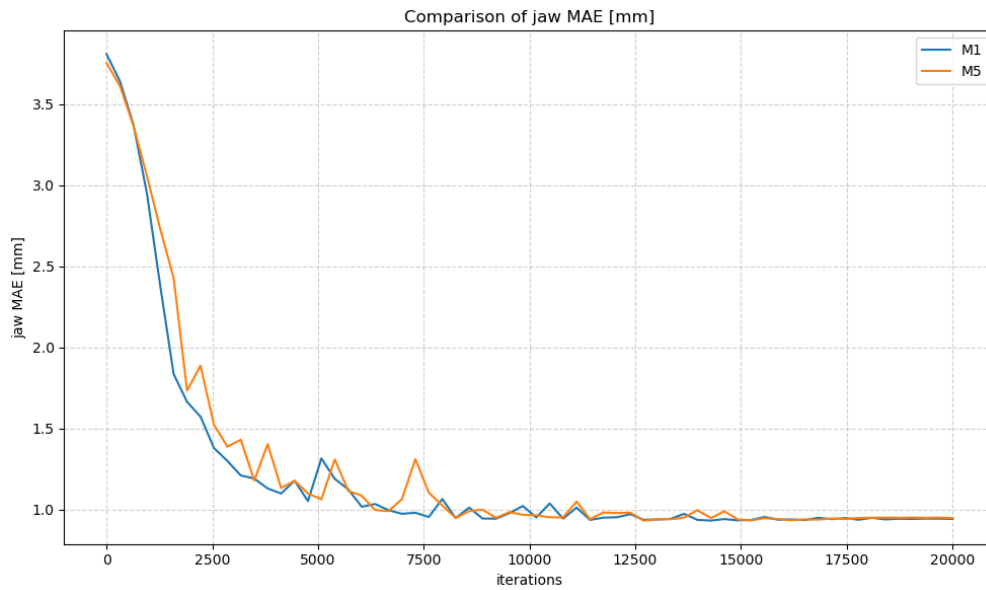


Figure 4.2: Comparison of the mean absolute error in the jaw positions from the validation for models M1 and M5. The errors were very similar as these models had the same architecture.




Figure 4.3: Comparison of the mean absolute error in the MU from the validation for all models.

dose loss plateaus after the initial few thousand training steps, suggesting that despite increasingly accurate machine parameter predictions, the model fails to further improve the actual physical dose distribution.

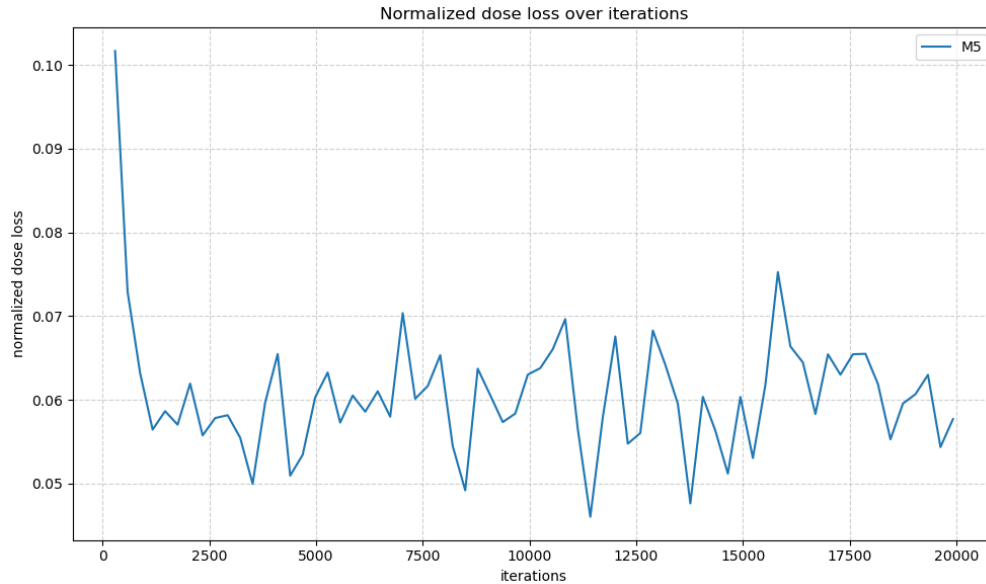


Figure 4.4: Graph depicting the pseudo dose loss over the iterations. While initial improvements are visible, the loss stagnates early, indicating that the predicted dose distributions do not significantly improve beyond the first training phase.

4.2 Geometric evaluation

A detailed view of all errors from the predictions on the test set is given in Table 4.3. Here, the mean absolute errors in the positions and MUs, averaged over all control points and leaves for a single patient, were further averaged over all patients of the test set. Models M4 and M5 were not used for this analysis. The former model was excluded because it could not reliably produce valid plans that could be loaded into the TPS. While the predicted machine parameters could have been analyzed, they would have distorted the comparison as they were unphysical. Model M5 in contrast was excluded, because it was in essence the same model as the base model M1, being built on the same architecture.

In terms of the mean absolute errors on the machine parameter predictions, the base model M1 performed the best for the leaf errors, slightly ahead of model M2, with the former achieving errors around 0.7 mm and the latter between 0.8 and 0.9 mm. In contrast to this the YOLO base model M3 had average errors more than twice as high at slightly above 2 mm. The jaw positions could only be evaluated for the base model, as they were removed from the network for the other two. The errors in the jaws positions

Table 4.3: Overview of the performance of models M1, M2 and M3 with their respective mean absolute errors and standard deviations over all patients of the test set. The metrics detailed below are in order: the leaf collimator position errors in mm, the jaw collimator positions errors in mm and errors for the MUs.

model	M1	M2	M3
leaf MAE [mm]	0.71 ± 0.23	0.85 ± 0.24	2.16 ± 0.87
jaw MAE [mm]	1.02 ± 1.44		
MU MAE	2.06 ± 0.25	2.16 ± 0.28	1.90 ± 0.28

were similar to those during validation at about 1 mm. All models performed similarly for the prediction of the MU values, achieving errors around 2 MU.

For the predictions of the leaf positions and the MUs a Friedman test was performed to assess whether the differences of the predicted parameters were significant. As this was the case for both metrics, a further pairwise Wilcoxon signed rank test was performed for the predictions of models M1, M2 and M3. These tests again yielded significant differences for all combinations of models, with p-values < 0.001 .

Boxplots of the errors of the predictions on the test set patients by models M1, M2 and M3 are shown in Figure 4.5 for the mean absolute errors in the MLC positions, in Figure 4.6 for the mean absolute errors in the jaw positions and in Figure 4.7 for the mean absolute errors in the MUs.




Figure 4.5: Boxplots depicting the mean absolute errors in the leaf positions for the patients of the test set as predicted by models M1, M2 and M3.

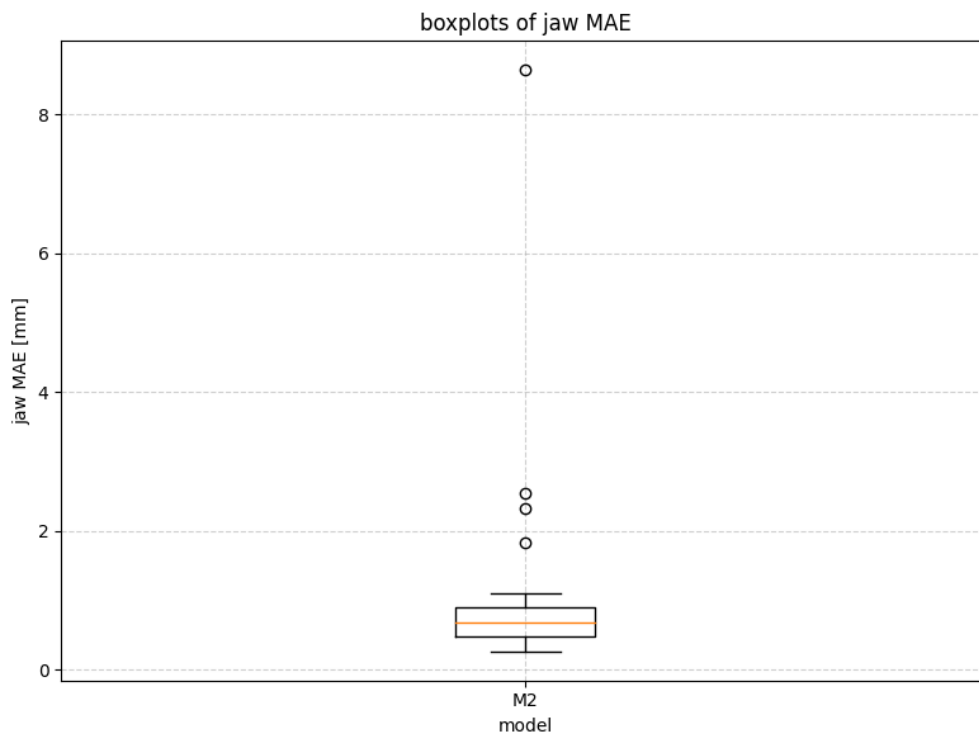


Figure 4.6: Boxplots depicting the mean absolute errors in the jaw positions for the patients of the test set as predicted by model M1. The jaw path was removed from the decoder for the other models.

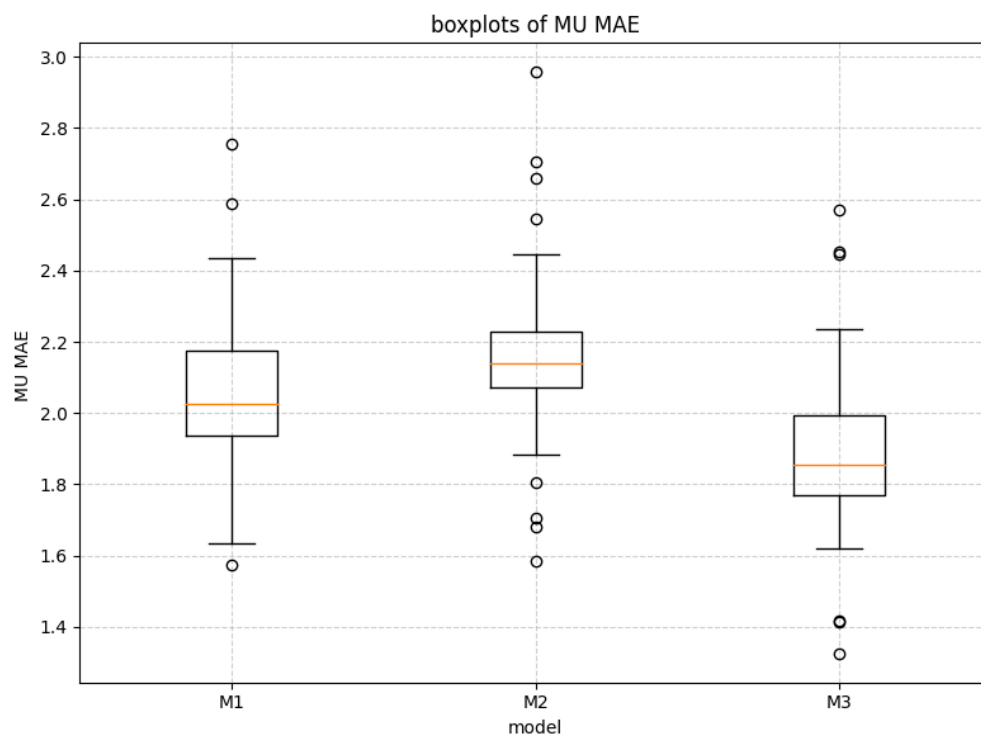


Figure 4.7: Boxplots depicting the mean absolute errors of the MUs for the patients of the test set as predicted by models M1, M2 and M3.

It can be seen that models M1 and M2 exhibit similar errors for the leaves around 1 mm, with latter's errors being consistently slightly higher. In contrast to this, the YOLO based model had much higher errors, averaging around 2 mm, and several extremely high errors above 4 mm.

Furthermore, it can be seen that while the base model could generally predict the jaw positions within one to two millimeters, it also had average errors of eight millimeters for one patient.

The performance in terms of the MUs was similar for all models.

4.3 Dosimetric evaluation

After the training and prediction on the test set patients. The collected machine parameters were postprocessed and stored as DICOM RT plan files. These were then loaded in to the RayStation TPS to calculate the dose distribution defined by the machine parameters. Afterwards, the DICOM dose files were extracted again and analyzed using Python scripts.

4.3.1 Gamma Pass Rate

After the predicted machine parameters were written back into DICOM RT plan files and loaded in to the RayStation TPS where the dose distribution created from by these new predicted plans was computed. These dose distributions were then used for further dosimetric analysis of the models.

The first thing that was calculated was the GPR for the predicted dose distributions of the three remaining models, M1, M2 and M3, using the original dose distribution from the clinical plans as ground truth. The results of this can be seen in Table 4.4.

Table 4.4: Overview of the average GPR for the three models used to predict machine readable DICOM RT plans.

model	GPR [%]
M1	53.67 ± 13.90
M2	39.38 ± 10.05
M3	40.13 ± 8.81

The base model M1 had the highest GPR at 53.67 ± 13.90 %, while the other models had lower rates of 39.38 ± 10.05 % for model M2 and 40.13 ± 8.81 % for model M3 respectively.

In figure Figure 4.8 boxplots for the GPRs over all the patients of the test set for the given models are presented. Form this it can be seen that the base model M1 had the

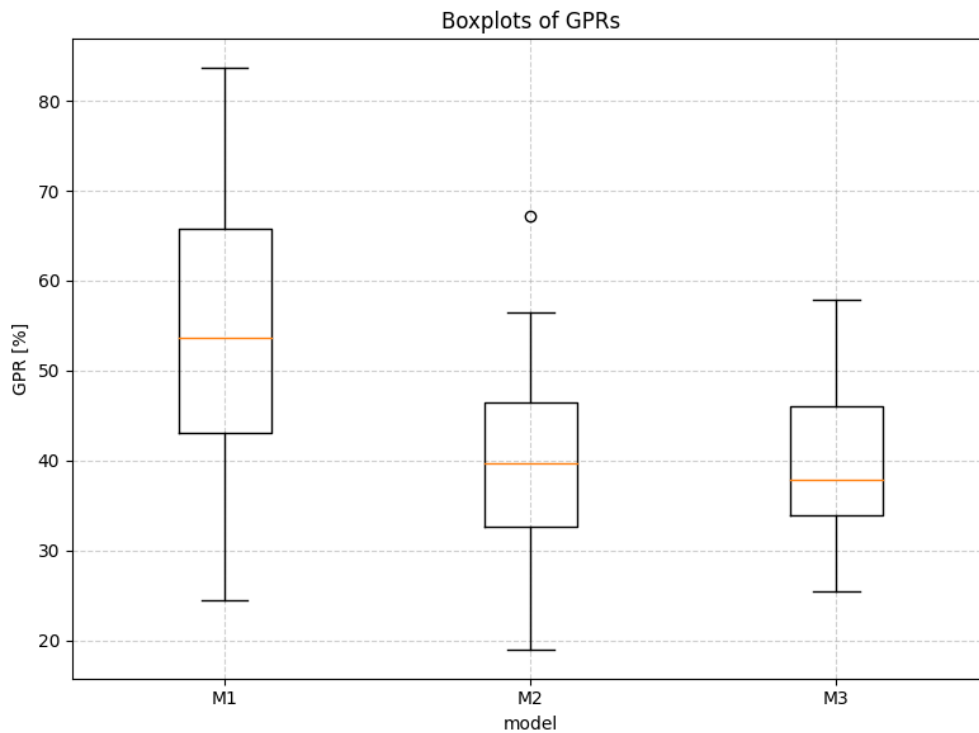


Figure 4.8: Boxplots of the GPRs compared to the original dose distribution from the original plan from the three different models.

biggest variance in the achieved GPRs, while model M3 had the lowest. Model M2 performed similar to model M3, both with a higher variance in the achieved GPRs.

4.3.2 Dose-Volume Histograms

For all predicted plans further dosimetric quantities were computed: D_{mean} , D_{98} , D_2 and V_{95} . The average of these metrics for all models and in the respective OAR can be seen in Table 4.5. The achieved DVH metrics were then used to perform a Friedman test for each organ and metric. If the test yielded p values < 0.05 , the data for this combination was used to perform a pairwise analysis using the Wilcoxon signed rank test for each model and the original plans. The results of this can be seen in Figure 4.9.

For the bladder it can be seen that all models produced similar amounts of dose for all three investigated metrics, achieving a mean dose slightly more than 20 % of the prescribed dose of 60 Gy. On average, these doses were marginally higher than those of the original plan. The results for D_{98} and D_2 were not significantly different with p-values higher than 0.05 in the Friedman test. The D_{mean} metric yielded a p-value below the threshold and subsequent testing using the Wilcoxon signed rank test showed that

Table 4.5: Overview of the average dosimetric quantities D_{mean} , D_{98} , D_2 and V_{95} . The doses are given as percentage of the prescribed dose for the full treatment, which was 60 Gy for all plans. The volume receiving 95 % of the dose is given as percentage of the full region volume.

OAR	Metric	Original [%]	M1 [%]	M2 [%]	M3 [%]
Bladder	D_{mean}	20.39 ± 8.64	21.29 ± 9.26	22.94 ± 10.30	23.06 ± 9.56
	D_{98}	1.80 ± 1.12	1.83 ± 1.16	2.01 ± 1.35	1.95 ± 1.27
	D_2	94.19 ± 6.94	94.99 ± 7.92	94.59 ± 9.22	93.42 ± 10.57
Femoral	D_{mean}	20.23 ± 5.17	17.17 ± 3.93	14.89 ± 3.04	17.25 ± 5.29
Head	D_{98}	1.83 ± 1.34	1.57 ± 1.04	1.90 ± 1.34	2.08 ± 1.56
Right	D_2	42.77 ± 7.99	36.57 ± 6.20	32.39 ± 6.35	36.69 ± 8.23
Femoral	D_{mean}	20.82 ± 4.99	17.85 ± 3.68	15.63 ± 3.24	18.01 ± 4.61
Head	D_{98}	1.71 ± 0.89	1.51 ± 0.76	1.81 ± 1.00	2.05 ± 1.29
Left	D_2	44.05 ± 7.13	38.03 ± 5.90	33.32 ± 6.70	38.61 ± 7.27
Rectum	D_{mean}	40.85 ± 5.71	42.67 ± 6.04	42.07 ± 5.26	41.23 ± 7.34
	D_{98}	3.78 ± 3.04	4.07 ± 3.51	4.62 ± 3.93	4.58 ± 3.84
	D_2	100.52 ± 1.52	99.99 ± 2.40	98.68 ± 2.49	99.42 ± 7.06
PTV	D_{mean}	100.63 ± 0.46	101.11 ± 0.51	100.87 ± 0.42	100.19 ± 0.84
	D_{98}	95.59 ± 0.82	90.34 ± 2.93	87.06 ± 3.33	71.13 ± 5.14
	D_2	105.00 ± 1.08	108.93 ± 2.28	110.54 ± 3.52	120.35 ± 4.79
	V_{95}	98.50 ± 0.71	90.20 ± 5.15	85.05 ± 5.04	67.77 ± 2.56

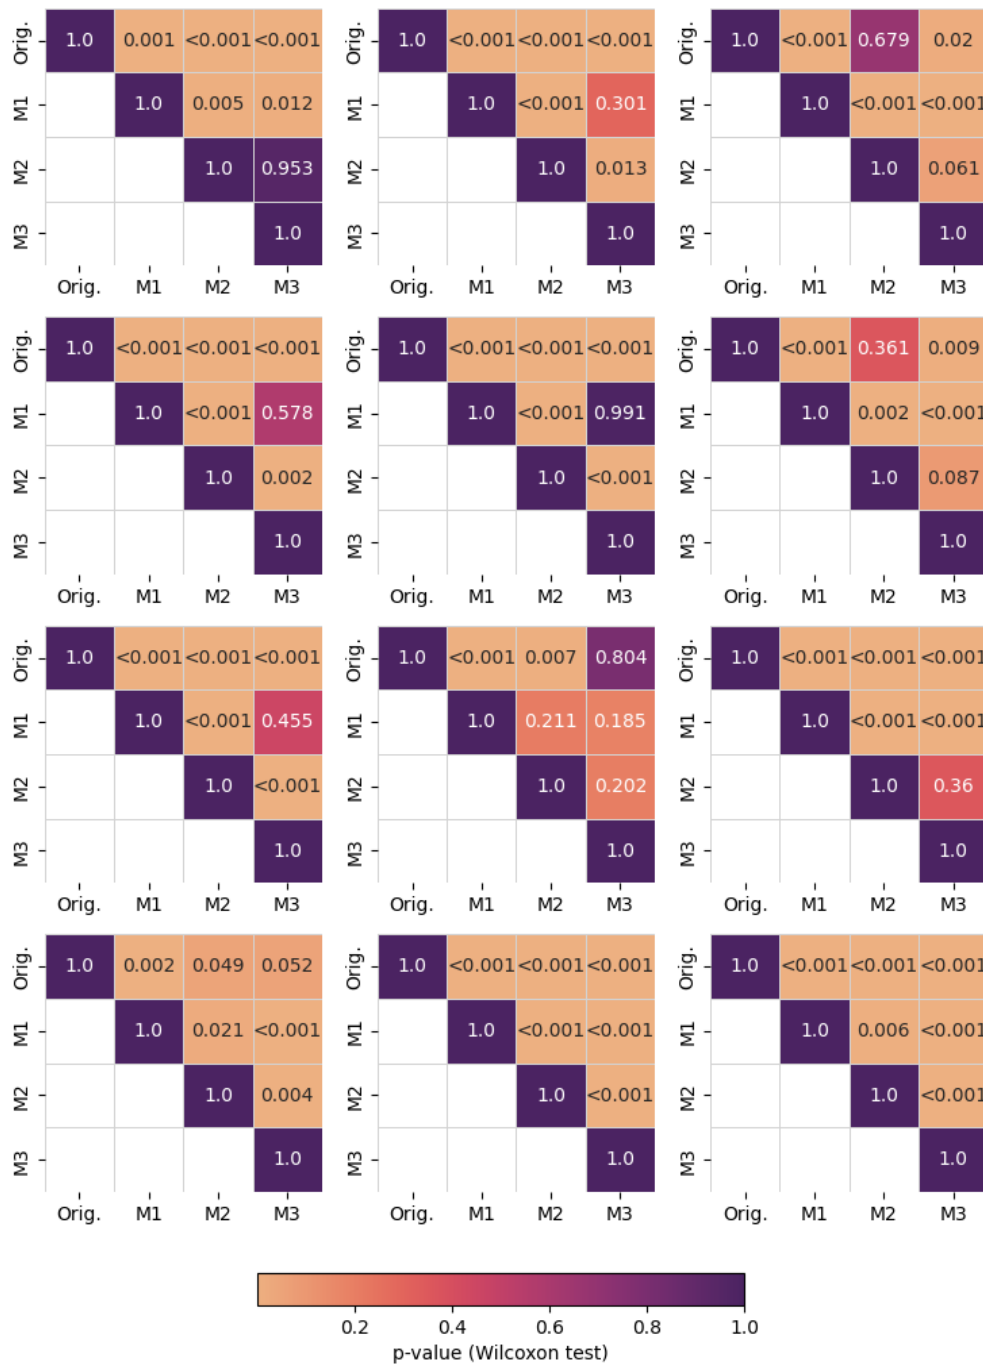


Figure 4.9: Heatmaps of p-values from the pairwise Wilcoxon signed rank tests. The Friedman test also yielded significant differences for the V_{95} metric, however the subsequent Wilcoxon signed rank test yielded p-values < 0.001 for all pairs and thus this matrix was omitted for clarity.

only the comparison between model M2 and M3 did not reach significance for this metric.

The results for the femoral heads showed that the models consistently produced lower dosages for D_{mean} and D_2 . Model M2 achieved the lowest mean dose with about 15 % of the prescribed dose followed by models M1 and M3 with around 17 % each, while the original plan had almost 21 %. The average dose received by 2 % of volume was slightly higher in the left than in the right femoral head, all doses being between 1 and 2 percentage points higher. Model M2 lead to the lowest dosage around 32 % in the right head, followed by models M1 and M3 with about 36 %. The highest doses were in the original plans at about 42 % of the prescribed dose. The average dose received by 98 % of the volume was between 1.5 and 2 % of the prescribed dose for all models and the original plan. The Friedman test yielded significant differences for all metrics of this organ. Further testing with the Wilcoxon signed rank test showed that only the differences in the predictions by model M1 and M3 for D_{mean} and D_2 did not reach significance. For D_{98} the differences between the predictions of model M2 and M3, and those between the prediction of model M2 and the original yielded p-values > 0.05 , albeit the latter only with $p = 0.087$.

For the rectum the achieved dose metrics were similar for all models and the clinical plans. The mean dose received by the rectum was around 40 % in the original plans, which slightly increased by about 1 to 2 percentage points in the model predictions. The dose received by 98 % of the rectal volume was around 4 % of the prescribed dose, with the original dose distribution again having the lowest value. The D_2 value was close to 100 % of the prescribed dose for all predicted and the original dose distribution. Further, the Friedman test did not show significant differences for this metric. In contrast to this, the p-values for the other two metrics were < 0.05 . The differences in the mean dose were significantly different for the original plan and each of those predicted by models M1 and M2. For D_{98} only the differences in the predictions between model M2 and M3 were not significant according to the Wilcoxon signed rank test.

Finally, the mean dose in the PTV was almost the same for all models and the original plan. A more pronounced difference is observed for the D_{98} value, which drops from about 0.96 % in the original slightly above and below 90 for models M1 and M2 respectively. The YOLO based model had a much lower average value achieving only 70 % of the dose in 98 % of the PTV. Conversely, the D_2 value increases going from the original plan to model M1 and further to model M2. Again, the object detection model had the worst result with 2 % of the PTV receiving 20 % more than the prescribed dose. Lastly, the volume of the PTV receiving at least 95 % of the dose was about 98 % percent for the clinical plan, indicating good coverage. This value dropped by almost 10 percentage points for the base model M1 and another 5 for the model with the reduced input M2. The YOLO based model only achieved a dose of 95 % of the prescription in two thirds of the PTV. The Friedman test yielded significant differences for all four metrics. Further testing with the Wilcoxon signed rank test showed that only the differences of model M3

and the original plans of the mean dose were not significantly different with $p = 0.052$, while the differences between model M2 and the original plans yielded $p = 0.049$.

In the following boxplots for the dosimetric quantities for all models and the clinical plans are shown to compare how homogeneous the predicted plans were compared to the original ones. In Figure 4.10 boxplots for the dosimetric quantities from the original plans are shown. The boxplots in Figure 4.11 show the results for these metrics from plans predicted by the base model M1. In Figure 4.12 and Figure 4.13 these metrics for models M2 and M3 can be seen.

Finally, in Figure 4.14 a comparison of the distributions for the volume receiving 95 % of the dose is presented. From this it can be seen, that the original plans are much more uniform in their high coverage of the PTV compared to the other models. Model M3 had the worst coverage, consistently irradiating only about two thirds of the volume with the required dose.

In the following, DVHs for the original and predicted plans are shown. The presented plots are not of the same patients and are meant to highlight the specific shortcomings and strengths of each model. These DVHs were chosen to be representative of the respective model's performance.

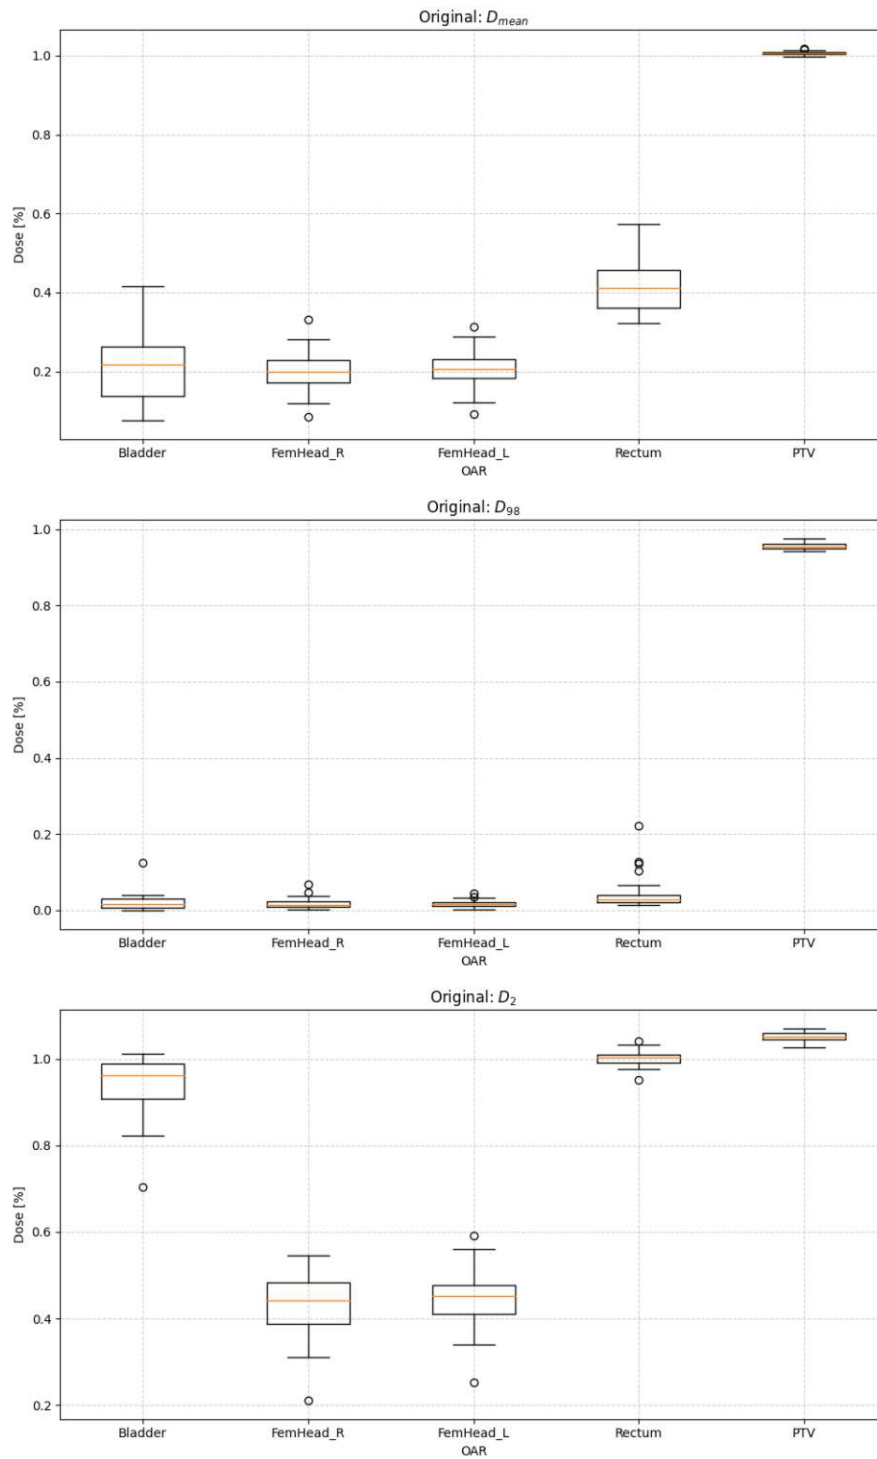


Figure 4.10: Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the original plans used as a reference for all models.

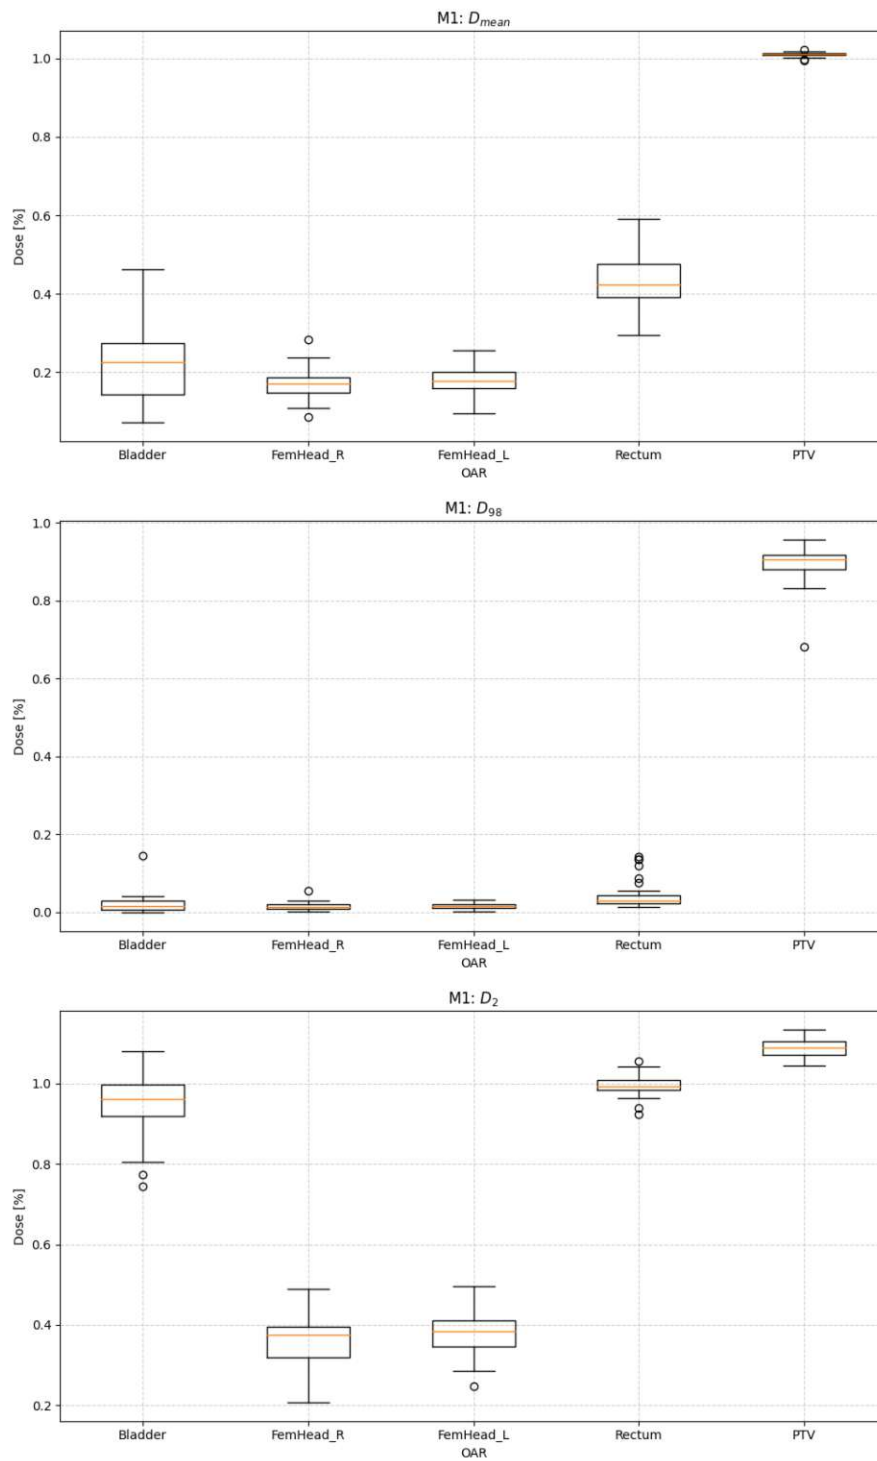


Figure 4.11: Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the base model M1.

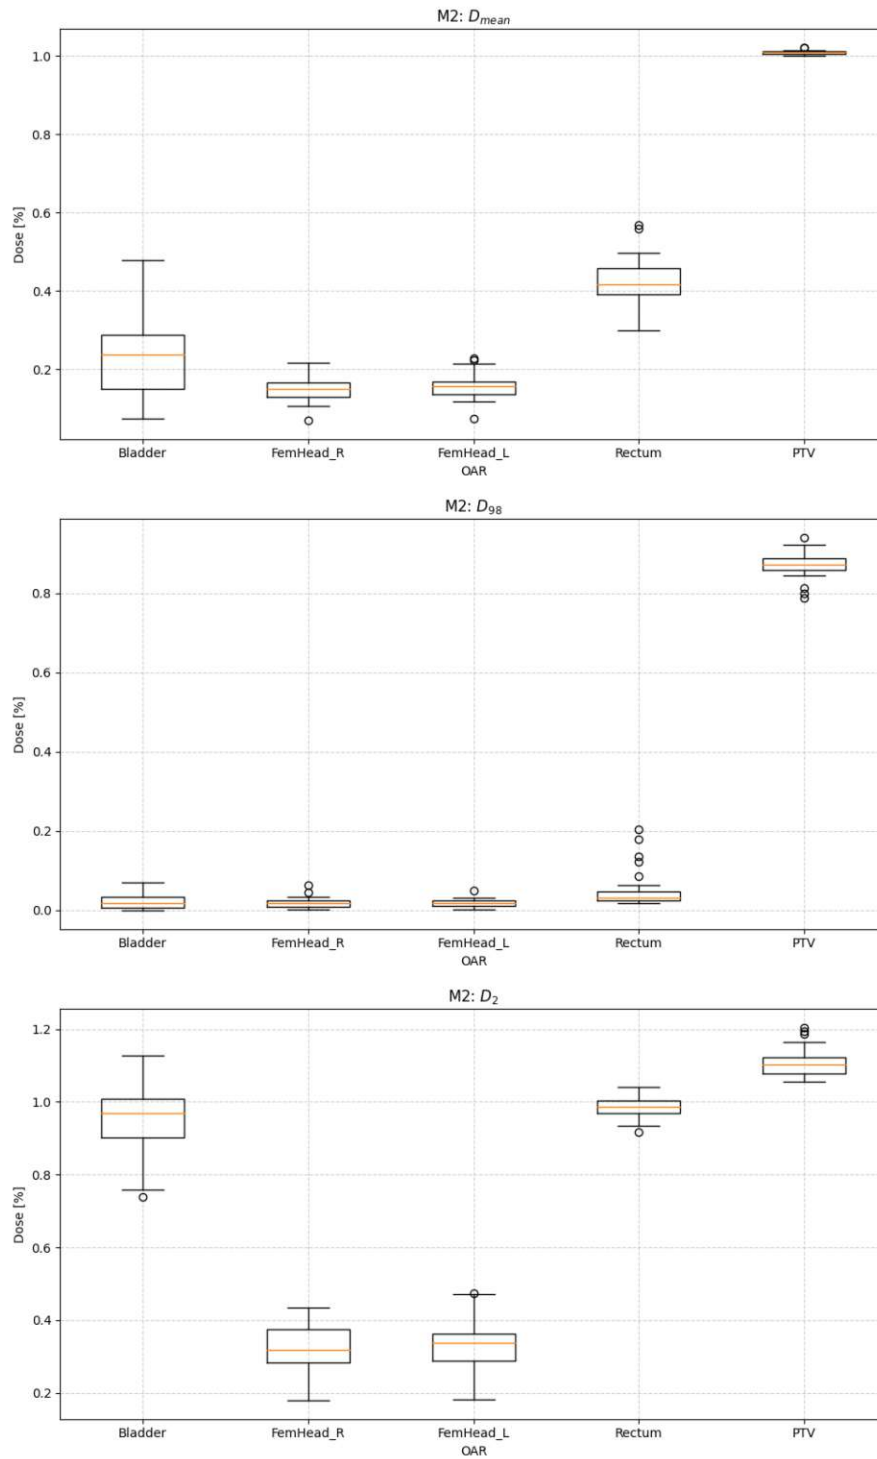


Figure 4.12: Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the model with reduced, “sliced”, input M2.

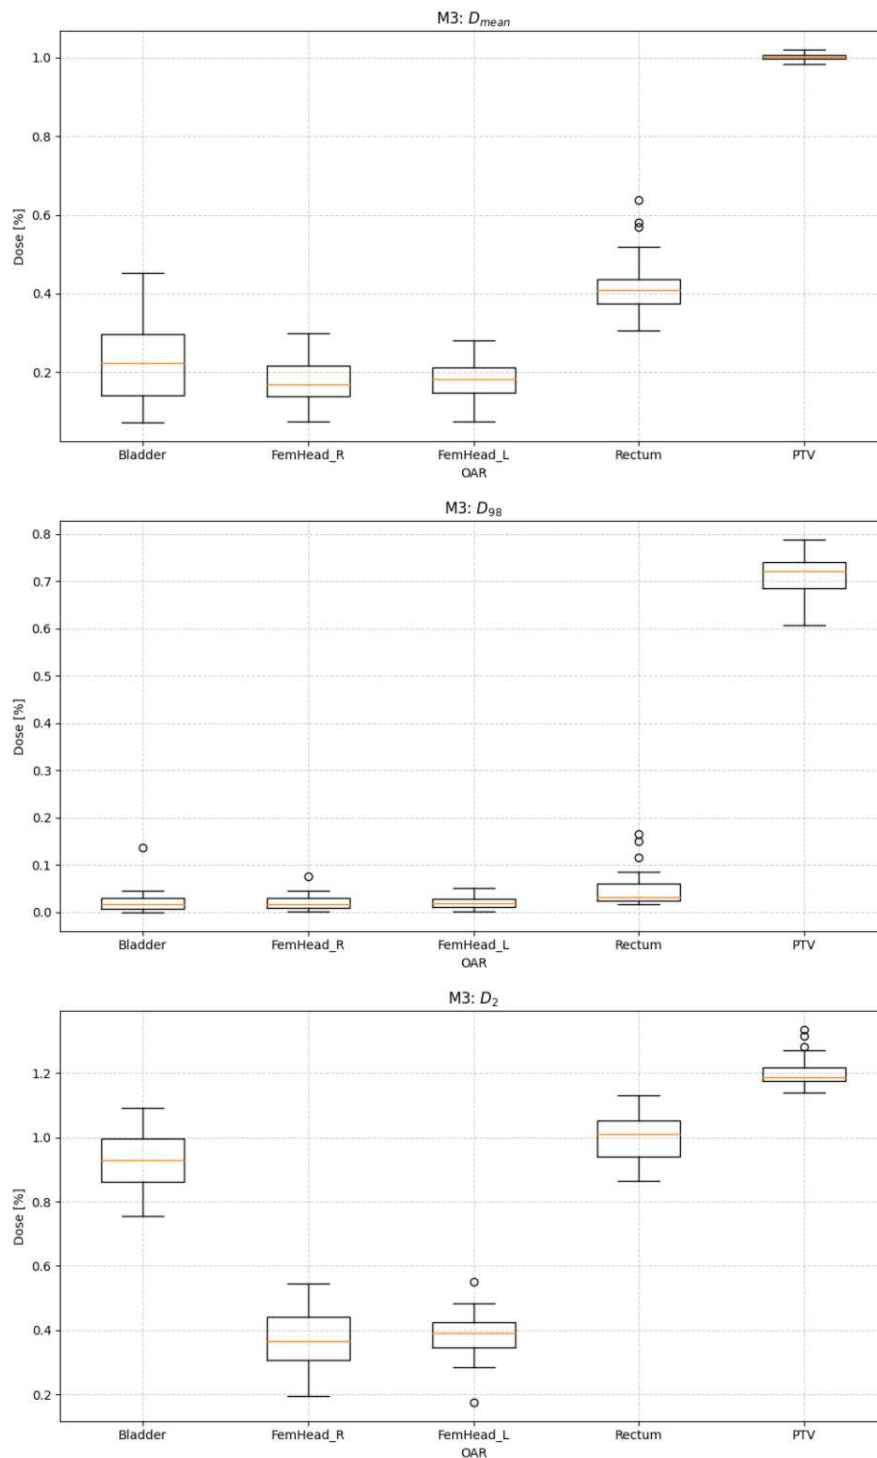


Figure 4.13: Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the YOLO object detection based model M3.

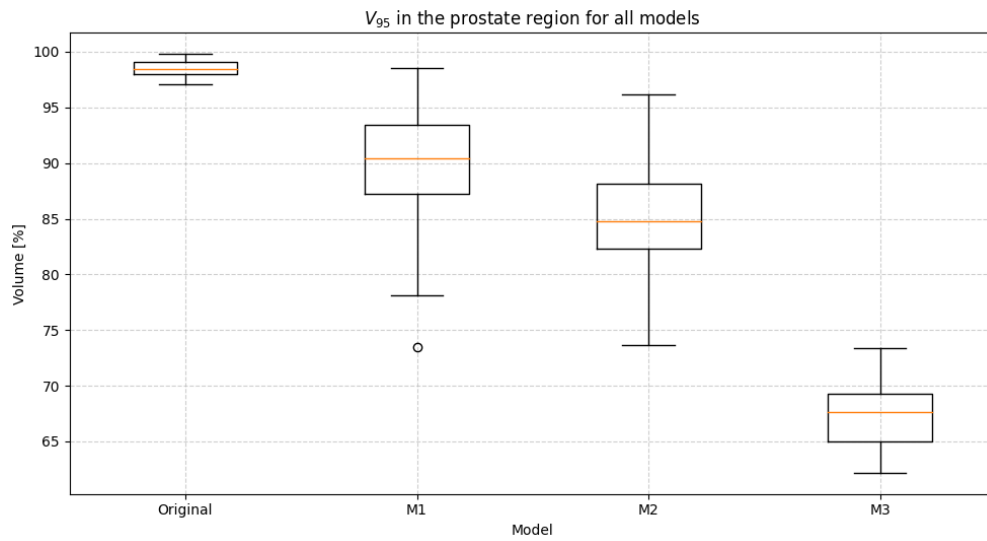


Figure 4.14: Boxplots for the volume receiving 95 % of the dose (V_{95}) calculated from the original plan and from the plans predicted by the base model M1, the model with reduced input M2 and the YOLO object detection based model M3.




Figure 4.15: DVH for the five OAR from the predictions using model M1. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.




Figure 4.16: DVH for the five OAR from the predictions using model M2, achieving a good fit for this model. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.

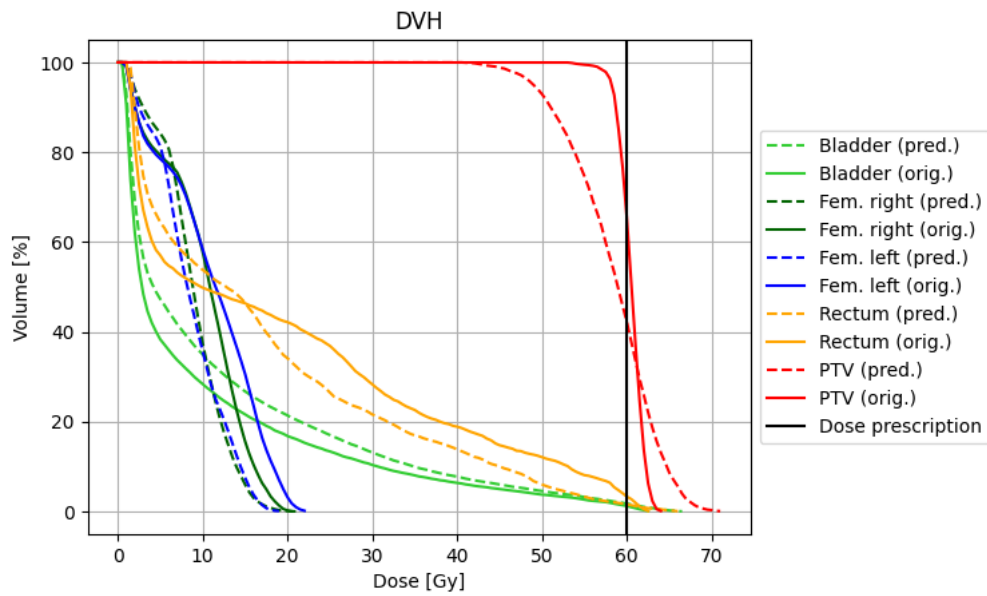


Figure 4.17: DVH for the five OAR from the predictions using model M2, achieving a bad fit for this model. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.




Figure 4.18: DVH for the five OAR from the predictions using model M3. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.

In Figure 4.15 a comparison of the DVH from the original plans and from the predicted plans of the base model M1 can be seen. It is evident that while the model could quite well reproduce the dose distributions of other OARs, except a slight overdosage in the rectum, it failed at producing the same steep decline at the prescription dose as in the original clinical plan.

In Figure 4.16 and Figure 4.17 comparisons of the DVH from the original plans and from the predicted plans of model M2 can be seen. The first DVH is similar to those produced by model M1, albeit with an even less steep decline at the prescribed dose. However, in the second a clear underdosage of the prostate region can be seen. While this model produced similar DVHs to those of model M1 for some patients, it was also more conservative in others.

In Figure 4.18 a comparison of the DVH from the original plans and from the predicted plans of the YOLO based model M3 can be seen. This model showed the most difference in the DVH as it consistently created dose distributions with less dose in the target region. It also produced the dose distributions with the least volume of the prostate receiving the full prescribed dose. However, the plans from this model also lead to the least dose in the other OAR, which can be seen best in the femoral heads and the rectum.

5 Discussion

In this work new strategies that aimed to improve a previously developed model for RT plan generation for VMAT of prostate radiotherapy were explored. The models were adapted, firstly, to take only a very restricted set of slices of the full 3D input data and secondly, to change the task from regression to object detection. While the hypothesis that these changes would improve the model performance and consequently plan quality could not be confirmed, some valuable insights were obtained. In the following, the detailed discussion of the results in terms of geometric and dosimetric performance will allow to put the presented treatment automation model in context with current state of the art DL approaches in the field and will provide useful information for future research to advance the model further.

5.1 Interpretation of results

In terms of the geometric results, i.e. the direct errors in the machine parameters, the base model M1 performed better than the new models. Especially the object detection based models M3 and M4 had much higher errors during training and validation, but so did model M2, which received only a slice of the input, but had a deeper backbone for feature extraction, increasing from ResNet18 to ResNet50. These errors in the machine parameters are only a surrogate for the quality of the actual underlying dose distribution.

Furthermore, the higher errors for models M2 and M4 using the slicing algorithm might be partially explained by the fact that this model was set up to always predict a central slice, i.e. at the height of the PTV, during validation. This focus was chosen deliberately during the setup of the model to emphasize anatomically relevant regions. This means that many of the leaves in focus show a lot of movement over the 180 control points, which is hard to predict, unlike the peripheral leaves which are continuously closed and hidden behind the jaw collimators. Furthermore, the ground truth for the MUs is still the value for the full view, which is not directly related to the sliced input. Compared to that, all models using the full view of all 80 leaves also predict many closed leaves at the top and bottom, which are easier to predict and thus decrease the mean error.

The addition of the dose engine to the loss function of model M5 showed that, while this model performed very similar as the base model as it had the same architecture, the predicted dose distributions did not improve after the first few thousand steps. This means that while the model learned to mimic the leaf positions better and better, it did not actually improve the underlying physical dose. This outcome further highlights the need for a fast, accurate and learnable dose calculation solution to improve the predicted plans.

The dosimetric analysis showed that the two new models used for in-depth analysis actually achieved lower GPRs compared to the base model, decreasing by about 10 percentage points, however also reducing the variance in the achieved rates. While the base model M1 achieved GPRs from about 25 % to more than 80 %, the YOLO based model M2

ranged from 25 % to about 60 %, indicating more uniform plans across the test set. Model M2 usually produced GPRs from 20 % to 60 %, even achieving 70 % in one case. However, clinical goals are generally set at around 90 % GPR and thus no model produced valid plans for clinical delivery.

Nevertheless, it has to be noted that the GPR also has limitations as a measure for the quality of a model. For example, it is very sensitive to low dose noise and as such a threshold of 20 % was defined and only points with a higher dose than this threshold were taken into account. Furthermore, the GPR is only of limited informative value on its own and is thus often combined with DVH analysis.

From the DVHs it could be seen that again the base model M1 produced the best fit to the clinical plans, however even these dose distributions did not achieve a consistent coverage of the PTV. This worsened for the other tested models and especially the YOLO model M3 lead to an underdosage of the prostate, with only two thirds receiving the required dose. In contrast to this, the doses in other OARs could be accurately reproduced by all models, even leading to less dose in some. Generally, it can be said that the predicted plans were more cautious regarding OAR sparing at the risk of underexposing the target. This could also be seen when analyzing the leaf sequences in the TPS, where especially the object detection based model M3 had leaves slightly shielding the PTV in the beams-eye-view. The other models tried to follow the outline of the PTV without overlapping, which was similar to the original plans.

However, it has to be noted that the informative value of DVH has its limits. As some information is lost when summing for the histograms, two plans with very similar DVHs might actually have very different spatial dose distributions within the respective regions.

There would have been possible strategies for further analysis of the created treatment plans, for example through actually delivering them with a LINAC and measuring the results. However, this would have been beyond the scope of this work and also not relevant for the current stage of the project as the plan quality is not yet high enough.

5.2 Comparison with state of the art

Recently, there has been increasing interest in automated treatment planning and several approaches have been investigated to survey the potential of automating various components of the RT planning workflow. Recently, a challenge to evaluate the degree of possible automation in RT treatment planning was published [53]. Earlier approaches included scripting solutions within commercial TPS [54], [55], [56] and dose prediction models which still require a TPS to create a fully usable DICOM RT plan [6], [7], [8]. Another approach uses fluence prediction models, which can provide fluence map without costly mathematical optimization procedures. However, these models rely on TPS as well to produce ready-to-treat plans [57], [58]. Reinforcement Learning (RL) represents an

recently investigated DL strategy that could be used for automating applications in treatment planning [59], [60]. However, the current body of evidence is limited, and further studies are required to determine its reliability, and integration potential into established clinical workflows. In contrast to this, the approach presented in this work departs significantly from these frameworks as it does not rely on conventional TPS except for potential QA after a machine-ready plan has already been generated. This further increases the plan generation speed, as no time consuming interfacing is needed with external software.

There have also been recent studies investigating the use of Particle Swarm Optimization (PSO) enable fully automatic VMAT treatment planning by iteratively optimizing planning constraints based on a defined plan quality score [61]. However, while this approach does not require a commercial TPS, it is computationally intensive and would currently impede the clinical workflow.

In contrast, the method proposed in this work eliminates the dependency on conventional TPS components altogether by sequentially combining two neural network models, where the first predicts a dose distribution that is then used by the second to generate plan by directly predicting machine parameters. This enables the direct production of deliverable plans that can be reviewed using secondary dose verification tools, significantly reducing the total time required for planning. Testing in [2] showed that the full pipeline could create a treatment plan in about 40 seconds, where half the time was need for pre- and postprocessing. This approach has particular advantages in the context of online adaptive radiotherapy, where plans must be adapted in seconds.

The long-term vision outlined in this thesis is the development of a fully integrated, DL-based auto-planning pipeline. Such a system would incorporate multiple sub-models capable of performing auto-segmentation, dose prediction based on CTs and structure sets, and final plan generation in a unified framework. This combination of AI-driven modules could produce clinically acceptable treatment plans within seconds, allowing for fully autonomous planning workflows. When combined with real-time imaging, this approach holds the potential to transform modern adaptive radiotherapy. However, several technical and clinical challenges remain and must be addressed before such a system can be deployed in routine clinical practice.

The long-term vision outlined in this thesis is the development of a fully integrated, DL-based auto-planning pipeline that is capable of enabling real-time adaptive radiotherapy. In this envisioned workflow, the patient arrives for a daily treatment sessions, at which point the current anatomical situation is captured via medical imaging, such as a CT. This image is then sent through a fully automated sequence: first, a segmentation model delineates the relevant anatomical structures; second, a dose prediction model suggests an optimal dose distribution tailored to the patient's current anatomy and tumor status; and finally, the model developed in this work translates the predicted dose distribution into a set of deliverable machine parameters in just a few seconds. This process would eliminate the need for manual plan adaptation and enable fully autonomous

treatment planning.

Realizing such a pipeline would require substantial modifications to current clinical workflows, including mechanisms for rapid plan approval and real-time QA procedures. A review of the potential ways in which AI could alter workflows is given in [62]. Nevertheless, if successfully implemented, this system would offer a significant leap forward in personalized cancer care as patients could receive treatment plans that are anatomically accurate for each session. This could lead to better target coverage, reduced dose to OARs, and ultimately, a lower risk of radiotherapy-related toxicities—thereby improving both treatment outcomes and the patient’s quality of life.

5.3 Future challenges and improvements

A few caveats have to be mentioned regarding the effectiveness of the applied modifications, as especially the YOLO did not achieve the expected performance. The implemented variant of the YOLO framework is a very early one and more sophisticated versions have been released recently [63]. These would require a more involved reworking of the network architecture, but might lead to better outcomes. Nevertheless, it could be shown that the model can actually create valid machine-readable plans, even though they still fall short of established clinical goals.

This modification might also especially profit from a further refinement of the reduction in the model input, i.e. the “slicing” of the 3D patient representation, as this might allow deeper models. Analysis of the trained encoder for the YOLO based model also showed that especially this model struggled with extracting the relevant spatial features from the input data. This indicates that it might be necessary to further deepen and adapt the model architecture in this regard.

Currently, the only learnable, i.e. directly implemented in the loss function, mechanism accounting for the absorption of the radiation in the patient’s anatomy is the linear gradient field that is applied after the encoder. From the current implementation of the dose calculation engine in the validation it can be seen that the model stops achieving better dose distributions very early on in the training, while still learning to mimic the MLC sequences. It is reasonable to assume that - while the geometric mimic of MLC sequence becomes better after 20k iterations, the dosimetric quality of the plan would not be better than that after 2k. This might indicate that the current architecture has hit a plateau in terms of the dose distribution and further highlights the necessity of creating a faster dose engine that is also trainable. Such a dose engine could also be used to specifically train the model to focus on the dose in specific OARs by weighing the voxels accordingly.

Furthermore, the existing $L1$ -regression loss function does not take the different impacts of the leaves on the final dose distribution into account. Currently, it equally weights all leaves, neglecting the fact that the outer leaves are often hidden behind the jaws and

have less influence on the plan quality.

Additionally, the model currently requires extensive vendor-specific postprocessing of illegal MLC positions before a deliverable plan can be created. This is necessary as some predicted parameters violate machine constraints, e.g. not taking into account the actual speed at which the leaves and jaws can be moved. In the future, it might be possible to implement these limitations into the training of the model reducing the need for time consuming postprocessing

Another possible approach to improve model performance would be to integrate transformers into the decoding network. Transformers, originally developed for NLP tasks, are particularly effective at capturing long-range dependencies and sequential patterns through their self-attention mechanisms [35]. Applied in this context, a transformer-based decoder could enable the model to more effectively learn the sequential nature of the machine parameters along the control point path. By attending to multiple control points simultaneously, the model may further gain a better understanding of how changes in earlier segments influence subsequent MLC positions and dose delivery. This could lead to more coherent and physically plausible machine parameter predictions, potentially enhancing both dosimetric accuracy and reducing the necessary postprocessing steps. Transformer based architectures have already previously been used for dose prediction [64], [65].

All ML models heavily depend on quality of the input data to produce acceptable and generalizable results. Currently, the model does not generalize well beyond the specific set of patients in terms of fractionation schemes, tumor extent, and anatomical differences. Furthermore, how exactly patient data is stored and used varies between clinics and departments. complicates efforts to aggregate the kinds of large scale varied data sets needed to train generalized models. One way to overcome this limitation is the establishment of multi-institutional collaborations aimed at pooling treatment data. This would not only increase dataset size but also introduce greater variability within the patient data. Another approach involves the use of more varied data augmentation techniques, for example simulating further anatomical variability or including imaging noise.

6 Conclusion and outlook

In this thesis several modifications to the proof-of-concept model from [1] for creating deliverable RT treatment plans for prostate cancer were implemented and tested. The model was modified to work with reduced input data, which allowed for a deeper network. An object detection based loss function was further implemented and according changes to the model architecture were made. Data augmentation strategies were employed to increase the training data pool and a pseudo-dose calculation engine was implemented in the model validation to provide further insight into the learning of the model. Even though the modified model failed to generate plans that achieved previously defined clinical goals, a number of valuable insights were gained for future implementations. To this end, potential strategies to further advance the model were analyzed, such as implementing a fast, accurate and trainable dose approximation algorithm, streamlining the postprocessing or even including it in the training, and integrating transformer methods into the network.

List of Tables

3.1	Overview of the number of patient data sets. Note that the patients from the test set were not subject to augmentations.	28
3.2	Overview of trained models and their respective characteristics. Note that the models sometimes also differed in more ways than these characteristics, but these chosen here for clarity's sake to distinguish the models. The full extent of the differences between the models can be found in the respective subsections of Section 3.3.	35
4.3	Overview of the performance of models M1, M2 and M3 with their respective mean absolute errors and standard deviations over all patients of the test set. The metrics detailed below are in order: the leaf collimator position errors in mm, the jaw collimator positions errors in mm and errors for the MUs.	44
4.4	Overview of the average GPR for the three models used to predict machine readable DICOM RT plans.	48
4.5	Overview of the average dosimetric quantities D_{mean} , D_{98} , D_2 and V_{95} . The doses are given as percentage of the prescribed dose for the full treatment, which was 60 Gy for all plans. The volume receiving 95 % of the dose is given as percentage of the full region volume.	50

List of Figures

2.1	Percentage depth-dose curves for different beams: ^{60}Co -beam, 6 MeV, 10 MeV and 18 MeV x-ray beams [14].	5
2.2	A schematic overview of the most important interactions and dose components for photon beams, the four most important are highlighted [16]. .	8
2.3	A graphic representation of the principle behind the PB model. The dose at a point P is calculated from the contributions of the individual pencil beams [18].	9
2.4	Treatment planning workflow: the patient is diagnosed and after creating a 3D representation of the patient anatomy using Imaging, the relevant regions are identified. The dose is then calculated and plans are created iteratively until it satisfies predefined clinical goals. After QA by a physician the patient can then be treated using this treatment plan [23].	13
2.5	The "Beams-Eye-View" in RayStation, showing how the vertically moving jaws and the horizontally moving leaves are modulated to irradiate the target area. The area left open by the jaws is given in orange and the affected leaves are shown in orange with their opening being the outlet for the beam during treatment. The differently colored outlines of relevant regions are used to define clinical goals, for example how much average dose should be in the prostate region in light pink [Screenshot from RayStation].	14
2.6	Example for the computation of a pooling operation on a 5×5 input matrix, reducing the output to a 3×3 matrix [29].	16
2.7	Sketch of the principle behind residual learning. Image taken from [31]. . .	18
2.8	Image of the U-Net architecture: the encoder can be seen as the path on the left side which reduces the spatial dimensions down towards the bottleneck at the bottom and then upsamples this representation back to be used. Image taken from [32].	19
2.9	The image is first split in to a grid on which bounding boxes and their respective confidence for containing an object of a given class are superimposed. These are then combined to detect the objects on the actual image [34].	20
3.1	Preprocessing pipeline as used for the patient data for the model training and testing.	27
3.2	Network architecture depicting the path of the initial input data through the encoder, latent space rotations, the three separate decoder heads and finally postprocessing to be assembled into a machine usable DICOM plan. Some of the specific dimensions in the figure have since changed, but the overall structure of the model remains the same. The image was taken from [1].	29
3.3	The reduction of the full set of 80 leaves to small subset of, in this example, five leaves between the red lines [Screenshot from RayStation].	31

4.1	Comparison of the mean absolute error in the leaf positions from the validation for all models.	41
4.2	Comparison of the mean absolute error in the jaw positions from the validation for models M1 and M5. The errors were very similar as these models had the same architecture.	42
4.3	Comparison of the mean absolute error in the MU from the validation for all models.	42
4.4	Graph depicting the pseudo dose loss over the iterations. While initial improvements are visible, the loss stagnates early, indicating that the predicted dose distributions do not significantly improve beyond the first training phase.	43
4.5	Boxplots depicting the mean absolute errors in the leaf positions for the patients of the test set as predicted by models M1, M2 and M3.	45
4.6	Boxplots depicting the mean absolute errors in the jaw positions for the patients of the test set as predicted by model M1. The jaw path was removed from the decoder for the other models.	46
4.7	Boxplots depicting the mean absolute errors of the MUs for the patients of the test set as predicted by models M1, M2 and M3.	47
4.8	Boxplots of the GPRs compared to the original dose distribution from the original plan from the three different models.	49
4.9	Heatmaps of p-values from the pairwise Wilcoxon signed rank tests. The Friedman test also yielded significant differences for the V_{95} metric, however the subsequent Wilcoxon signed rank test yielded p-values < 0.001 for all pairs and thus this matrix was omitted for clarity.	51
4.10	Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the original plans used as a reference for all models.	54
4.11	Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the base model M1.	55
4.12	Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the model with reduced, “sliced”, input M2.	56
4.13	Boxplots for dosimetric quantities D_{mean} , D_{98} and D_2 in the plans predicted by the YOLO object detection based model M3.	57
4.14	Boxplots for the volume receiving 95 % of the dose (V_{95}) calculated from the original plan and from the plans predicted by the base model M1, the model with reduced input M2 and the YOLO object detection based model M3.	58
4.15	DVH for the five OAR from the predictions using model M1. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.	58

4.16 DVH for the five OAR from the predictions using model M2, achieving a good fit for this model. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.	59
4.17 DVH for the five OAR from the predictions using model M2, achieving a bad fit for this model. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.	60
4.18 DVH for the five OAR from the predictions using model M3. The dashed lines are the predicted plan values and the solid lines are the values from the clinical plan. The prescribed dose of 60 Gy is marked with a vertical black line.	60

Abbreviations

- AI** Artificial Intelligence. 1, 2, 21, 65
- ATP** Automated Treatment Planning. 21
- CCC** Collapsed Cone Convolution. 7, 8, 37
- CNN** Convolutional Neural Network. 15, 16, 18, 21
- CPE** Charged Particle Equilibrium. 6
- CT** Computed Tomography. i, ii, 5, 11, 25, 26, 28, 36, 64
- DD** Dose Difference. 37
- DICOM** Digital Imaging and Communications in Medicine. 24, 25, 29, 36, 48, 63, 68, 69
- DL** Deep Learning. i, ii, 1–3, 15, 21, 22, 24, 28, 29, 62, 64
- DNN** Deep Neural Network. 1
- DTA** Distance-To-Agreement. 37
- DVH** Dose Volume Histogram. 12, 21, 37, 38, 49, 53, 58–61, 63, 70, 71
- EBRT** External Beam Radiation Therapy. 4, 11, 25
- GPR** Gamma Passing Rate. i, ii, 37, 38, 48, 49, 62, 63, 68, 70
- IMRT** Intensity Modulated Radiation Therapy. 5, 12, 13
- KERMA** Kinetic Energy Released per unit MAAss. 6
- LINAC** Linear Accelerator. i, ii, 2, 12, 14, 22, 24, 63
- MC** Monte Carlo. 7, 9, 10
- ML** Machine Learning. 2, 21, 24, 66
- MLC** Multi-Leaf Collimator. i, ii, 12, 22, 26, 28, 33, 34, 41, 44, 65
- MRI** Magnetic Resonance Imaging. 5, 11, 12, 25
- MU** Monitor Unit. i, ii, 14, 26, 29, 33–37, 40–44, 47, 48, 62, 68, 70
- NLP** Natural Language Processing. 20, 66

- NN** Neural Network. 2, 15, 17
- OAR** Organ At Risk. 12, 15, 38, 49, 58–61, 63, 65, 70, 71
- PB** Pencil Beam. 7–10, 69
- PET** Positron Emission Tomography. 5, 11, 12
- PSA** Prostate-Specific Antigen. 11
- PTV** Planning Target Volume. i, ii, 26, 33, 38, 52, 53, 62, 63
- QA** Quality Assurance. 1, 12, 37, 64, 65
- ReLU** Rectified Linear Unit. 31–33
- ResNet** Residual Network. 17, 18, 29, 31, 34, 62
- RL** Reinforcement Learning. 63
- RNN** Recurrent Neural Network. 15, 20
- RT** Radiotherapy. 22, 36, 48, 62, 63, 67, 68
- TPS** Treatment Planning System. i, ii, 1, 2, 11, 12, 14, 22, 24, 25, 27, 37, 43, 48, 63, 64
- TRUS** Transrectal Ultrasound. 11
- VMAT** Volumetric Modulated Arc Therapy. i, 2, 5, 12, 13, 24, 26, 62
- YOLO** You Only Look Once. 19, 30–33, 35, 36, 40, 41, 43, 48, 52, 57, 58, 61–63, 65, 70

References

- [1] G. Heilemann, L. Zimmermann, R. Schotola, W. Lechner, M. Peer, J. Widder, G. Goldner, D. Georg, and P. Kuess. “Generating deliverable DICOM RT treatment plans for prostate VMAT by predicting MLC motion sequences with an encoder-decoder network”. In: *Medical Physics* 50 (8 2023), pp. 5088–5094. DOI: 10.1002/mp.16545.
- [2] G. Heilemann, L. Zimmermann, T. Nyholm, A. Simkó, J. Widder, G. Goldner, D. Georg, and P. Kuess. “Ultra-fast, one-click radiotherapy treatment planning outside a treatment planning system”. In: *Physics and Imaging in Radiation Oncology* 33 (2025). DOI: 10.1016/j.phro.2025.100724.
- [3] N. D. James, I. Tannock, J. N’Dow, F. Feng, S. Gillesen, S. A. Ali, B. Trujillo, B. Al-Lazikani, G. Attard, F. Bray, E. Compérat, R. Eeles, O. Fatiregun, E. Grist, S. Halabi, Á. Haran, D. Herchenhorn, M. S. Hofman, M. Jalloh, S. Loeb, A. MacNair, B. Mahal, L. Mendes, M. Moghul, C. Moore, A. Morgans, M. Morris, D. Murphy, V. Murthy, P. L. Nguyen, A. Padhani, C. Parker, H. Rush, M. Sculpher, H. Soule, M. R. Sydes, D. Tilki, N. Tunariu, P. Villanti, and L. P. Xie. “The Lancet Commission on prostate cancer: planning for the surge in cases”. In: *The Lancet* 403 (10437 2024), pp. 1683–1722. DOI: 10.1016/S0140-6736(24)00651-2.
- [4] World Cancer Research Fund International. *Prostate Cancer Statistics*. <https://www.wcrf.org/cancer-trends/prostate-cancer-statistics/>. 2024. Accessed: 18.06.2025.
- [5] E. Huynh, A. Hosny, C. Guthier, D. S. Bitterman, S. F. Petit, D. A. Haas-Kogan, B. Kann, H. J. Aerts, and R. H. Mak. “Artificial intelligence in radiation oncology”. In: *Nature Reviews Clinical Oncology* 17 (12 2020), pp. 771–781. DOI: 10.1038/s41571-020-0417-8.
- [6] A. Babier, R. Mahmood, B. Zhang, V. G. Alves, A. M. Barragán-Montero, J. Beaudry, C. E. Cardenas, Y. Chang, Z. Chen, J. Chun, K. Diaz, H. D. Eraso, E. Faustmann, S. Gaj, S. Gay, M. Gronberg, B. Guo, J. He, G. Heilemann, S. Hira, Y. Huang, F. Ji, D. Jiang, J. C. J. Giraldo, H. Lee, J. Lian, S. Liu, K. C. Liu, J. Marrugo, K. Miki, K. Nakamura, T. Netherton, D. Nguyen, H. Nourzadeh, A. F. Osman, Z. Peng, J. D. Q. Muñoz, C. Ramsel, D. J. Rhee, J. D. Rodriguez, H. Shan, J. V. Siebers, M. H. Soomro, K. Sun, A. U. Hoyos, C. Valderrama, R. Verbeek, E. Wang, S. Willems, Q. Wu, X. Xu, S. Yang, L. Yuan, S. Zhu, L. Zimmermann, K. L. Moore, T. G. Purdie, A. L. McNiven, and T. C. Chan. “OpenKBP-Opt: An international and reproducible evaluation of 76 knowledge-based planning pipelines”. In: *Physics in Medicine and Biology* 67 (18 2022). DOI: 10.1088/1361-6560/ac8044.
- [7] D. Nguyen, X. Jia, D. Sher, M. H. Lin, Z. Iqbal, H. Liu, and S. Jiang. “3D radiotherapy dose prediction on head and neck cancer patients with a hierarchically densely connected U-net deep learning architecture”. In: *Physics in Medicine and Biology* 64 (6 2019). DOI: 10.1088/1361-6560/ab039b.

- [8] J. Fan, J. Wang, Z. Chen, C. Hu, Z. Zhang, and W. Hu. “Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique”. In: *Medical Physics* 46 (1 2019), pp. 370–381. DOI: 10.1002/mp.13271.
- [9] Y. Sun, Z. Sun, and W. Chen. “The evolution of object detection methods”. In: *Engineering Applications of Artificial Intelligence* 133 (2024). DOI: 10.1016/j.engappai.2024.108458.
- [10] N. G. Zaorsky, T. Shaikh, C. T. Murphy, M. A. Hallman, S. B. Hayes, M. L. Sobczak, and E. M. Horwitz. “Comparison of outcomes and toxicities among radiation therapy treatment options for prostate cancer”. In: *Cancer Treatment Reviews* 48 (2016), pp. 50–60. DOI: 10.1016/j.ctrv.2016.06.006.
- [11] National Cancer Institute. *Brachytherapy to Treat Cancer*. <https://www.cancer.gov/about-cancer/treatment/types/radiation-therapy/brachytherapy>. 2019. Accessed: 18.06.2025.
- [12] Prostate Cancer Foundation. *Radiation Therapy for Prostate Cancer*. <https://www.pcf.org/patient-support/treatment/localized-treatment/radiation-therapy/>. 2024. Accessed: 18.06.2025.
- [13] K. Koka, A. Verma, B. S. Dwarakanath, and R. V. Papineni. “Technological Advancements in External Beam Radiation Therapy (EBRT): An Indispensable Tool for Cancer Treatment”. In: *Cancer Management and Research* 14 (2022), pp. 1421–1429. DOI: 10.2147/CMAR.S351744.
- [14] P. Mayles. “From Measurements to Calculations”. In: *Handbook of Radiotherapy Physics*. Ed. by P. Mayles, A. Nahum, and J.-C. Rosenwald. Taylor & Francis, 2007. Chap. 20, pp. 433–438.
- [15] International Atomic Energy Agency. *Radiation Oncology Physics*. Non-serial Publications. Vienna, 2005.
- [16] A. Ahnesjö and M. M. Aspradakis. “Dose calculations for external photon beams in radiotherapy”. In: *Physics in Medicine and Biology* 44 (11 1999), R99–R155. DOI: 10.1088/0031-9155/44/11/201.
- [17] J.-C. Rosenwald. “Principles of Patient Dose Computation”. In: *Handbook of Radiotherapy Physics*. Ed. by P. Mayles, A. Nahum, and J.-C. Rosenwald. Taylor & Francis, 2007. Chap. 25, pp. 549–559.
- [18] J.-C. Rosenwald, I. Rosenberg, and G. Shentall. “Patient Dose Computation for Photon Beams”. In: *Handbook of Radiotherapy Physics*. Ed. by P. Mayles, A. Nahum, and J.-C. Rosenwald. Taylor & Francis, 2007. Chap. 26, pp. 559–587.
- [19] E. J. Schafer, M. Laversanne, H. Sung, I. Soerjomataram, A. Briganti, W. Dahut, F. Bray, and A. Jemal. “Recent Patterns and Trends in Global Prostate Cancer Incidence and Mortality: An Update.” In: *European urology* (2024). DOI: 10.1016/j.eururo.2024.11.013.
- [20] G. Goldner. *Therapieprotokoll Primäres Prostata - Karzinom*. Department of Radiation Oncology (AKH Vienna, MedUni Vienna). Unpublished. 2022.

- [21] RaySearch Laboratories. *RayStation Treatment Planning System*. <https://www.raysearchlabs.com/raystation/>. 2024. Accessed: 18.06.2025.
- [22] Elekta. *Radiation Therapy*. <https://www.elekta.com/products/radiation-therapy/>. 2024. Accessed: 18.06.2025.
- [23] M. Peer. *Generating Treatment Plans for Prostate VMAT using Deep Learning (MSc thesis)*. 2021.
- [24] D. Dearnaley, I. Syndikus, H. Mossop, V. Khoo, A. Birtle, D. Bloomfield, J. Graham, P. Kirkbride, J. Logue, Z. Malik, J. Money-Kyrle, J. M. O’Sullivan, M. Panades, C. Parker, H. Patterson, C. Scrase, J. Staffurth, A. Stockdale, J. Tremlett, M. Bidmead, H. Mayles, O. Naismith, C. South, A. Gao, C. Cruickshank, S. Hassan, J. Pugh, C. Griffin, and E. Hall. “Conventional versus hypofractionated high-dose intensity-modulated radiotherapy for prostate cancer: 5-year outcomes of the randomised, non-inferiority, phase 3 CHHiP trial”. In: *The Lancet Oncology* 17 (8 2016), pp. 1047–1060. DOI: 10.1016/S1470-2045(16)30102-4.
- [25] Y. Lecun, Y. Bengio, and G. Hinton. “Deep learning”. In: *Nature* 521 (7553 2015), pp. 436–444. DOI: 10.1038/nature14539.
- [26] M. Wang, Q. Zhang, S. Lam, J. Cai, and R. Yang. “A Review on Application of Deep Learning Algorithms in External Beam Radiotherapy Automated Treatment Planning”. In: *Frontiers in Oncology* 10 (2020). DOI: 10.3389/fonc.2020.580919.
- [27] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016.
- [28] C. M. Bishop and H. Bishop. *Deep Learning*. Springer International Publishing, 2024. DOI: 10.1007/978-3-031-45468-4.
- [29] V. Dumoulin and F. Visin. *A guide to convolution arithmetic for deep learning*. 2018. arXiv: 1603.07285 [stat.ML].
- [30] N. Lang. *Autoencoders: An Ultimate Guide for Data Scientists*. <https://towardsdatascience.com/autoencoders-an-ultimate-guide-for-data-scientists-dca3e56a070e/>. 2024. Accessed: 18.06.2025.
- [31] K. He, X. Zhang, S. Ren, and J. Sun. *Deep Residual Learning for Image Recognition*. 2015. arXiv: 1512.03385 [cs.CV].
- [32] O. Ronneberger, P. Fischer, and T. Brox. “U-net: Convolutional networks for biomedical image segmentation”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 9351. Springer Verlag, 2015, pp. 234–241. DOI: 10.1007/978-3-319-24574-4_28.
- [33] F. Nazem, F. Ghasemi, A. Fassihi, and A. M. Dehnavi. “3D U-Net: A voxel-based method in binding site prediction of protein structure”. In: *Journal of Bioinformatics and Computational Biology* 19 (2 2021), p. 2150006. DOI: 10.1142/S0219720021500062.

- [34] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. *You Only Look Once: Unified, Real-Time Object Detection*. 2016. arXiv: 1506.02640 [cs.CV].
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. *Attention Is All You Need*. 2023. arXiv: 1706.03762 [cs.CL].
- [36] R. Merritt. *What Is a Transformer Model?* <https://blogs.nvidia.com/blog/what-is-a-transformer-model/>. 2022. Accessed: 18.06.2025.
- [37] A. Sadeghnejad-Barkousaraie, O. Ogunmolu, S. Jiang, and D. Nguyen. “A fast deep learning approach for beam orientation optimization for prostate cancer treated with intensity-modulated radiation therapy”. In: *Medical Physics* 47 (3 2020), pp. 880–897. DOI: 10.1002/mp.13986.
- [38] A. Sadeghnejad-Barkousaraie, G. Bohara, S. Jiang, and D. Nguyen. “A reinforcement learning application of a guided Monte Carlo tree search algorithm for beam orientation selection in radiation therapy”. In: *Machine Learning: Science and Technology* 2 (3 2021). DOI: 10.1088/2632-2153/abe528.
- [39] X. Li, C. Wang, Y. Sheng, J. Zhang, W. Wang, F.-F. Yin, Q. Wu, Q. J. Wu, and Y. Ge. “An artificial intelligence-driven agent for real-time head-and-neck IMRT plan generation using conditional generative adversarial network (cGAN)”. In: *Medical Physics* 48 (6 2021), pp. 2714–2723. DOI: 10.1002/mp.14770.
- [40] L. Zimmermann, E. Faustmann, C. Ramsel, D. Georg, and G. Heilemann. “Technical Note: Dose prediction for radiation therapy using feature-based losses and One Cycle Learning”. In: *Medical Physics* 48 (9 2021), pp. 5562–5566. DOI: 10.1002/mp.14774.
- [41] Y. Ni, S. Chen, L. Hibbard, and P. Voet. “Fast VMAT planning for prostate radiotherapy: dosimetric validation of a deep learning-based initial segment generation method”. In: *Physics in Medicine and Biology* 67 (15 2022). DOI: 10.1088/1361-6560/ac80e5.
- [42] R. Schotola. *Deep Learning based Prediction of Deliverable Treatment Plans in Radiation Oncology (MSc thesis)*. 2023.
- [43] Medical Imaging & Technology Alliance. *Digital Imaging and Communications in Medicine. Supplement 11: Radiotherapy Objects*. <https://www.dicomstandard.org/News-dir/ftsup/docs/sups/sup11.pdf>. 1997. Accessed: 18.06.2025.
- [44] Varian. *Radiotherapy*. <https://www.varian.com/products/radiotherapy>. 2024. Accessed: 18.06.2025.
- [45] T. Nyholm, J. Olofsson, A. Ahnesjö, and M. Karlsson. “Photon pencil kernel parameterisation based on beam quality index”. In: *Radiotherapy and Oncology* 78 (3 2006), pp. 347–351. DOI: 10.1016/j.radonc.2006.02.002.
- [46] X. Chen, C. Liang, D. Huang, E. Real, K. Wang, Y. Liu, H. Pham, X. Dong, T. Luong, C.-J. Hsieh, Y. Lu, and Q. V. Le. *Symbolic Discovery of Optimization Algorithms*. 2023. arXiv: 2302.06675 [cs.LG].

- [47] W. Falcon and The PyTorch Lightning team. *PyTorch Lightning*. Version 1.4. 2019. DOI: 10.5281/zenodo.3828935.
- [48] J. Ansel, E. Yang, H. He, N. Gimselshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski, G. Chauhan, A. Chourdia, W. Constable, A. Desmaison, Z. DeVito, E. Ellison, W. Feng, J. Gong, M. Gschwind, B. Hirsh, S. Huang, K. Kalambarkar, L. Kirsch, M. Lazos, M. Lezcano, Y. Liang, J. Liang, Y. Lu, C. Luk, B. Maher, Y. Pan, C. Puhersch, M. Reso, M. Saroufim, M. Y. Siraichi, H. Suk, M. Suo, P. Tillet, E. Wang, X. Wang, W. Wen, S. Zhang, X. Zhao, K. Zhou, R. Zou, A. Mathews, G. Chanan, P. Wu, and S. Chintala. “PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation”. In: *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS ’24)*. ACM, 2024. DOI: 10.1145/3620665.3640366.
- [49] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Y. Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015.
- [50] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. “Array programming with NumPy”. In: *Nature* 585 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- [51] D. Mason, scaramallion, mrbean-bremen, rhaxton, J. Suever, D. P. Orfanos, Vanes-sasaurus, G. Lemaitre, A. Panchal, A. Rothberg, M. D. Herrmann, J. Massich, J. Kerns, K. van Golen, C. Bridge, S. Biggs, T. Robitaille, moloney, M. Shun-Shin, B. Conrad, pawelzajdel, M. Mattes, Y. Lyu, T. Cogan, Z. Baratz, F. C. Morency, Taylor, and T. Sentner. *pydicom/pydicom: pydicom 3.0.1*. Version v3.0.1. 2024. DOI: 10.5281/zenodo.13824606.
- [52] L. Biewald. *Experiment Tracking with Weights and Biases*. Software available from wandb.com. 2020. Accessed: 18.06.2025.
- [53] M. J. Gooding, S. Aluwini, T. G. Urbano, Y. McQuinlan, D. Om, F. H. Staal, T. Perennec, S. Azzarouali, C. E. Cardenas, A. Carver, S. S. Korreman, and J. E. Bibault. “Fully automated radiotherapy treatment planning: A scan to plan challenge”. In: *Radiotherapy and Oncology* 200 (2024). DOI: 10.1016/j.radonc.2024.110513.

- [54] C. Church, M. Yap, M. Bessrou, M. Lamey, and D. Granville. “Automated plan generation for prostate radiotherapy patients using deep learning and scripted optimization”. In: *Physics and Imaging in Radiation Oncology* 32 (2024). DOI: 10.1016/j.phro.2024.100641.
- [55] J. Harms, J. A. Pogue, C. E. Cardenas, D. N. Stanley, R. Cardan, and R. Popple. “Automated evaluation for rapid implementation of knowledge-based radiotherapy planning models”. In: *Journal of Applied Clinical Medical Physics* 24 (10 2023). DOI: 10.1002/acm2.14152.
- [56] Z. Lou, C. Cheng, R. Mao, D. Li, L. Tian, B. Li, H. Lei, and H. Ge. “A novel automated planning approach for multi-anatomical sites cancer in Raystation treatment planning system”. In: *Physica Medica* 109 (2023). DOI: 10.1016/j.ejmp.2023.102586.
- [57] X. Li, J. Zhang, Y. Sheng, Y. Chang, F. F. Yin, Y. Ge, Q. J. Wu, and C. Wang. “Automatic IMRT planning via static field fluence prediction (AIP-SFFP): A deep learning algorithm for real-time prostate treatment planning”. In: *Physics in Medicine and Biology* 65 (17 2020). DOI: 10.1088/1361-6560/aba5eb.
- [58] Z. Yuan, Y. Wang, P. Hu, D. Zhang, B. Yan, H. M. Lu, H. Zhang, and Y. Yang. “Accelerate treatment planning process using deep learning generated fluence maps for cervical cancer radiation therapy”. In: *Medical Physics* 49 (4 2022), pp. 2631–2641. DOI: 10.1002/mp.15530.
- [59] W. T. Hrinivich, M. Bhattacharya, L. Mekki, T. McNutt, X. Jia, H. Li, D. Y. Song, and J. Lee. “Clinical VMAT machine parameter optimization for localized prostate cancer using deep reinforcement learning”. In: *Medical Physics* 51 (6 2024), pp. 3972–3984. DOI: 10.1002/mp.17100.
- [60] C. Shen, L. Chen, Y. Gonzalez, and X. Jia. “Improving efficiency of training a virtual treatment planner network via knowledge-guided deep reinforcement learning for intelligent automatic treatment planning of radiotherapy”. In: *Medical Physics* 48 (4 2021), pp. 1909–1920. DOI: 10.1002/mp.14712.
- [61] L. A. Künzel, S. Leibfarth, O. S. Dohm, A. C. Müller, D. Zips, and D. Thorwarth. “Automatic VMAT planning for post-operative prostate cancer cases using particle swarm optimization: A proof of concept study”. In: *Physica Medica* 69 (2020), pp. 101–109. DOI: 10.1016/j.ejmp.2019.12.007.
- [62] C. Gianoli, E. D. Bernardi, and K. Parodi. “‘Under the hood’: artificial intelligence in personalized radiotherapy”. In: *BJR/Open* 6 (1 2023). DOI: 10.1093/bjro/tzae017.
- [63] J. Terven, D. M. Córdova-Esparza, and J. A. Romero-González. “A Comprehensive Review of YOLO Architectures in Computer Vision: From YOLOv1 to YOLOv8 and YOLO-NAS”. In: *Machine Learning and Knowledge Extraction* 5 (4 2023), pp. 1680–1716. DOI: 10.3390/make5040083.

- [64] C. Hu, H. Wang, W. Zhang, Y. Xie, L. Jiao, and S. Cui. “TrDosePred: A deep learning dose prediction algorithm based on transformers for head and neck cancer radiotherapy”. In: *Journal of Applied Clinical Medical Physics* 24 (7 2023). DOI: 10.1002/acm2.13942.
- [65] Z. Jiao, X. Peng, Y. Wang, J. Xiao, D. Nie, X. Wu, X. Wang, J. Zhou, and D. Shen. “TransDose: Transformer-based radiotherapy dose prediction from CT images guided by super-pixel-level GCN classification”. In: *Medical Image Analysis* 89 (2023). DOI: 10.1016/j.media.2023.102902.

Eidesstattliche Erklärung

Hiermit erkläre ich, daß ich die vorliegende Arbeit selbstständig angefertigt, nicht anderweitig zu Prüfungszwecken vorgelegt und keine anderen als die angegebenen Hilfsmittel verwendet habe. Sämtliche wissentlich verwendete Textausschnitte, Zitate oder Inhalte anderer Verfasser wurden ausdrücklich als solche gekennzeichnet. Diese Arbeit wurde bisher weder im In- noch im Ausland in gleicher oder in ähnlicher Form in anderen Prüfungsverfahren vorgelegt.

Wien, 23.06.2025

Simon Glatzer