Towards EPID-based 3D in vivo Dosimetry for Modern Radiation Therapy

Juliana Cristina Martins

München 2023

Towards EPID-based 3D in vivo Dosimetry for Modern Radiation Therapy

Juliana Cristina Martins

Dissertation

an der Fakultät für Physik

der Ludwig-Maximilians-Universität München

vorgelegt von

Juliana Cristina Martins

geboren in Londrina

München, den 25. April 2023

Erstgutachterin: Prof. Dr. Katia Parodi
Zweitgutachter: Prof. Dr. Jan Wilkens
Tag der mündlichen Prüfung: 07.06.2023

Zusammenfassung

Moderne Verfahren der Strahlentherapie, wie beispielsweise die intensitätsmodulierte Strahlentherapie (IMRT) oder Volumen Intensitätsmodulierte Arc Therapie (VMAT), sind in der Lage Dosisverteilungen abzubilden, welche hohe Konformität und steile Gradienten zwischen dem zu bestrahlendem Tumorgewebe und umliegenden, zu schonenden Organen bieten. Diese Verfahren stellen hohe Anforderungen an Qualitätssicherung und Dosisverifikation sowohl in der Vorbereitung aber auch während der Bestrahlung des Patienten (in vivo). Diese Arbeit stellt eine Methode der EPID-basierten in vivo Dosimetrie vor, welches die Genauigkeiten der Monte-Carlo (MC) Methoden für Dosissimulationen in gegebenen Patientengeometrien mit der Zeiteffizienz tiefer neuronaler Netze verbindet. Das "Deep Dose Estimation" (DDE) Netzwerk, ursprünglich zur Abschätzung von Dosisverteilungen für Computertomographie (CT) Aufnahmen entwickelt, wurde erweitert und trainiert um 3D Dosisverteilungen vorherzusagen, welche durch IMRT Bestrahlungsfelder innerhalb des Patienten entstehen. Die Genauigkeit ist vergleichbar zu MC-Methoden.

Das DDE verwendet eine CT Aufnahme des Patienten und eine angenäherte Dosisverteilung als Eingabe, welche "First Order Dose approximation" (Dosisverteilung erster Ordnung, bzw. FOD) genannt wird und aus simulierten EPID Messungen erstellt wird. Das Netzwerk wurde trainiert um diese beiden Eingaben in genaue Dosisverteilungen ("Accurate Dose Distribution", ADD) zu übertragen. Diese Dosisverteilungen wurden auf dem selben Patienten CT mit MC Methoden simuliert. Die FODs sind vereinfachte 3D Dosisverteilungen, die durch Rückprojektion der simulierten EPID Signale unter Berücksichtigung ihrer Magnifikation, ihrer Inverse-Square-Law Korrekturen, und Schwächung innerhalb des virtuellem Patientenmodells erzeugt werden. Die FODs berücksichtigen jedoch nicht alle bekannten Umstände. So wird unter anderem der Aufbaueffekt, die Strahlhärtung und die Streuung innerhalb des Patienten nicht im FOD, jedoch im ADD entsprechend berücksichtigt.

Daraus folgt, dass diese Methodik stark abhängig ist von den Ergebnissen des MC Modells, welches sowohl ADD als auch das übertragene EPID Signal errechnet. Ein verlässliches MC Modell eines Linearbeschleunigers wurde für diese Arbeit eigens erstellt und ausgiebig validiert. Der patientenspeziefische Teil des Linearbeschleunigerkopfes, namentlich der Multilamellenkollimator (MLC), wurde nur auf Grundlage verfügbarer Literatur angelegt. Ein virtuelles Modell des EPIDs wurde ebenfalls im patientenspezifischen Teil

eingebunden um gleichzeitig das transmittierte Signal durch den virtuellen Patienten aufzunehmen. Der patientenunabhängige Teil, das sind die statischen Teile des simulierten Linearbeschleunigerkopfes, wurde auf Grundlage von zur Verfügung gestellten vertraulichen Herstellerinformationen erzeugt. Dieser Teil wurde verwendet um Phasenraumdateien (PhSp) zu generieren. Diese PhSp Dateien wurden anschließend als Primärteilchenerzeuger verwendet um ADDs und EPID Signale zu simulieren. Eine alternative Methodik um existierende IAEA PhSp Daten zu optimieren, ist im Rahmen eines Nebenprojektes entstanden und kann ebenfalls als patientenunabhängiger Teil des Linerbeschleunigerkopfes verwendet werden, falls die andernfalls benötigten vertraulichen Produktinformationen nicht verfügbar sind.

Die ADDs für klinische Prostata IMRT Bestrahlungsfelder, und die entsprechenden übertragenen EPID Signale wurde mit 83 Becken-CT-Aufnahmen simuliert mit Gantry bei 0°. Insgesamt wurden 581 unterschiedliche ADD-FOD Kombinationen mit jeweils 7 verschiedenen Bestrahlungsfeldern pro Patient erzeugt. Das DDE Netzwerk wurde mit 67 Patientendatensätzen trainiert (Trainigsdatensatz). Die verbleibenden 16 Patientendatensätze wurden zur Validierung verwendet (Testdatensatz). Ein weiterer Datensatz mit 8 Bestrahlungsfeldern bei Gantry 90° (Lateraldatensatz) wurde zur Auswertung der Performance des DDE bei unterschiedlichen Bestrahlungsrichtungen benutzt.

Die Übereinstimmung der durch das DDE vorhergesagten Dosisverteilungen (DDEP) für die Validierungs- und Lateraldatensätze mit den ADD wurde mithilfe der Gamma Analyse (3%, 2 mm Kriterium) quantifiziert. Um die Verbesserungen durch das DDE Netzwerk auszuwerten wurde dieselbe Gamma Analyse auch auf die FODs und entsprechenden ADDs angewandt. Die Gammaakzeptanzraten zwischen FODs und ADDs waren bei 46%, die Akzeptanzraten für die DDEPs waren dahingegen bei über 97% für alle Bestrahlungsfelder des Testsatzes. Für den Lateraldatensatz konnte das DDE die Akzeptanzraten von 88% auf über 95% zu steigern. Die hohen Akzeptanzraten für DDEPs unterstreichen, dass das DDE unter angemessener Berücksichtigung aller fehlender Effekte in der Lage ist, FODs in ADDs zu konvertieren. Darüber hinaus können Dosen innerhalb des Patienten CTs in 0,6 Sekunde pro Bestrahlungsfeld (Berechnungen mit GPU) vorhergesagt werden, sobald das DDE einmal trainiert ist. Dies stellt eine deutliche Verbesserung gegenüber 14 Stunden Berechnungszeiten für MC Simulationen (Berechnungen auf CPU Clustern) dar. Die zugestellte Dosis für eine vollständige Prostatabestrahlungsitzung eines Patienten kann damit in weniger als einer Minute vorhergesagt werden. 3D in Vivo Dosisverteilungen bei klinischen Patientenbestrahlungen können mit der hier vorgestellten Methodik innerhalb von Sekunden erzeugt werden, was ein möglicher wichtiger Baustein auf dem Weg zu einer klinisch verwendbaren EPID-basierten in Vivo Echtzeitdosimetrie ist.

Abstract

Modern radiotherapy techniques, such as Intensity Modulated Radiation Therapy (IMRT) and Volumetric Modulated Arc Therapy (VMAT), can deliver highly conformal dose distributions, with steep dose gradients between the target and organs at risk. This increases the demands on proper quality assurance and dose verification before (pre-treatment) and during (in vivo) patient irradiation. This project proposes a methodology for EPID-based in vivo dosimetry, combining the accuracy of Monte Carlo (MC) methods for dose simulation in patient geometry, with the time-efficiency of deep neural networks. The Deep Dose Estimation (DDE) network, originally developed for dose estimation in radiological computed tomography (CT) exams, has been extended and trained to predict 3D dose distributions due to IMRT fields, inside a patient, with accuracy comparable to MC methods.

The DDE uses as input a patient CT image and an approximated dose distribution, called first order dose approximation (FOD), reconstructed from simulated EPID signals. The network was trained to map this two-channel input to an accurate dose distribution (ADD) inside the same patient CT, simulated using MC methods. The FODs are simplified 3D dose distributions produced as backprojections of the simulated EPID signals, accounting for magnification and inverse square law corrections, and attenuation through the virtual patient model. The FODs do not account for several effects, such as the build-up, beam hardening and scattering within the patient, all of which were properly considered in the ADDs.

Hence, the methodology relies strongly on the MC model used to produce both the ADD and the transmitted EPID signals. A reliable MC model of the linac considered in this work was constructed and extensively validated. The patient-dependent part of the linac head, namely the multi-leaf collimator (MLC) system, was produced based entirely on information available in the literature. A virtual model of the EPID was also included in the patient-dependent part, to simultaneously record the transmitted signal through the virtual patient. The patient-independent part, i.e. the static parts of the linac head, was constructed based on confidential information provided by the vendor, and used to produce phase space (PhSp) files. These PhSp files were subsequently used as primary particle generators to simulate the ADDs and EPID signals. An alternative methodology for optimization of existing IAEA PhSp files was developed as a side project, which can be used to model the patient-independent part of the linac head when confidential vendor information is not available.

The ADDs for clinical prostate IMRT fields, and respective transmitted EPID signals, were simulated inside 83 pelvic CTs, with gantry at 0°. In total, 581 different ADD-FOD sets were produced, with seven different fields per patient CT. The network was trained using the data sets of 67 patients (training set). The data of the remaining 16 patients were used for validation (test set). An additional dataset with eight fields simulated with gantry at 90° (lateral set) was used for evaluating the performance of the trained DDE for other irradiation directions.

The quality of the DDE-predicted dose distributions (DDEP) on the test and lateral sets was quantified in terms of the gamma analysis with respect to the ADD (3%, 2 mm criteria). To evaluate the improvement obtained with the DDE, the same evaluation was performed for FODs and respective ADDs. The gamma passing rates between FODs and ADDs were as low as 46%, while for DDEPs the passing rates were above 97% for all fields on the test set. For the fields in the lateral set, the DDE was able to improve the passing rates from 88% to above 95%. The high passing rates for DDEPs indicate that the DDE was able to convert the FODs into ADDs, properly accounting for all missing effects. Moreover, once trained, the DDE can predict the dose inside a patient CT within 0.6 s per field (using a GPU), in contrast to 14 h needed for MC simulations (using a CPU-cluster). The dose delivered to a patient due to an entire prostate treatment session can therefore be predicted in less than one minute. With the proposed methodology, 3D in vivo dose distributions due to clinical patient irradiation can be obtained within seconds, potentially paving the way towards a clinically viable, real-time EPID-based in vivo dosimetry.

Contents

List of Figures

List of Tables

CHAPTER 1

Introduction and Motivation

Since the discovery of X-rays by Wilhelm Conrad Röntgen in 1895 [1], the natural radioactivity by Henri Becquerel in 1896, and the isolation of radium and polonium by Marie and Pierre Curie in 1898, the field of radiation science has evolved rapidly [2]. Within months after their discovery, X-rays were already being used for diagnostic purposes [3], and the successful use of radioactivity to treat cancer was reported in 1899 by T. Stenbeck and T. Sjögren [4]. Not long after, in the early 20^{th} century, it was discovered that radiation could also be harmful [5], and therefore should not be used indiscriminately.

In this context, the focus of radiation therapy is to eliminate cancerous tissue, usually called targets, while minimizing the potential harmful irradiation of surrounding healthy organs and tissue, also referred to as organs at risk (OARs). Ideally, the target should be irradiated with a prescribed dose without affecting the surrounding OARs. With conventional radiotherapy, performed either with photons or electron beams, this is not possible for several reasons, but mainly because of the dose deposition[1] at different depths of a tissue, characterized by the depth dose curve or percentage depth dose (PDD) profile. As can be seen in Fig. 1.1, photons and electrons deposit energy increasingly until a maximum value (at the depth of maximum dose d_{max}), beyond which the deposition follows approximately an exponential decrease, steeper for electrons than for photons. This means that in order to irradiate a tumor at a certain depth, the surrounding tissue before and after the tumor depth will also receive some dose. Moreover, in external beam radiotherapy, the beam penetrates healthy tissue in order to reach the target. This problem could be alleviated by using particle therapy with e.g. protons or carbon ions to irradiate the tumor, as protons and ions release most of their energies at a fixed energy-dependent depth, the so-called Bragg peak, hence better sparing the healthy tissue before and after the tumor compared to electromagnetic radiation (Fig. 1.1).

---

[1] Radiation dose is the mean energy dE imparted by any ionizing radiation to the matter of mass dm in a finite volume V: $D = dE/dm$, given in gray Gy (Gy = J/kg)

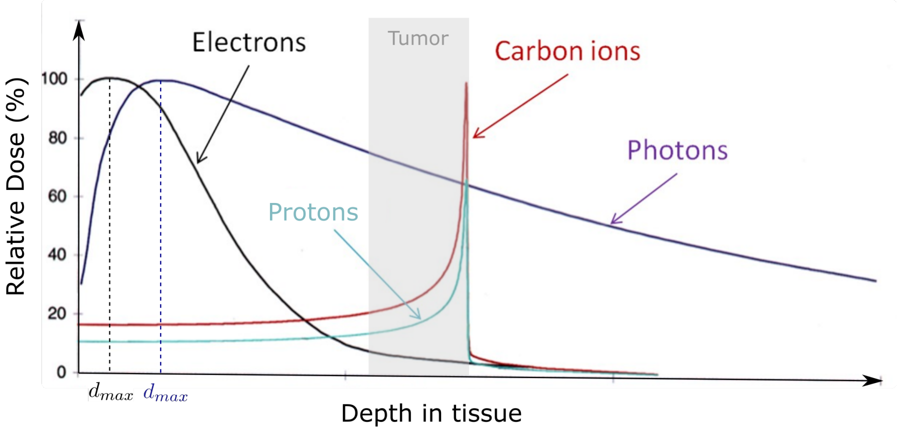

Figure 1.1: Depth dose deposition curve, or percentage depth dose (PDD) profile for photons (dark blue), electrons (black), protons (light blue) and carbon ions (red). Modified and reprinted with permission from S. Liprandi [6].

Particle therapy has clear advantages regarding dose deposition with respect to photon radiotherapy, therefore the interest in the former has increased in the last decades. By January 2023, around 107 proton and 14 carbon ion therapy centers are operational in 21 countries, and 35 new facilities are under construction [7]. However, these particle therapy centers are mainly distributed in developed countries, with 74 of them located in North America and Europe, while there is only one facility under development in South America (Argentina), and one facility in a planning stage in Africa (Egypt) [7]. In contrast, photon radiotherapy centers are present at over 156 different countries, in 7736 facilities with over 15000 operating machines [8], indicating that, despite the advantages and rapid expansion of particle therapy, the great majority of patients worldwide are treated with photon therapy. Moreover, the costs for particle therapy are considerably higher than the costs for photon therapy [9, 10], which might be an issue in implementing particle therapy in developing countries. Therefore, studies and research towards improving the quality of conventional photon radiotherapy are still very relevant and should be encouraged, as it affects a large number of patients.

With over 100 years of existence, photon radiation therapy has evolved and become more conformal to the target tumor volume. Modern treatment techniques, such as Intensity Modulated Radiation Therapy (IMRT) and Volumetric Modulated Arc Therapy (VMAT), deliver highly conformal dose distributions, aiming at improving coverage of the target volumes while sparing the surrounding normal tissue and OARs. These requirements result in steep dose gradients between the target and the OARs. Hence, the quality of the treatment depends on highly accurate machine operation, proper patient positioning during irradiation and reproducible patient anatomy on each treatment fraction, increasing the demands for quality assurance of the clinical linear accelerators (linac) and dosimetric verifications.

To verify the proper functioning of the equipment prior to patient irradiation, several pre-treatment verification protocols have been established [11–14]. However, pre-treatment verifications are limited and cannot detect machine malfunctions during patient irradiation, errors in patient positioning, or changes in the patient anatomy in the course of treatment, all of which can seriously influence the final dose distribution. Therefore, despite its great value, pre-treatment verification cannot guarantee the technical success of patient irradiation [15].

1.1 EPID-based in vivo dosimetry

In the context of modern high precision and adaptive photon therapy, the verification of the dose delivered at each treatment fraction has gained increased attention, and much effort has been put into developing in vivo dosimetry methodologies that can be integrated into the clinical practice. Some approaches use point detectors such as plastic scintillators, thermoluminescent and optically stimulated detectors [16–19], that offer only point dose measurements, mostly at the skin and/or at beam entrance/exit, which limits their use for 3D dose verification [18, 20, 21]. Other groups propose introducing point detectors into the tumor or OARs [22, 23], however these approaches are subject to large uncertainties due to detectors' mispositioning and/or changes in patient anatomy, which combined with their invasiveness, make them difficult to implement in a clinical routine. To overcome the limitations of point detectors, other techniques involving film dosimetry [24, 25] and two-dimensional detector arrays [26, 27] have been investigated, but no other method has gained more attention than the dosimetric use of portal dose imagers.

Electronic Portal Imaging Devices (EPID) are an intrinsic part of modern clinical linacs, primarily developed for patient positioning verification. After their clinical implementation in the late 1980s, it was quickly realized that EPIDs could be useful for dosimetric verifications, encouraging several groups to investigate the dosimetric properties of different types of EPID [28]. A detailed overview of the development of the EPID technology is out of the scope of this thesis, but more information can be found in Boyer et al. [29], L. Antonuk [30] and Kirby and Glendinning [31]. Since the beginning of the 21st century, amorphous-silicon (a-Si) EPIDs are the most common type of EPID available in the clinics, and their properties will be described in detail elsewhere in this thesis.

The use of modern a-Si EPIDs as a dosimetric tool has been encouraged due to their sub-millimeter spatial resolution, temporal resolution in the order of milliseconds, approximate linear response to radiation, digital and real-time readout, and large sensitive area [28, 32, 33], with several methods proposed for both pre-treatment and in vivo patient-specific verification. This thesis will focus on in vivo methodologies. For an overview of current pre-treatment EPID-based methods, please refer to [28, 32, 34–37] and references therein.

EPID in vivo dosimetry can be divided into forward and backprojection approaches [38]. In the forward approach, the measured EPID signal is compared to a predicted

EPID signal, at the EPID level, either as an intensity image or as calibrated dose-to-water distributions [39–43]. The predicted signal is obtained by forward dose calculation engines, based on analytical or Monte Carlo (MC) methods [40, 42, 44]. As the forward approach is limited to verifications at the EPID level and not inside the patient geometry, the comparison between prediction and measurement cannot be translated to clinically relevant metrics, such as dose volume histograms of targets and OARs, which is its main disadvantage.

The backprojection approach, on the other hand, provides dose distributions inside a patient model, typically computed tomography (CT) images. The backprojection approach can be further divided into *direct backprojection* and *indirect backprojection.* In the *direct* approach, the acquired EPID signal is converted into a primary fluence map at the EPID level, after corrections for scatter within the EPID and from the patient. The dose distribution can be determined by backprojecting the 2D fluence map to a plane inside the patient model, using dose deposition kernels [45, 46]. Further corrections might be necessary when backprojecting to different depths inside the patient, with numerous approaches proposed by different research groups [28]. For instance, Wendling et al. [45] extend a previously validated 2D EPID-based method [47] to 3D by backprojecting the primary fluence through the entire patient volume. In addition to the corrections inherent to their 2D method, such as the beam attenuation from the patient and the scatter within the patient at the reconstruction plane, extra corrections are needed for the beam hardening and build-up effect at beam entrance. As both 2D and 3D methods convert the EPID signal into dose-to-water, the resultant 3D dose distribution is given in terms of dose-to-water and does not account for tissue inhomogeneities, for which some solutions have been presented [48, 49]. By comparing the backprojected 3D distribution to distributions calculated with the treatment planning system (TPS) for a water-filled patient (all voxels inside the patient CT used for calculation overwritten with water values), the methodology has been proven able to detect changes in patient positioning and anatomy, however it does not estimate the dose delivered to the patient [49, 50]. Besides, this requires an extra TPS dose distribution calculation for the water-filled patient, which represents an additional workload on the planning staff. Furthermore, the methodology relies on in-air EPID acquisitions for transmission corrections, meaning that for each patient irradiation, one extra in-air EPID irradiation is necessary, increasing not only the workload on the clinical staff, but also on the treatment time assigned for each patient. Finally, Olaciregui-Ruiz et al. [50] have recently incorporated dose inhomogeneity conversion maps, defined as the ratio between patient and water-filled patient dose distributions computed with MC, yielding dose distributions inside patient geometries. Ultimately, the in-air EPID measurements could be replaced by a calculated transmission model [46].

Alternatively, in the *indirect* approach, the incident fluence at the entrance of the patient is obtained with in-air EPID measurements (prior to patient irradiation), or by backprojecting the primary fluence through a patient model, obtained from transmitted

EPID signals (during patient irradiation). This is subsequently given as input to dose calculation engines to generate dose distributions inside the patient model, using either analytical or MC methods [51–55], whereby MC-based approaches yield better accuracy, especially considering the inhomogeneous patient geometries [28]. Steciw et al. [56] used the clinical TPS, also used in the planning process, with the advantage that clinical TPSs are constantly checked and commissioned to the equipment in use. The pitfall is that possible errors in the TPS dose calculation remain unchecked [28].

For both forward and backprojection methods, the raw EPID signals should be pre-processed to account for dead pixels, pixel sensitivity, pixel-value-to-dose conversion (if applicable) and energy response [32, 35, 45, 57]. A comprehensive review of EPID dosimetry can be found in Olaciregui-Ruiz et al. [38], van Elmpt et al. [28], McCurdy et al. [58], and references therein.

The ultimate goals of an in vivo dosimetry method is to determine the dose delivered to the patient and detect deviations between planned and delivered dose distributions. Mans et al. [59] report on several errors in patient irradiation identified by their clinically applied EPID in vivo dosimetry routine, which would not have been detected only by standard pre-treatment verification methods. Bojechko et al. [15] have evaluated numerous incident reports with high potential severity at their institute, concluding that a majority of errors could have been detected with in vivo EPID dosimetry. Finally, EPID-based in vivo dosimetry has the potential to detect dose deviations arising from different sources, such as (but not limited to) [28, 38]:

- *Machine-related errors*: multi-leaf collimator (MLC) malfunctioning[2], discrepancies in the collimator and gantry angles, beam flatness and symmetry;
- *Plan-related errors*: TPS beam model, dose calculation engine, delivery of wrong treatment plans, problems with plan data transfer;
- *Patient-related errors*: patient geometric errors (anatomical changes, movement during treatment and/or mispositioning), missing or wrong bolus material, wrong patient.

EPID-based in vivo dosimetry brings clear benefits to the patient treatment outcome, being able to identify inaccuracies during patient irradiation that would not have been detected by other methods. Moreover, a myriad of commercial systems is available, and a list of the current commercial solutions can be found in Olaciregui-Ruiz et al. [38]. However, despite its clear benefits and commercial availability, there are no specific guidelines and recommendations for EPID-based in vivo dosimetry methods, which does not encourage its expansion. Hence, very few radiotherapy centers have it implemented in

---

[2] Examples of MLC malfunctioning are wrong positioning of leaves and/or diaphragms, wrong speed during leaves and/or diaphragms movement, among others.

the clinical routine. Finally, an optimal methodology for EPID-based in vivo dosimetry should be accurate and fast, providing reliable estimates of the dose delivered to the patient within seconds, and should not increase the workload on the clinical staff.

In the past decade, the use of artificial intelligence in medical physics has boomed, and the potential of deep learning (DL) methods has been explored for a wide range of applications [60, 61], such as image segmentation [61], treatment planning [62, 63], outcome prediction [64, 65], detection and diagnosis [66–69] and dose estimation [70, 71]. EPID-based data and DL methods have been combined to propose solutions for machine quality assurance [72, 73], pre-treatment IMRT dosimetry [74, 75] and tumor motion prediction [76]. Wolfs et al. [77] demonstrate how powerful the combination of convolutional neural networks (CNNs) and forward EPID-based in vivo dosimetry can be, by presenting a proof-of-concept of a fast and automated method for plan delivery errors detection. Moreover, CNNs have been used to establish direct mapping from coarse to accurate dose distributions [78], and ultimately for EPID-based in vivo dosimetry for hybrid magnetic resonance imaging (MRI) linac machines [70, 79], in which the external magnetic field disturbs the EPID signal and consequently the resultant dose reconstruction.

1.2 Objective and outline of this thesis

This thesis describes a methodology that combines the accuracy of MC methods with the efficiency of CNN, taking into consideration the requirements for a clinically feasible EPID-based in vivo dosimetry, and inspired by the recent achievements of CNNs in medical physics, especially the Deep Dose Estimate (DDE) in estimating patient-specific dose distributions for diagnostic CT acquisitions [71, 80]. The DDE network is trained to predict dose distributions inside a patient CT for sub-fields of step-and-shoot IMRT treatments, giving as input the patient's CT image and a first-order dose (FOD) approximation. The FOD is reconstructed from a MC simulation of EPID signals, scored behind the patient CT. An accurate dose distribution (ADD), namely the MC simulated dose inside the patient CT, is given as the training target. Once trained, the network is able to correct for beam hardening, build-up and scatering effects, bypassing the complex corrections typically required for 3D EPID-based in vivo dosimetry [45, 46]. Moreover, the predicted distribution is given as dose-to-medium, properly accounting for inhomogeneities of the patient, so no conversion from dose-to-water is necessary [50]. Predictions in terms of dose-to-water are also feasible depending on the training data. Finally, since the EPID was modeled as a water box [32, 81], the EPID signal is scored as dose-to-water and no further corrections are needed to the scored signal prior to FOD production, in contrast to other methodologies [70].

The thesis is organized as follows: Chapter 2 presents the theoretical background of all topics relevant to the thesis, such as the physical interaction of radiation with matter, a brief overview of radiotherapy techniques focusing on Intensity Modulated Radiation

Therapy (IMRT), an overview of modern clinical linear accelerators (linac) and amorphous-silicon EPIDs, a brief description of Monte Carlo (MC) methods and finally an overview on neural networks.

The proposed methodology relies on MC simulations of dose distributions inside a virtual patient and at a virtual EPID model, and therefore a reliable MC model of the linac considered is essential. Chapter 3 describes the production of a detailed MC model of the Elekta Synergy® linac head coupled to an Elekta Agility$^{\text{TM}}$ MLC collimator, using a combination of vendor-provided geometric parameters and information freely available in the literature. The model has been extensively validated to represent the equipment installed at the Klinikum der Universität München (Klinikum Grosshadern).

Linac specific geometric information is not always available, as these are often a commercial secret and not easily disclosed by the vendors. Hence, Chapter 4 describes a methodology for optimization of freely available IAEA phase space (PhSp) files. The energy and angular direction of the particles stored in the PhSp files were manipulated until they can represent the Elekta Agility head. This offers an alternative for MC linac simulations when both the geometric details and validated IAEA PhSp files are unavailable for the linac of interest.

Chapter 5 is dedicated to the Deep Dose Estimate (DDE) neural network and all the necessary steps for network training and validation. The selection and preparation of the patient CT database used for both MC simulations and DDE training is described in detail, as well as all the aspects involved in the production of the first order dose approximations (FODs). The network implementation and training process are also described. Finally, the promising results obtained with the trained network are presented.

Chapter 6 presents a general discussion of all results obtained, followed by possible extensions of the methodology and future perspectives. Finally, a summary of the work developed throughout this thesis is given in Chapter 7, together with general conclusions.

CHAPTER 2

Theoretical Background

2.1 X-ray production

In a very simplified explanation, X-rays are a type of electromagnetic radiation generated when an electron beam, produced by a cathode, is accelerated towards a target called anode, thanks to a high voltage applied between cathode and anode. The accelerated electrons reach high velocities before hitting the target, where they are suddenly deflected or decelerated by the nuclei of the target material. A schematic representation of an X-ray tube used for X-rays production is depicted in Fig. 2.1. X-rays are produced by two different mechanisms, resulting in *Bremsstrahlung* and *characteristic* X-rays, described briefly in the following.

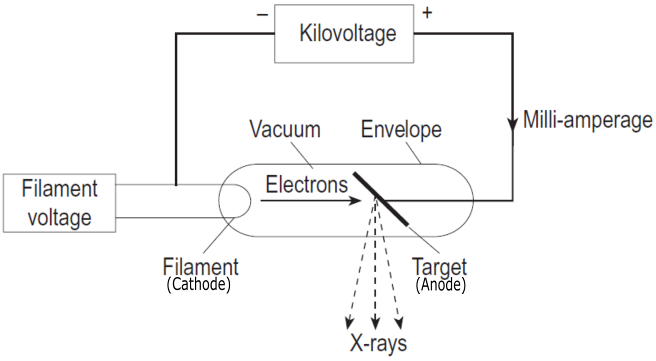

Figure 2.1: Schematic representation of an X-ray tube. Reprinted with permission from [82].

The bremsstrahlung radiation is generated when a high-speed electron, traveling in the vicinity of a nucleus, interacts with the Coulomb forces of the nucleus and changes its trajectory; the deflected electron therefore loses part of its energy in the form of electromagnetic radiation (i.e. photons) (Fig. 2.2). The incident electron can have more than one bremsstrahlung interactions inside the material, and each interaction might result in partial or total energy loss. Hence, the resultant photon can have any energy, limited only to the maximum value of the initial energy of the electron. Therefore, the bremsstrahlung X-rays form a *polyenergetic* beam. Moreover, the direction of emission of the bremsstrahlung photons depends on the energy of the incident electrons: at energies below 100 keV, the photons are emitted more or less homogeneously in all directions; the direction of the emitted photons becomes increasingly forward with increasing electron energies (Fig. 2.3). Hence, for clinical linear accelerators (linac), which operate in the megavoltage range, *transmission type* targets are used, in which the electron beam strikes the target from one side and the produced X-rays travel to the other side.

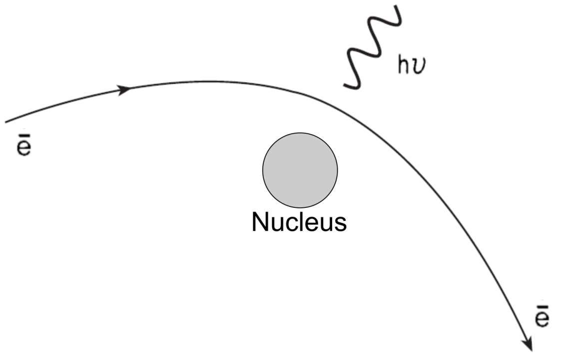

Figure 2.2: Illustration of the bremsstrahlung process, where the incident electron is deflected by the nucleus, emitting a photon of energy $h\nu$. h is the Planck constant and ν is the frequency of the emitted electromagnetic wave (photon).

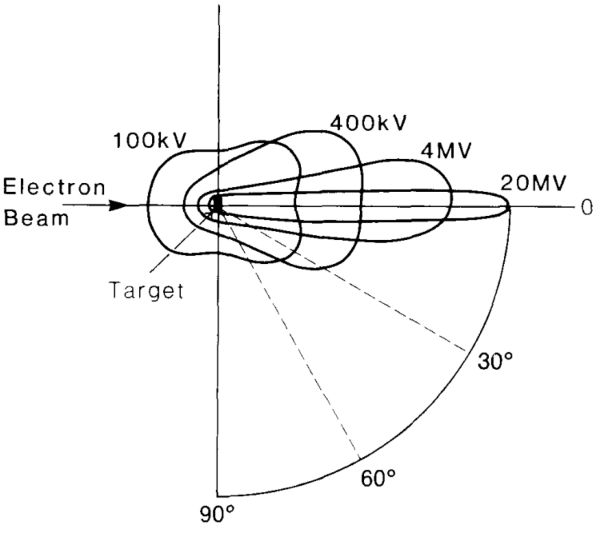

Figure 2.3: Illustration of the angular distribution of bremsstrahlung radiation for different energies of the incident electron beam. Reprinted with permission from [83].

Alternatively, the incident electron can eject one orbital electron from the atoms of the material. One electron from a higher orbital (or shell) fills the resultant vacancy, and the energy difference between the two shells is emitted in the form of a *characteristic* photon. In this interaction, the resultant photons have discrete energies that are characteristic to the particular atom in which this shell transition has occurred, hence the name of *characteristic* X-rays. The original electron is deflected and continues with a lower energy (original energy minus the binding energy from the orbital electron and the energy given to it). Fig. 2.4 illustrates the production of characteristic X-rays.

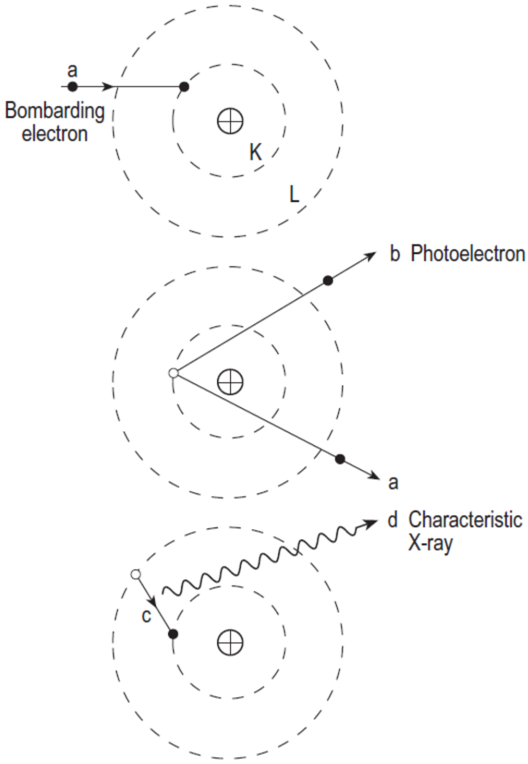

Figure 2.4: Schematic representation of the production of characteristic X-rays. Reprinted with permission from [82].

2.2 Interaction of radiation with matter

In this thesis, radiotherapy performed with photon beams of 6 MV nominal energy was considered. Therefore, the relevant physical aspects regarding this radiation type and energy will be covered. A more comprehensive description of all aspects of radiation interaction with matter can be found elsewhere [83, 84].

2.2.1 Photon beam attenuation

Photon beams, also denoted as X-rays, are classified as *indirectly* ionizing radiation. This type of radiation can only deposit energy to the traversed medium in more than one step, via secondary products (hence *indirectly*): the incident photons eject electrons from the surrounding atoms, which will then transfer their energy by producing ionization and excitation of other atoms along their path. In this process, the photons can be completely absorbed by the medium, and hence the number N of photons that travel further deep in the medium reduces with the thickness of the medium. In other words, the intensity I of the photon beam decreases, with respect to the initial incident intensity I_0, when the thickness x of transversed medium (or absorber) increases. This process is called *attenuation* and can be described by Eq. 2.1:

$$I(x) = I_0 e^{-\mu x} \tag{2.1}$$

where μ is the linear attenuation coefficient, given by Eq. 2.2:

$$\mu = \rho \frac{N_A}{A} \sigma(h\nu, Z) \tag{2.2}$$

with ρ being the density of the absorber material, N_A is the Avogadro's number (amount of atoms in one mol of element or compound), A is the number of grams per mol of material (also called molar mass) and σ is the cross section. The cross section $\sigma(h\nu, Z)$ depends on the photon energy $h\nu$ and the atomic number Z of the absorber, and represents the probability of interaction of the photon with the atoms of the absorber.

The attenuation of a photon beam by an absorber is caused by five major types of interactions, namely coherent (Rayleigh) scattering, photoelectric effect, Compton effect, pair production and photonuclear disintegration. Each interaction has its own cross section, i.e. own probability of interaction. Hence, the probability for each interaction depends on the energy $h\nu$ of the photon and on the atomic number Z of the material[1], as illustrated in Fig. 2.5.

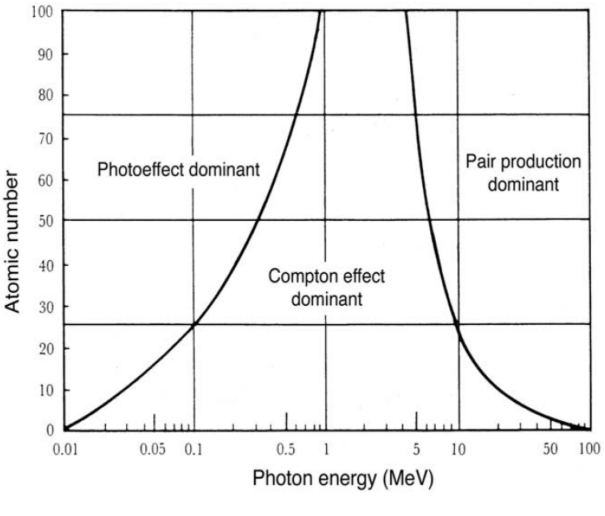

Figure 2.5: Relative predominance of the three major types of photon interactions, with respect to photon energy ($h\nu$) and atomic number (Z) of the medium. The three different regions denominate the values where each effect is predominant. The curves represent the regions where the neighboring predominant effects have the same probability, i.e. the same cross section. Reprinted with permission from [85].

---

[1] For compounds and mixtures, the effective atomic number Z_{eff} has to be considered

Coherent scattering is mostly relevant at very low photon energies ($<$ 10 keV) and for materials with high atomic number Z, being often disregarded for therapeutic beams. In contrast, photonuclear disintegration occurs mainly at very high photon energies ($>$ 10 MeV), being out of the energy range used in this work [83, 85]. Therefore, the descriptions of both processes are omitted in this thesis. The other three interactions are fundamental for radiation therapy with 6 MeV, and are described in the following.

Each one of these processes can be represented by its own linear attenuation coefficient, such that the total attenuation coefficient μ can be defined as the sum of the individual coefficients:

$$\mu = \tau + \sigma_c + \kappa \tag{2.3}$$

where τ, σ_c and κ are the linear attenuation coefficients for the photoelectric effect, Compton effect and pair production, respectively. For convenience, μ can be divided by the density of the medium ρ, resulting in the mass attenuation coefficient μ/ρ, leading to Eq. 2.4:

$$\frac{\mu}{\rho} = \frac{\tau}{\rho} + \frac{\sigma_c}{\rho} + \frac{\kappa}{\rho} \tag{2.4}$$

where τ/ρ, σ_c/ρ and κ/ρ are the mass attenuation coefficients for the photoelectric effect, Compton effect and pair production, respectively.

It is important to mention that the exponential attenuation described by Eq. 2.1 only applies to monoenergetic beams. Clinical linear accelerators however produce polyenergetic photon beams, with an energy spectrum characterized by the target material and limited by the maximum energy of the electron beam incident on the target (section 2.1). In a polyenergetic beam, photons with lower energies are attenuated more rapidly than those with higher energies, increasing the mean energy of the beam as it goes through the medium, an effect known as *beam hardening*, and Eq. 2.1 needs to be modified to accommodate a polyenergetic beam[2]. However, the underlying mechanisms of photon interaction with matter are the same.

2.2.2 Photoelectric effect

In the photoelectric effect, a photon interacts with one atom and ejects one of the tightly bound orbital electrons, typically in the K, L, M or N orbital shells from the atom [83] (Fig. 2.6).

The photon is completely absorbed, and the ejected electron, called *photoelectron*,

---

[2] For a polyenergetic beam, $I(x) = \int_0^{E_{max}} I_0(E) e^{-\int \mu dx} dE$, where E is the energy.

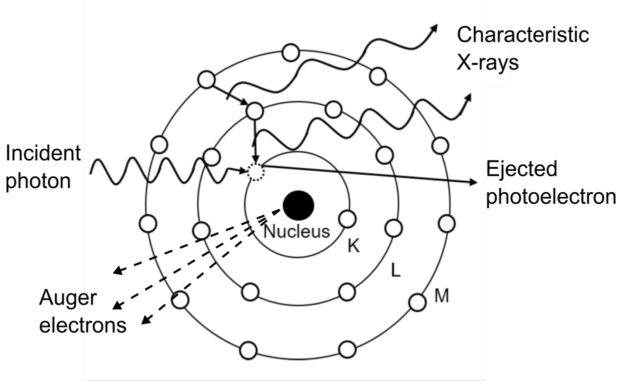

Figure 2.6: Schematic representation of the photoelectric interaction. Modified and reprinted with permission from S. Liprandi [6].

has a kinetic energy E_e given by Eq. 2.5

$$E_e = h\nu - E_B \tag{2.5}$$

where E_B is the binding energy of the electron. The ejected electron leaves a vacancy in one atomic shell, which is thus in an excited state. This vacancy is filled by an electron from an outer orbital, and the energy difference is released as characteristic X-rays and/or Auger electrons.

The photoelectric effect is dominant for low energy photons, as illustrated in Fig. 2.5. Fig. 2.7 shows the mass attenuation coefficient for the photoelectric effect τ/ρ for water and lead, and it can be observed that the probability increases with increasing values of the atomic number Z of the absorber. There is no analytical expression for τ/ρ that covers the entire range of photon energy and Z. An approximation for the photoelectric cross section $\sigma_{p.e.}$ can be expressed as Eq. 2.6, with $n \approx 4$ and $m \approx 3$ for $h\nu < 0.1$ MeV, where it dominates [84, 86].

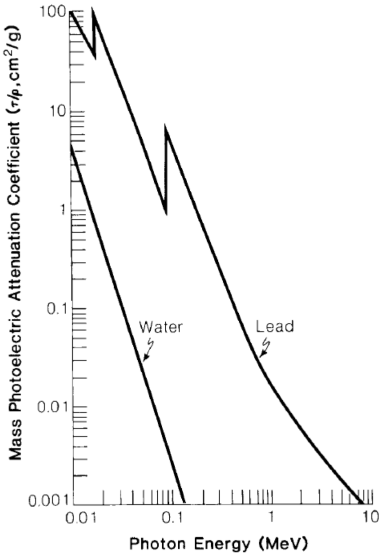

Figure 2.7: Mass attenuation coefficients for the photoelectric interaction τ/ρ for water ($Z_{eff} = 7.42$) and lead ($Z = 82$). The absorption edges for the K and L shells of lead (88 and 15 keV respectively) are also depicted. Reprinted with permission from [83].

$$\sigma_{p.e.} \propto \frac{Z^n}{(h\nu)^m} \tag{2.6}$$

Consequently, τ/ρ dependency to Z and $h\nu$ follows Eq. 2.7, with $m \approx 3$.

$$\frac{\tau}{\rho} \propto \left(\frac{Z}{h\nu}\right)^m \tag{2.7}$$

The curve of τ/ρ for lead has some edges at around 15 and 88 keV (Fig. 2.7), which correspond to the binding energies of the L and K shells of lead, respectively. If the energy of the photon is lower than 15 keV, it does not have enough energy to eject an electron from the L shell, limiting the interactions to electrons at higher shells. If the photon has exactly the binding energy of an electron at the L shell, resonance occurs and there is a high probability of photoelectric effect in the L shell. Beyond this point, the probability decreases with increasing photon energy, until 88 keV, which is the binding energy of the K shell and the probability spikes again. These probabilities spikes are called *absorption edges* and are material-dependent. The absorption edges for water are at very low photon energies (K shell ∼ 0.5 keV) and are therefore not shown in Fig. 2.7.

2.2.3 Compton effect

As observed in Fig. 2.5 the Compton effect is the dominant effect in the energy range of radiotherapy. The Compton effect can be described as an inelastic interaction between the incident photon and one electron from the outer shells of the absorber atom, sometimes referred to as *free* electron. This means that the binding energy of the electron is much lower than the energy of the incident photon. The energy of the photon is partially transferred to the electron, which is ejected at an angle ε with respect to the original photon direction (*recoil electron)*; the photon, with the remaining energy, is scattered at an angle θ. Fig. 2.8 illustrates the Compton effect.

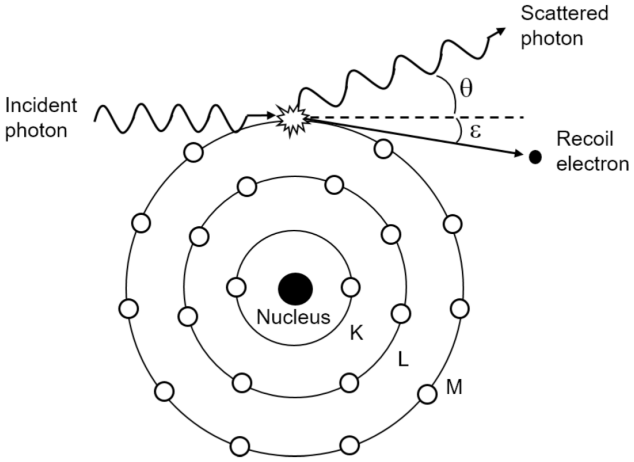

Figure 2.8: Schematic representation of the Compton interaction. Reprinted with permission from S. Liprandi [6].

Considering the laws of conservation of energy and momentum, the energy of the scattered photon $h\nu'$ can be derived, described in terms of its scattering angle θ by Eq. 2.8:

$$h\nu' = \frac{h\nu}{1 + \frac{h\nu}{m_e c^2}(1 - cos\theta)} \tag{2.8}$$

where $m_e c^2$ is the rest mass energy of the electron (0.511 MeV). Using Eq. 2.8, the kinetic energy of the recoil electron E_e is then given by Eq. 2.9. The binding energy of the electron E_B can be neglected, as it is much smaller in comparison to the photon and recoil electron energies.

$$E_e = h\nu - h\nu' = h\nu \frac{\frac{h\nu}{m_e c^2}(1 - cos\theta)}{1 + \frac{h\nu}{m_e c^2}(1 - cos\theta)} \tag{2.9}$$

From Eq. 2.9 the minimum energy E_e transferred to the recoil electron can be interpreted as the interaction when $h\nu \simeq h\nu'$. This implies that the scattered photon has a scattering angle $\theta \simeq 0°$, being only slightly deflected by the interaction. However, a direct strike between the photon and the electron will cause the ejected electron to travel forward ($\varepsilon = 0°$), while the scattered photon travels backward ($\theta = 180°$). This represents the scenario when the scattered photon is left with a minimum energy $h\nu'_{min}$, while the recoil electron gets maximum energy $E_{e,max}$. Replacing $\theta = 180°$ in Eqs. 2.8 and 2.9,

$h\nu'_{min}$ and $E_{e,max}$ are thus given by Eqs. 2.11 and 2.10, respectively.

$$h\nu'_{min} = h\nu \frac{1}{1 + \dfrac{2h\nu}{m_e c^2}} \tag{2.10}$$

$$E_{e,max} = h\nu \frac{\dfrac{2h\nu}{m_e c^2}}{1 + \dfrac{2h\nu}{m_e c^2}} \tag{2.11}$$

For the Compton effect, the energy of the incident photon needs to be large in comparison to the binding energy of the electron. On the contrary, the probability of the photoelectric effect reaches a maximum when the energy of the incident photon is similar to the binding energy of the electron. Thus, as the energy of incident photons becomes larger than the binding energy of the electrons in the K shell, the probability of the photoelectric effect decreases rapidly, following Eq. 2.7, while the probability of the Compton effect increases, as illustrated by Fig.2.5. However, the probability for the Compton effect decreases gradually for higher photon energies, approaching a $(h\nu)^{-1}$ dependency [86] (Fig. 2.9).

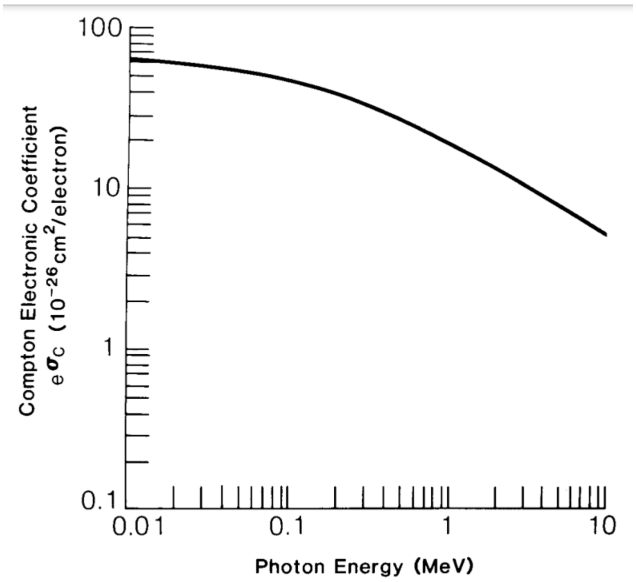

Figure 2.9: Variation of the Compton electronic cross section $_e\sigma_c$ with respect to incident photon energy. The mass attenuation coefficient σ_c/ρ can be obtained by multiplying $_e\sigma_c$ to the number of electrons per gram for a given material. Reprinted with permission from [83].

Since the electron binding energy can be considered zero for the Compton effect, the electronic cross section $_e\sigma_c$[3] is independent of the atomic number Z of the material. Hence, the cross section per atom of any Z (atomic cross section, $_a\sigma$, equivalent to the

---

[3] The electronic cross section $_e\sigma$ is the cross section per electron, i.e. the probability of interaction per electron.

cross section) can be determined from Eq. 2.12 as:

$$ {}_a\sigma_c = Z_e\sigma_c \tag{2.12}$$

Following Eq. 2.2 on the relationship between the linear attenuation coefficient and cross section, the mass attenuation coefficient for the Compton effect σ_c/ρ can be written as Eq. 2.13:

$$\frac{\sigma_c}{\rho} = \frac{N_A}{A} {}_a\sigma_c = \frac{N_A Z}{A} {}_e\sigma_c \tag{2.13}$$

For a neutral atom, the number of electrons in the atom equals the atomic number Z. Hence, the quantity $(N_A Z)/A$ gives the number of electrons per gram of material [86]. With exception of Hydrogen, which has no neutrons in its nucleus and therefore has an electron density approximately twice higher than other elements, most elements can be considered as having approximately the same number of electrons per gram. Table 2.1 shows the electron density (number of electrons per gram) for different materials and compounds. The value of Z/A varies between 0.5 and 0.4, with a decreasing tendency with increasing Z (Hydrogen is the exception, with $Z/A = 1$). By assuming Z/A to be roughly constant, the Compton mass attenuation coefficient can be considered independent of Z, and thus, σ_c/ρ is nearly the same for all materials.

Table 2.1: Electron density (number of electrons per gram) for different materials and compounds [83].

Material or Compound	Density (g/cm^3)	Z or Z_{eff}	Electron density ($\times$ 10^{23}/g)
Hydrogen	0.0000899	1	6.00
Carbon	2.25	6	3.01
Oxygen	0.001429	8	3.01
Aluminum	2.7	13	2.90
Copper	8.9	29	2.75
Lead	11.3	82	2.38
Fat	0.916	5.92	3.48
Muscle	1.0	7.42	3.36
Water	1.0	7.42	3.34
Air	0.001293	7.64	3.01
Bone	1.85	13.8	3.00

2.2.4 Pair production

When the energy of the incident photon is higher than 1.02 MeV, the photon might interact with the electromagnetic field of the nuclei of the absorber material. The photon completely transfers its energy to create an electron-positron pair. Since the rest mass energy of electrons and positrons $m_e c^2 = 0.511$ keV, this process has an energy threshold and can only occur if the photon has an energy higher than $2m_e c^2 = 1.02$ MeV. For photons with energy higher than this threshold, the exceeding energy is shared, on average equally, between the electron and positron as kinetic energy, E_e^- and E_e^+, respectively (Eq. 2.14).

$$h\nu - 2m_e c^2 = E_e^- + E_e^+ \tag{2.14}$$

The positron produced loses its energy as it travels through the material by the same mechanisms of an electron. The positron combines to an atomic electron in its vicinity ("positronium") and finally annihilates, most likely, in two photons of 0.511 keV each. The annihilation photons travel in opposite directions due to momentum conservation. Fig. 2.10 illustrates the pair production and subsequent annihilation process.

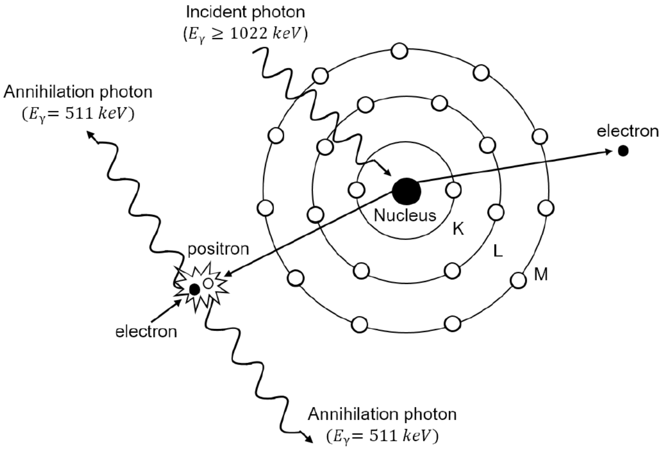

Figure 2.10: Schematic representation of the pair production process. Reprinted with permission from S. Liprandi [6].

As pair production occurs in the electromagnetic fields of the nuclei[4], the probability of this process rapidly increases with the atomic number Z, as described in Eq. 2.15 for the pair production cross section ($\sigma_{p.p.}$). For the same material, the probability increases approximately with the logarithm of the incident photon energy $h\nu$.

$$\sigma_{p.p.} \propto Z^2 ln(h\nu) \tag{2.15}$$

of the absorber resulting in the photon to disappear and three light charged particles are released: the original orbital electron and the electron-positron pair produced in the interaction.

---

[4] A similar phenomenon can occur, with lower probability, in the field of an atomic electron, called *triplet production.* Triplet production can occur for incident photons with energy higher than $4m_e c^2 =$ 2.04 MeV, resulting in three particles: the atomic electron (released after the photon interaction) and the electron-positron pair.

2.3 Medical linear accelerators

Medical linear accelerators (linac) are the most widely used radiation source in radiotherapy. In a linac, electrons are accelerated inside a linear tube in vacuum, called waveguide, using high-frequency electromagnetic waves. The high-energy electron beam itself can be used for patient treatment, or it can be used to generate X-rays (section 2.1). Fig. 2.11 shows a schematic diagram of a typical modern medical linac, however some details might differ among machines from different manufacturers.

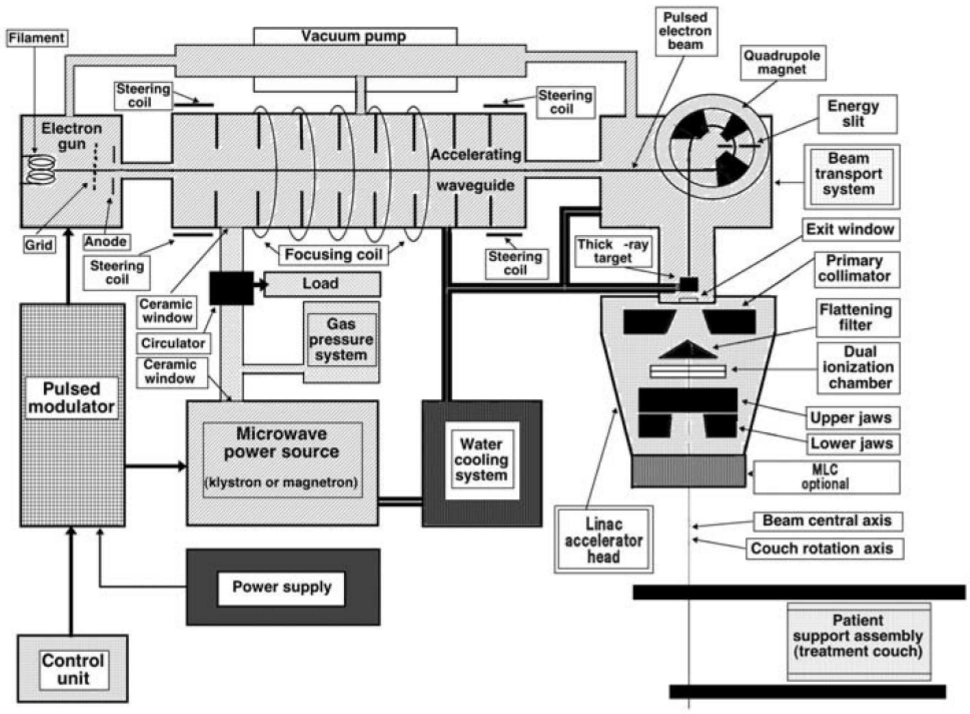

Figure 2.11: Schematic representation of a typical medical linear accelerator. Some components can vary between equipments from different manufacturers. Reprinted with permission from [87].

The components of a linac can be divided into five groups, which will be briefly described, as follows:

1. Injection system
2. Radio-frequency (RF) power generation system
3. Acceleration waveguide
4. Beam transport system
5. Beam collimation and beam monitoring systems

The injection system is the source of electrons, containing a heated cathode, which will emit electrons by thermionic emission, and a perforated anode. The emitted electrons are focused into a pencil beam and, thanks to a potential difference applied between the cathode and the anode, are accelerated towards the perforated anode, drifting through the hole into the acceleration waveguide. The high power microwave radiation, responsible for accelerating the electrons in the waveguide, is produced by the RF power generating system.

As the electrons enter the acceleration waveguide, they interact with the electromaginetic field of the microwaves set up in the waveguide, gaining energy and acceleration. Generically explained, the waveguide is a copper tube with its interior divided by copper disks with circular apertures, positioned at equal spacing inside the tube; the waveguide is thus "divided" into a series of cylindrical cavities that form the basic structure of the acceleration waveguide of a linac. The electron beam transport system brings the pulsed, high-energy electron beams from the waveguide onto the target (for X-ray irradiations) or onto a scattering foil (for electron irradiations).

The beam monitoring and beam collimator systems ensure that the radiation is delivered to the patient as prescribed. Finally, the medical linac contains several other components that are not directly involved in the beam generation, however they make it possible: the vacuum pumping system of the waveguide, the water cooling system, the air pressure system and the shielding against leakage of radiation.

It is important to notice the possible configurations of medical linacs, regarding the beam transport. For electrons with megavoltage energies, the bremsstrahlung X-rays produced by the target are mainly forward-peaked, i.e. the clinical photon beam is produced in the direction of the incident electron beam on the target (section 2.1). The simplest configuration possible consists of the electron gun and target directly aligned with the isocenter, as illustrated in Fig. 2.12a (direct isocenter mounting). This configuration is only possible for the production of nominal photon beam energies of 4 to 6 MV. For intermediate (8 to 15 MeV) and high (15 to 30 MeV) electron energies, the accelerating waveguides have to be longer, becoming too long for direct isocenter mounting. The waveguides are then located in the gantry parallel to the axis of rotation as in Fig. 2.12b, or in the gantry stand as in Fig. 2.12c. A beam transport system is used to bring the electron beam from the waveguide onto the target.

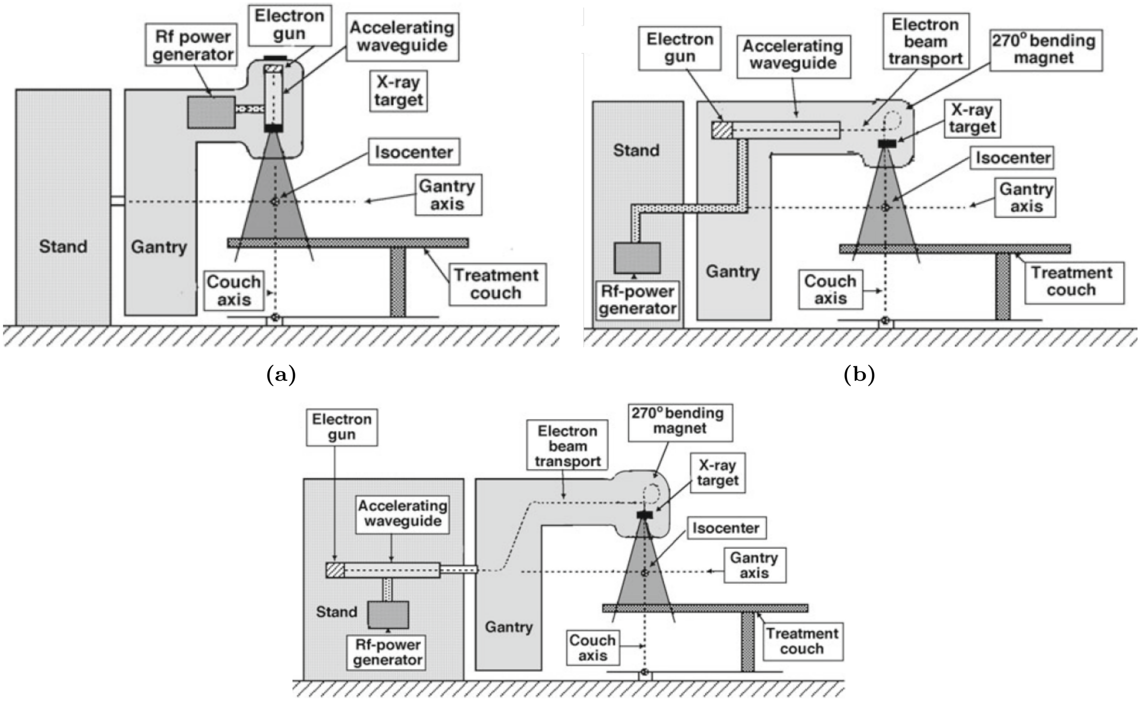

(a) **(b)**

(c)

Figure 2.12: Different design configurations of medical linacs. In (a), the gun and target are directly aligned with the isocenter, and fixed to the acceleration waveguide. In (b), the acceleration waveguide is mounted in the gantry while in (c) the waveguide is in the gantry stand. For both (b) and (c) configurations, the electron beam is steered to the target via bending magnets from the beam transport system. Reprinted with permission from [87]

.

The beam transport system is composed of evacuated drift tubes and bending magnets. Additionally, energy slits are used to remove electrons with energies that differ from the nominal electron beam energy (within a given threshold value). Steering and focusing coils are used for steering and focusing the accelerated electron beam from the waveguide into the target. Three different systems for electron bending have been developed, namely 90° bending, 270° bending (achromatic) and 112.5° bending (slalom), illustrated in Fig. 2.13.

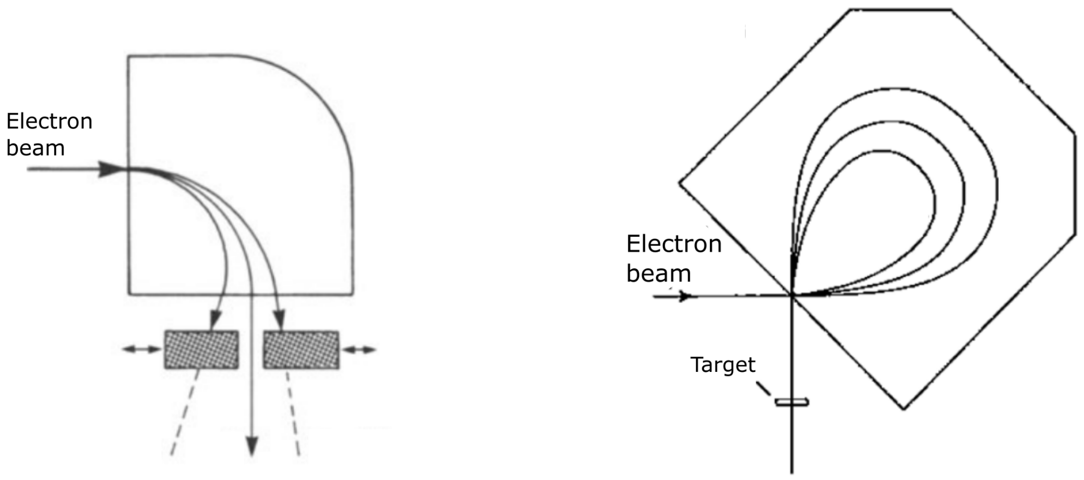

(a) 90° bending **(b)** 270° bending

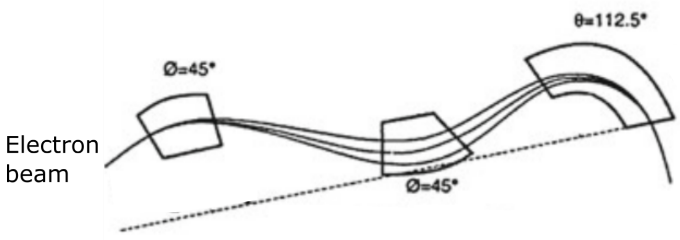

(c) 112.5° bending

Figure 2.13: Illustration of the three different bending magnet systems: (a) 90° bending, (b) 270° bending (achromatic) and (c) 112.5° bending (slalom). Reprinted with permission from [88–90]

.

2.3.1 The linac head

The linac treatment head consists of several components that influence the production, shaping and monitoring of the clinical photon and electron beams. Typical components of a linac head are the target, the primary collimator, the flattening filter and electron scattering foils (for electron treatment mode), ionization chambers, a light localizing system and a secondary collimator system. Fig. 2.14 displays the typical components of a linac head. Moreover, the linac head is encompassed by shielding materials such as lead, tungsten or lead-tungsten alloy, providing shielding against radiation leakage [83].

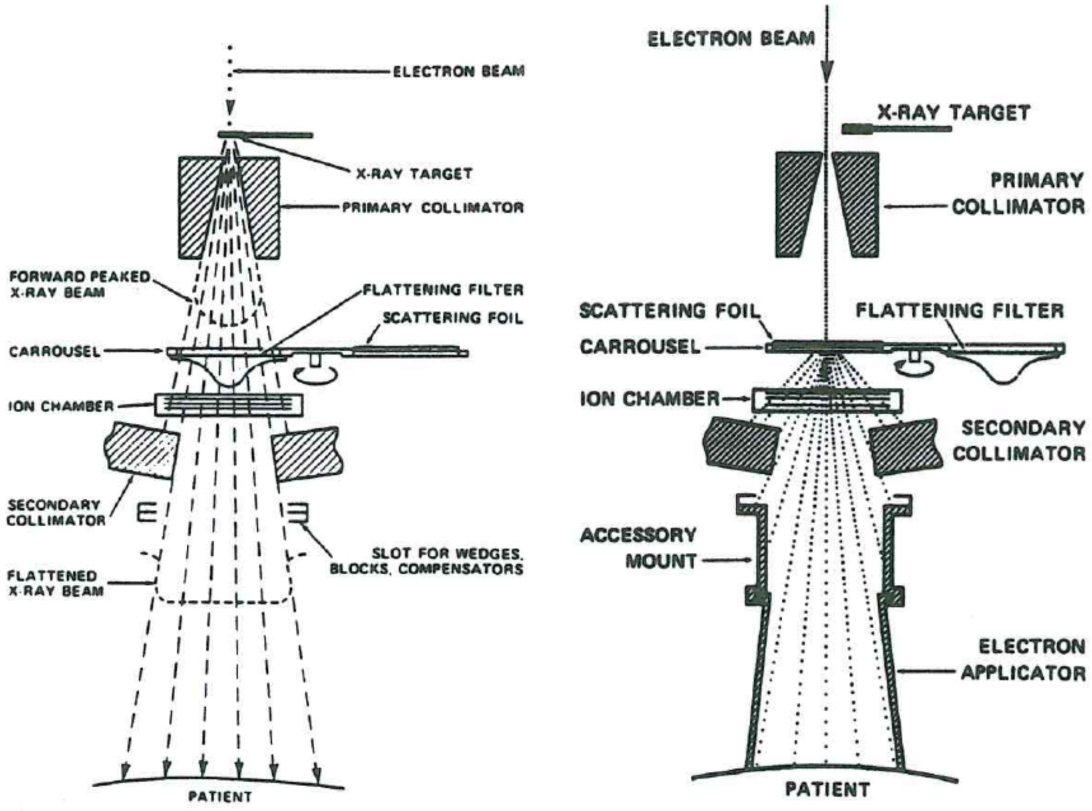

(a) **(b)**

Figure 2.14: Schematic representation of the typical components of a linac head. (a) Photon therapy mode, where the flattening filter is used. (b) Electron therapy mode, where the flattening filter is retracted, and a scattering foil is used in its place instead. Reprinted with permission from [91]

.

The clinical X-ray beams produced in the target are mainly forward peaked, as observed in Fig. 2.14a. To make the beam fluence uniform throughout the field, a flattening filter is used[5]. For each electron beam energy used to produce photons, i.e. for each nominal photon beam energy, a specific target/flattening filter combination is used [87]. Fig 2.14b shows the components used for electron therapy, where the flattening filter is replaced by a scattering foil.

The primary collimator is located right after the target, and it defines a maximum circular field size. Different from the secondary collimator system (subsection 2.3.2), the primary collimator is static and does not shape the field delivered to the patient, being therefore independent of patient irradiation. The ion chambers monitor the dose rate, integrated dose and field symmetry [83], while the light localizing system projects a light beam as if emerging from the X-ray focal spot, so the light field is representative of the radiation field.

---

[5] In some treatment cases the irradiation is performed without the flattening filter, called flattening filter free (FFF) irradiation.

2.3.2 The Multi-leaf collimator system

After the photon beam passes through the primary collimator, the resultant irradiation field covers a wide circular area at the isocenter level, located at a fixed 100 cm distance from the X-ray target. A secondary set of rectangular collimators (called *jaws*), mostly two pairs of lead or tungsten blocks positioned perpendicularly to each other and downstream from the beam traveling direction, blocks and shapes the radiation field to rectangular or squared fields from a wide range of size. In this configuration only squared and rectangular fields can be delivered to the patient, which might not result in optimal target coverage and/or sparing of healthy tissues and OARs. Custom-made shielding blocks of a variety of shapes and sizes can be attached to the linac head, providing extra collimation. These extra blocks are however patient- and field-specific, meaning that each irradiated field requires the production of a unique block. Several disadvantages on the use of such beam blocks have led to the development of a new collimation system: the Multi-leaf collimator (MLC) [87, 92].

The MLC system is an intrinsic part of modern clinical linacs. It consists of several narrow collimator blocks, called *leaves*, that can move independently from each other and block just a fraction of the radiation beam. Each leaf is automatically and separately controlled by its own motor, which enables the production of arbitrary, irregularly-shaped radiation fields. For instance, the leaves can be set such that the resultant radiation field is conformal to the tumor shape from a certain direction, sparing the surrounding OARs. Generally, narrower leaves can generate more conformal fields, but with a trade-off for complexity. Moreover, the leaves' movement can act as a dynamic filter to modulate the intensity of the beam [92]. Hence, the introduction of MLC systems played a decisive role in the development of modern, highly conformal and modulated radiation therapy techniques (section 2.5).

The MLC system can either replace the upper jaws, the lower jaws, or be added as a tertiary collimator. In this thesis, an Elekta Synergy® linac head coupled to an Elekta Agility^{TM} MLC collimator was considered. In the Agility collimator system the MLC replaces the upper jaws, hence this configuration will be considered hereon. The detailed description, advantages and pitfalls of all different configurations can be found elsewhere [92]. The Agility MLC is composed of 160 leaves, organized in two leaf banks with 80 pairs of interdigitating leaves[6], with 5-mm projected leaf width at the isocenter. The maximum field size at isocenter is 40×40 cm^2, with 5-mm leaf resolution across this area [93]. More details on the Agility MLC can be found in Chapter 3.

The leaves of the MLC must provide acceptable attenuation and for that, the material and design of the leaf shape are decisive, as both the material and the leaf height (leaf

---

[6] Interdigitating leaves refers to the MLC design in which leaves from one leaf bank can travel across the neighboring leaves of the opposing leaf banks without collision.

dimension along the direction of the photon beam propagation) influence the beam attenuation. More specifically, the combination of leaf height and density will determine the *intraleaf transmission*, i.e. the amount of radiation that transverses the leaf without being attenuated. Finding the optimal leaf shape and material is an engineering challenge, and different manufacturers have developed different approaches to this problem. Hence, leaf heights vary over different MLCs, however tungsten alloys seem to be the optimal material, with small variations in the alloy composition and density among different manufacturers [92].

The leaves in a MLC are positioned side-by-side following the beam divergence, with a fixed space (or gap g) between adjacent leaves, just enough to avoid friction. However, incoming radiation that travels through these gaps will not be attenuated: this is called *interleaf transmission*. To minimize the interleaf transmission, the leaf design includes an overlap between the adjacent leaves, called tongue-and-groove (T&G): in one side of the leaf there is a small step (called *tongue*), that fits perfectly inside a small retraction (or *groove*) present in the side of the adjacent leaf, as illustrated in Fig. 2.15. Different manufacturers have developed different designs, but the tongue-and-groove design is the most common nowadays [92].

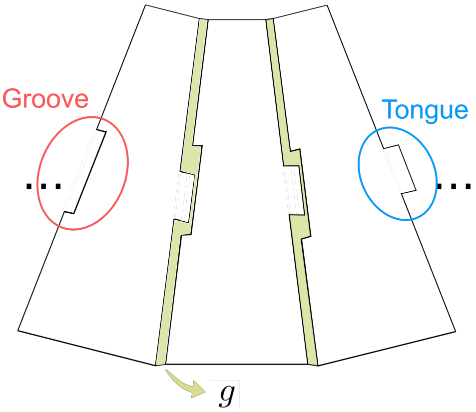

Figure 2.15: Schematic representation of the tongue-and-groove (T&G) design of adjacent leaves, with the tongue and the groove indicated in blue and red, respectively. The gap g between adjacent leaves is depicted in green. The figure is not to scale.

One problem arising from the T&G design is an underdosage of up to 20%, observed when one field is divided into two sub-fields (or two segments). This happens because, at the borders of the two fields, the radiation goes through either the tongue or the groove. The attenuation is therefore slightly different, causing an underdosage at the match-line of the two fields, as illustrated in Fig.2.16. This effect, called T&G effect is well known and is incorporated in clinical treatment planning systems (TPSs), so that this underdosage is taken into consideration for patient treatment.

Finally, the leaves of the Agility MLC have rounded ends, such that the radiation field is always tangent to the leaf end throughout the leaf motion. However, the penumbra of the field defined by a rounded leaf end is not independent of the leaf position [92]. Such

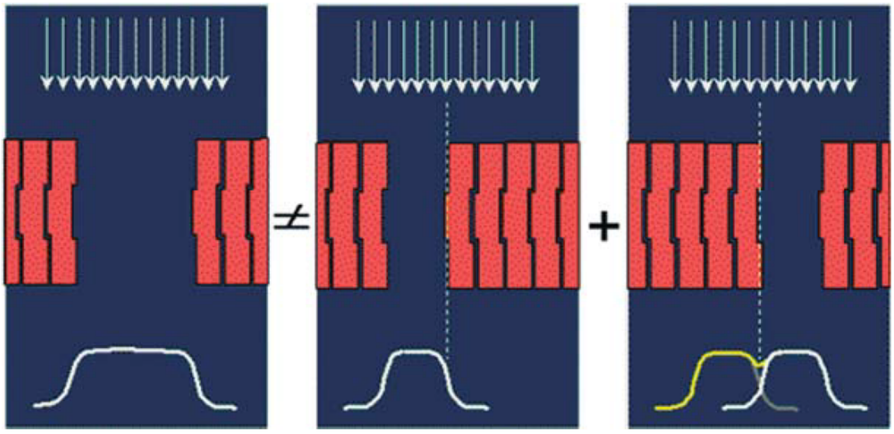

Figure 2.16: Illustration of the T&G effect, an underdosage caused by the T&G leaf design. Reprinted with permission from [94].

variations should also be implemented in the clinical TPSs.

2.4 Electronic Portal Imaging Devices (EPID)

Amorphous silicon (a-Si) Electronic Portal Imaging Devices (EPIDs) are an intrinsic part of modern medical linacs, with different designs available from different vendors. The Elekta Synergy® linac installed at Klinikum Grosshadern uses the PelkinElmer XRD 1640 AL5 P a-Si EPID (Elekta iView GT$^{\mathrm{TM}}$), which will be described in the following.

The PelkinElmer a-Si EPID is a flat panel sensor, with a scintillation layer connected to a pixelated photodiode array based on hydrogen-doped amorphous silicon (a-Si:H) semiconductors. A diagram of one a-Si EPID is depicted in Fig. 2.17. Each pixel in the photodiode array is connected to a switching thin film transistor (TFT). The scintillation layer is typically made of terbium doped gadolinium (Gd_2O_2S:Tb) or cesium iodide (CsI), with CsI being the material used in the PelkinElmer a-Si EPID. A 1-mm thick copper build-up plate is positioned on top of the scintillator. The entire structure is enclosed in a low-density plastic cover and connected to a read-out system. When the incident photon beam hits the EPID, secondary electrons might be produced either in the build-up plate or in the scintillator. These electrons interact with the scintillator, producing a shower of visible light (optical photons) which are absorbed by the photodiode array, generating an electric charge (indirect detection). Some photons might reach the photodiodes directly without interacting with the build-up plate and scintillator (direct detection), however the indirect detection gain is one order of magnitude higher than the direct detection. The charge is stored in the capacitor of the photodiode and controlled by the TFT switch, which transfers the signals of all pixels in parallel to the signal read-out system.

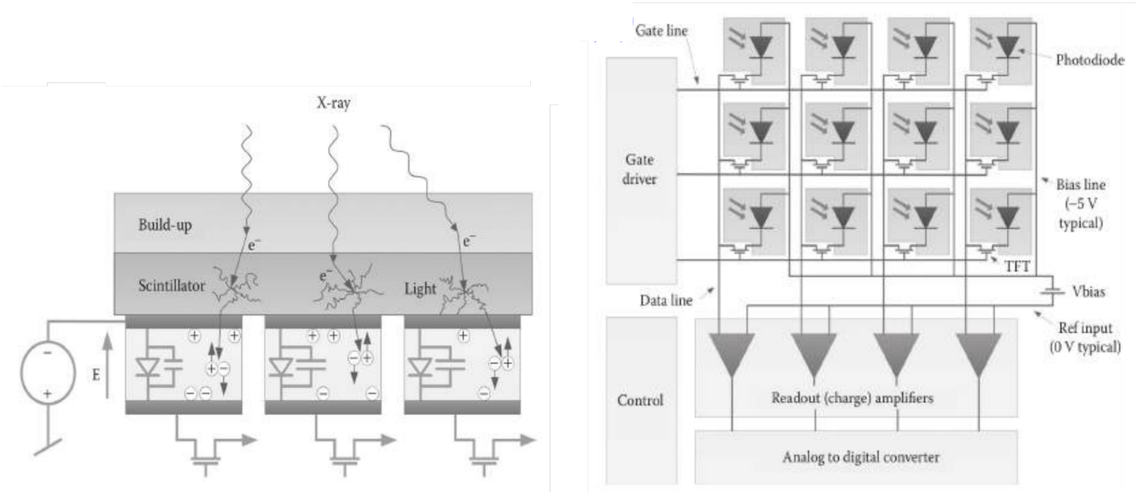

(a) **(b)**

Figure 2.17: (a) Schematic representation of the structure of an a-Si EPID, and the indirect conversion of ionizing radiation to optical photons in the scintillator. (b) An exemplary array of an a-Si EPID and associated electronics. Reprinted with permission from [95]

.

The EPID used in this work is located at 160 cm from the X-ray target. With an active detection area of 41 × 41 cm^2, field sizes of up to 25 × 25 cm^2 (at isocenter) can be detected. Finally, it has 1024 × 1024 pixels, with 0.4 mm pixel size. The dosimetric properties of the PelkinElmer a-Si EPID are described below.

2.4.1 Pixel response linearity

In theory, the signal generated from each pixel should be proportional to the intensity of the incident radiation of a certain radiation quality. In practice however, this is approximately true, but some deviations might occur due to photodiode saturation, pixel defect, incomplete signal acquisition, ghosting effect[7] or dose rate dependency. Such variations depend on the EPID model and acquisition software used. Corrections or calibrations can be applied to accommodate the deviations, as proposed by several authors [32, 96–98].

2.4.2 Energy-dependent response

In comparison to water equivalent detectors, a-Si EPIDs have a higher response to low-energy photons (< 0.5 MeV). This happens due to the increased number of photoelectric interactions in the copper plate and at the scintillator [32, 99, 100]. Scatter radiation coming from the linac head increases the low-energy component of the beam. This is specially relevant for regions of low dose areas, such as the penumbra regions on the edge

---

[7] Ghosting effect occurs when the signal acquired in one frame is not completely read out by the time subsequent frames are acquired. The subsequent frames will then be affected by the signal acquired on previous frames.

of the irradiated field or at the out-of-field regions. Several approaches have been proposed to model and correct for this dependency [32, 35, 99, 100].

2.4.3 Pixel sensitivity

The pixels in the a-Si array have individual offset signal and gain response to radiation. The signal variation is due to intrinsic differences in the pixels, combined to differences in the read out system. This can be corrected by irradiating the entire EPID sensitive area with an open field[8] to create a pixel sensitivity map (PSM), and then normalizing each pixel response [32, 97].

2.4.4 Optical glare

The incident radiation interacts with the copper plate of the EPID mostly by Compton scattering. The secondary electrons produced move further to the scintillator and deposit energy within it, generating optical photons. Some of these optical photons might be absorbed by the photodiode, resulting in a blurring effect of the signal recorded by the photodiode array. Contributions can also come from other parts of the EPID panel, such as the covers or supporting materials. The blurring effect, or optical glare, can be corrected by applying point spread function ("blur kernel") to the recorded image [32, 101, 102].

2.5 Intensity Modulated Radiation Therapy (IMRT)

Different modalities of radiation therapy are currently available for patient treatment, characterized by the location of the radiation source with respect to the patient (external or internal), the type of radiation (photons or particles) and the delivery approach. In this work, we focused on Intensity Modulated Radiation Therapy (IMRT), a type of external photon radiation therapy that can be delivered using the medical linacs and MLC system described previously. Details on the other types of radiotherapy can be found elsewhere [83, 85, 94]. Regardless of the modality, the treatment planning usually starts with a CT image of the patient, where the target and OARs are delineated. Recently, functional and/or anatomical imaging techniques have been used in combination with CTs to complement the information provided by the CT images, such as Positron Emission Tomography (PET) and/or Magnetic Resonance Imaging (MRI).

With the introduction of MLC systems, the irradiation field could be shaped to match the projected target contours in the beams-eye-view (BEV). Different beams can be delivered to the target from different gantry angles (i.e. at different BEV), and the dose distribution at the target is the sum of the contributions from each beam. This is

---

[8] In an open field, the radiation beam is not collimated by the MLC, i.e. the leaves and jaws are retracted (or open) to their maximum available position.

the principle of three-dimensional conformal radiation therapy (3D-CRT), illustrated in Fig. 2.18a. In 3D-CRT, dose homogeneity and conformality are achieved by changing the relative contribution (or weight) of each field with respect to the others. Most fields are delivered with flattened beams with uniform intensity across the field, although wedges or compensators can be used occasionally to modify the intensity [83]. This planning strategy is known as *forward planning*, in which the radiation fields are matched to the targets and the resultant dose distribution is subsequently calculated and "tuned" by applying different weighting factors to each field, also called a segment.

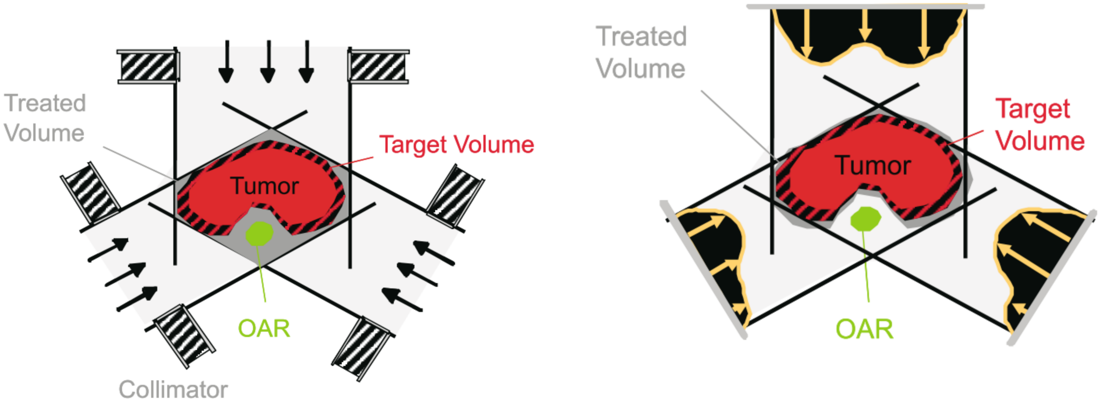

(a) 3D-CRT **(b)** IMRT

Figure 2.18: Simplified representation of the fields and dose delivered to a target volume (in red) from (a) a 3D-CRT treatment and (b) an IMRT treatment. In the former, the OAR (in green) is inside the treatment volume and irradiated with the same dose as the tumor, while in the IMRT plan, the treated volume is more conformal to the tumor, and the OAR is spared. By varying the intensity of the radiation beam across each treatment field, IMRT can deliver more conformal dose distributions in comparison to 3D-CRT. Reprinted with permission from [103].

.

In addition to using different projection angles and different MLC fields, the fluence of the photon beams can be changed or modulated, such that the beams intentionally deliver a non-uniform dose distribution to the target. This is the principle of IMRT, which enables even better dose conformity to the target in combination with better sparing of OARs, as illustrated in Fig. 2.18b. For IMRT planning, the user defines the desired prescription to the target, constraints to surrounding OARs and a fixed number of beam angles. These parameters are passed to an optimizer, which calculates the modulated fields required to achieve the input constraints and prescriptions: this is known as *inverse planning* strategy [83, 87, 94], widely used in clinical TPSs. The optimizer generates an arbitrary number of sub-fields within these beams, also called segments or *control points* (CP), such that the fluence distribution within each beam is no longer homogeneous.

The intensity is modulated using the MLC, in two different modes: static (or *step-and-shoot*) and *dynamic*. In the step-and-shoot mode, first demonstrated in 1994 by Bortfeld et al. [104] and used for this work, one segment is shaped by the MLC, which remains static while the dose corresponding to this segment is delivered. The irradiation stops while the MLC moves to shape the next segment, and the process starts again, until all the segments of one beam are delivered. The sum of the dose increments from each segment creates the intensity modulated beam planned by the TPS, as illustrated in Fig. 2.19 [83]. The gantry then rotates to the next beam angle and the dose delivery follows as described. In the dynamic delivery mode, the field shaping and dose delivery are performed simultaneously for the fixed beam angles. The intensity modulation is achieved by varying the individual velocities of the moving leaves, as illustrated in Fig. 2.20.

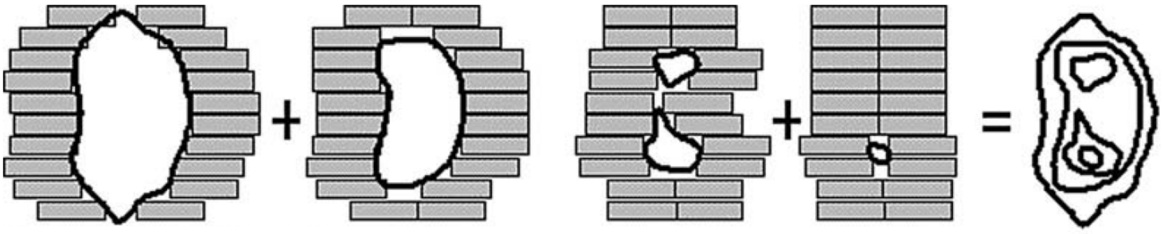

Figure 2.19: Schematic representation of the idea behind step-and-shoot IMRT delivery, for a beam composed of four segments. The leaves are static during irradiation, and the radiation stops while the MLC changes to the next sub-field. The intensity modulation of the beam is a result of the superposition of the four irregularly shaped, partially overlapping sub-fields. Reprinted with permission from [94].

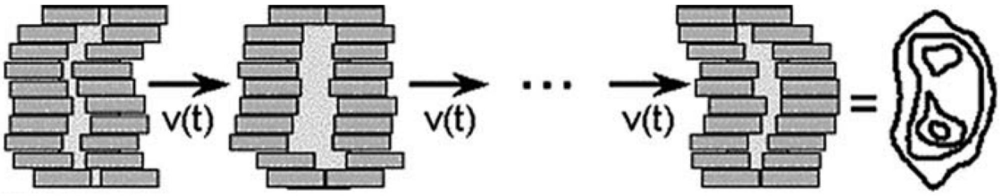

Figure 2.20: Principle of the dynamic IMRT delivery. The field shaping and irradiation are performed simultaneously. The intensity modulation is achieved by varying the velocity of the leaves. Reprinted with permission from [94].

For IMRT treatments, the gantry remains static at one certain angle until all the sub-fields corresponding to that beam are delivered. The extension of IMRT is the Volumetric Modulated Arc Therapy (VMAT), first introduced in 2008 by K. Otto [105], where the gantry rotates continuously during beam delivery. In addition to the MLC leaves motion, the beam modulation can also be performed by changes in the dose rate, gantry speed and/or angular spacing. VMAT treatments are usually delivered in one or more, full or partial gantry rotations (called arcs), and the irradiation is no longer restricted to a few projection angles as in IMRT. This results in more conformal dose distributions to the target and better sparing of OARs. Fig. 2.21 shows the dose distributions obtained with 3D-CRT, IMRT and VMAT for a gastric tumor [106].

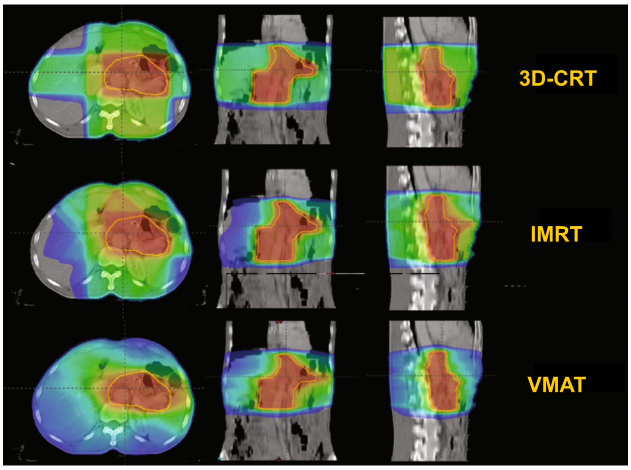

Figure 2.21: Comparison of dose distributions for a gastric tumor, obtained with 3D-CRT, IMRT and VMAT. Reprinted with permission from [106].

2.6 Monte Carlo Method

The Monte Carlo method is a numerical approach to the solution of problems that can be described as a temporal evolution of objects interacting with other objects upon specific probabilities [107]. Such interactions are processed randomly and repeatedly, following a sequence of random (or pseudo-random) numbers and the known probability distribution of the underlying phenomenon. Following the *Law of Large Numbers*, the results from a MC simulation will converge to a reliable mean value only if a sufficiently high number of trials are performed. Finally, due to its statistical nature, MC methods are notably suited for the stochastic process of radiation transport through matter [108]. A detailed mathematical description of the MC method is out of the scope of this thesis, but can be found elsewhere [107].

MC methods have been proven to be efficient in problems with high complexity, like the ones involving the radiation transport and dose deposition in radiotherapy [107]. Moreover, compared to traditional analytical methods, MC simulations consider the detailed geometry and composition of the human body [109, 110]. Therefore, MC methods are particularly useful in the field of medical physics, and have been extensively used to address different problems over 50 years. An overview of the history of MC applications to medical physics is presented in several review articles [107, 111, 112].

In this work, all simulations were performed using the Geant4 (Geometry and Tracking) toolkit. Geant4 is a software toolkit for simulation of particle transport through matter using MC methods, based on the object-oriented programming language C++. This object-oriented approach allows the user to develop highly customized applications, using the set of provided C++ class libraries that describe the geometries, physics models, evaluation tools, among others. Geant4 was developed and is maintained by a worldwide collaboration of scientists and engineers at the European Council for Nuclear Research

(CERN) [113–115]. Geant4 has been used in a wide range of applications, from particle and nuclear physics, to astrophysics and space engineering, along with several topics in medical physics, such as conventional photon radiation therapy, particle therapy, imaging and linac development, among others [111, 116–118].

2.7 Artificial Neural Networks

This section presents a brief introduction to the ideas behind artificial neural networks (ANN) and convolutional neural networks (CNN). A comprehensive introduction and mathematical descriptions can be found elsewhere [119].

An artificial neural network (ANN) is a type of artificial intelligence framework inspired by the propagation of information in the human brain: many complex, interconnected neurons transfer information among each other and work together to perform all types of tasks. Briefly, each neuron receives input signals on the dendrites, and if the sum of all signals exceeds a certain threshold, an output signal is transmitted to other neurons via the axon (Fig. 2.22a). Similarly, ANNs rely on an ensemble of artificial neurons designed to perform certain tasks in analogy to their biological counterpart. Here, one artificial neuron receives a given number of inputs a_i and a bias b, assigns a weight w_i specific to each input a_i, calculates the weighted sum of inputs and bias, then generates an output by applying an activation function σ to the weighted sum, and passes this output to subsequent neurons (Fig. 2.22b). The weighting factor w_i defines the relative importance in the given connection with respect to others, while the activation function introduces non-linearity in the output, required to solve non-linear problems. A variety of activation functions have been used in ANN, however recent works rely on the rectified linear unit (ReLU) functions [119].

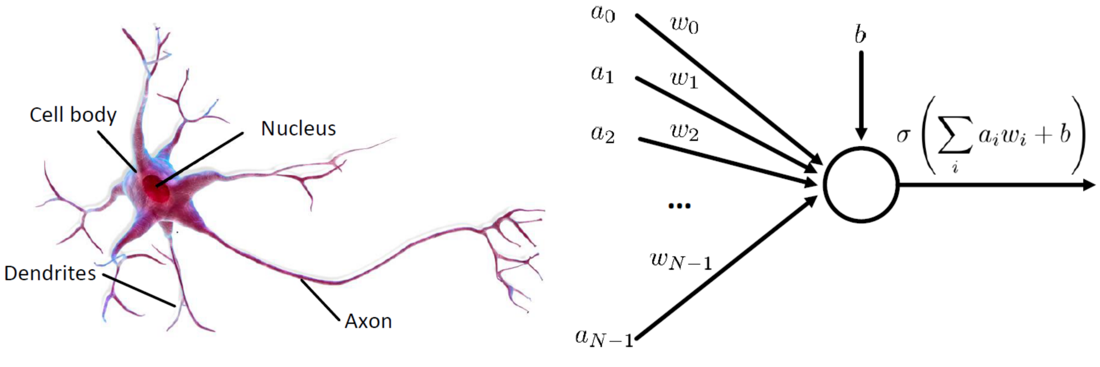

(a) Biological neuron

(b) Artificial neuron

Figure 2.22: Representation of (a) biological and (b) artificial neurons. Reprinted with permission from J. Maier [120].

To describe some of the properties of neural networks, the Feedforward ANN will be considered. In Feedforward ANN the artificial neurons are arranged in different layers, which are connected to the neighboring layers in such a way that the information is propagated only in one direction [121]. The first layer, receiving the external input information, is the input layer, and it passes the information to the next layer without performing any computation. The last layer, called the output layer, provides the output of the ANN. All the layers between them are referred to as hidden layers. Feedforward ANN can have only one input and one output layers, however the number of hidden layers is variable and depends on the complexity of the problem. A schematic representation of a Feedforward ANN architecture with one hidden layer is presented in Fig. 2.23, together with the respective nomenclature [120].

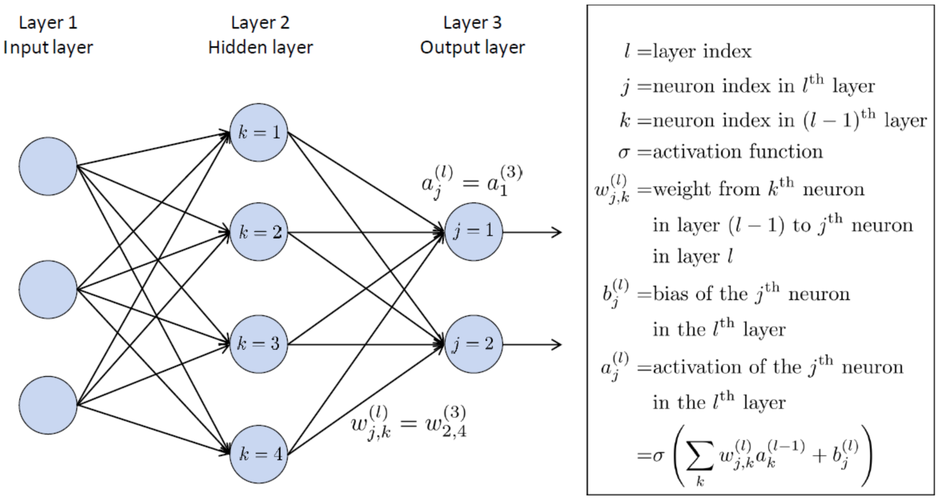

Figure 2.23: An example of the architecture of a feedforward ANN with one hidden layer. The input layer has 3 neurons, while the hidden layer and the output layers have four and two neurons, respectively. The box on the right describes the nomenclature commonly used in ANN. Reprinted with permission from J. Maier [120].

The goal of ANNs is to find a function $\boldsymbol{g}$ able to map from one set of information, given in the input layer, towards another set of information: $\boldsymbol{g} : \boldsymbol{X} \longrightarrow \boldsymbol{Y}$. Let $\boldsymbol{X}$ be the vector containing the input information $\boldsymbol{x}$, $\boldsymbol{Y}$ be the output vector with the output $\boldsymbol{y}$, and $\boldsymbol{\theta}$ the vector with the weights and biases. The network defines a function $\boldsymbol{y} = f(\boldsymbol{x}, \boldsymbol{\theta})$ and optimizes $\boldsymbol{\theta}$ until $f(\boldsymbol{x})$ is a suitable approximation of $\boldsymbol{g}$. In analogy to the human brain, the ANN "learns" the optimal set of parameters $\boldsymbol{\theta}$, based on a subset of examples, in a process called *training*. The subset used for the training process is called the training dataset $\boldsymbol{X}_{train}$.

ANNs can be further categorized regarding the training process in *supervised* or *unsupervised* learning. In the former, the training dataset is labeled, and every training example i is a pair $(f(\boldsymbol{x}^{(i)}), (\boldsymbol{y}_{true}^{(i)}))$, where $\boldsymbol{y}_{true}^{(i)} = \boldsymbol{g}(\boldsymbol{x}^{(i)})$ is the ground true. In contrast, the training dataset for unsupervised learning comprises only $f(\boldsymbol{x}^{(i)})$, with an infinite number of possible solutions which need to be constrained in some other way rather than the use of labels. In this work, supervised learning was used and therefore will be considered hereon. A description of unsupervised learning and its applications can be found elsewhere [119].

To find the optimal values of the parameter vector $\boldsymbol{\theta}$, i.e. to find a suitable $\boldsymbol{y} \approx \boldsymbol{g}$, the neural network works on minimizing (or maximizing) an objective function known as loss function L. The loss function provides a quantitative value for the discrepancies between the output $\boldsymbol{y}^{(i)}$ of the mapping $\boldsymbol{Y}$ and the true value $\boldsymbol{y}_{true}^{(i)}$, as described by Eq. 2.16, where N is the total number of examples in the training set and $L^{(i)}$ is the loss value for each example i. The appropriate choice of the loss function depends on the problem to be solved. In this work the mean square error (MSE) function was chosen. The MSE loss function is then described by Eq. 2.17.

$$L(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=0}^{N-1} L^{(i)}(\boldsymbol{x}^{(i)}, \boldsymbol{y}_{true}^{(i)}, \boldsymbol{\theta}) \tag{2.16}$$

$$L_{MSE}(\boldsymbol{\theta}) = \frac{1}{N} \sum_{i=0}^{N-1} L_{MSE}^{(i)}(\boldsymbol{x}^{(i)}, \boldsymbol{y}_{true}^{(i)}, \boldsymbol{\theta}) = \frac{1}{N} \sum_{i=0}^{N-1} (\boldsymbol{y}^{(i)} - \boldsymbol{y}_{true}^{(i)})^2 \tag{2.17}$$

The loss function can be minimized using the gradient descent algorithm [119]. In gradient descent, the values of the parameters $\boldsymbol{\theta}$ are updated towards lower values of the loss function, determined by the negative derivative of the L with respect to $\boldsymbol{\theta}$. Hence, for every time the values of $\boldsymbol{\theta}$ are updated during the training process, i.e. for every *epoch* n, the update scheme for $\boldsymbol{\theta}$ can be defined by Eq. 2.18.

$$\boldsymbol{\theta}_{n+1} = \boldsymbol{\theta}_n - \alpha \nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta}_n) \tag{2.18}$$

Eq. 2.18 can be interpreted as follows: for the epoch $n+1$, the values $\boldsymbol{\theta}_{n+1}$ will be the values of the previous epoch n, incremented by the derivative of the loss function L with respect to each parameter $\boldsymbol{\theta}$ ($\nabla_\theta L(\boldsymbol{\theta})_n$), scaled by a factor α. The derivative $\nabla_{\boldsymbol{\theta}} L(\boldsymbol{\theta})$ gives the slope of the function, and small increments of $\boldsymbol{\theta}$ in the opposite direction of the slope will reduce the values of L. Hence, for a suitable number of epochs n, the value of the function L will converge to a minimum. The constant α is known as learning rate and controls the speed of convergence.

As the training set $\boldsymbol{X}_{train}$ is a subset of $\boldsymbol{X}$, the training process leads to a mapping function that is optimal for the examples in the training set. Therefore, the performance of the network in generalizing the mapping function to examples unseen during the training process needs to be evaluated with a different dataset, called the validation or test set $\boldsymbol{X}_{test}$. The test set contains examples that were not included in $\boldsymbol{X}_{train}$. A good performance on the test set does not guarantee that the network performs well for any unknown input, however it is a good indication that a suitable mapping has been learned.

2.7.1 Convolutional Neural Networks

Convolutional neural networks (CNN) are a type of ANN specialized in processing data with grid-like topology, such as 2D and 3D images [119]. The idea behind CNN was inspired by how the human brain processes an image in the visual cortex. Each neuron works in a particular region of the visual field, called the receptive field, and different neurons are connected to each other in such a way to cover the entire visual field. Two types of cells are involved in the process: the first type has a maximum activation when the receptive field presents patterns with a particular orientation, while the second type responds similarly in a larger receptive field, but with lower sensitivity in the position of the pattern. CNNs are able to perform a similar task using two special types of layers, namely the *convolution* and the *pooling* layers.

A convolution layer is formed by individual neurons with learnable weights and biases. The neurons are arranged in many small squared templates, called “convolution kernels”, that restrict the receptive field of a single neuron to only a small region of the previous layer. The kernels move across the entire input with a certain step size (stride), looking for patterns. All neurons in the same layer use the same weights, which can be interpreted as the entries of the kernel K. Hence, the output D of one single neuron applied to a 2D input S can be interpreted as a convolution, defined in Eq. 2.19, where i and j are the pixel indices, σ is the activation function and b is the bias. Fig. 2.24 illustrates one convolution layer.

$$D_{i,j} = \sigma((S * K)_{i,j} + b) \tag{2.19}$$

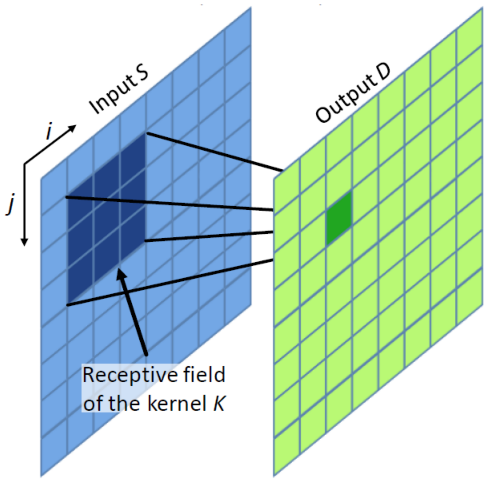

Figure 2.24: Schematic representation of one convolution layer with a 3 × 3 kernel. Reprinted with permission from J. Maier [120].

The output D, also known as feature map, indicates if and where a certain feature, encoded by the kernel K, is present in the input. By stacking several convolution layers, the network is able to recognize sophisticated patterns from the input, with deeper layers able to extract more complex features based on the patterns identified by previous layers. Finally, as the weights of the kernels are learnable parameters, the kernels are not predefined but rather learned during network training.

The pooling layer, usually placed between successive convolution layers, reduces the spatial dimension of the input by reducing the number of parameters, while extracting only the dominant features. Similar to the convolution layer, a pooling filter travels across the input and applies an aggregation function on the values within its receptive field. The most common types of pooling operations are the max pooling, which selects the maximum value of the receptive field as the output of the operation, and the average pooling, which outputs the average value. Fig. 2.25 illustrates a max pooling layer with a 2 × 2 filter.

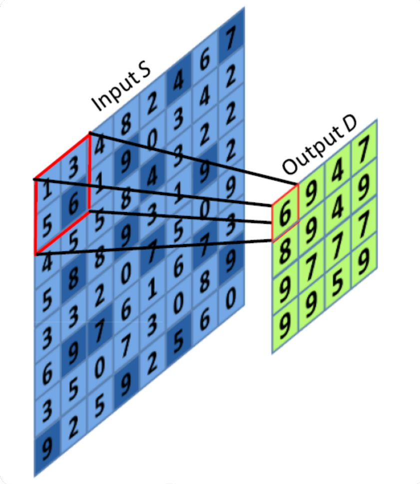

Figure 2.25: Schematic representation of one max pooling layer with a 2 × 2 filter. Reprinted with permission from J. Maier [120].

CHAPTER 3

Linac Monte Carlo Model

For training the machine learning network to reproduce dose distributions inside a complex patient geometry from EPID signals, a reliable dose distribution is needed to serve as the target for the training process. In this work, Monte Carlo (MC) techniques have been used to generate the required target dose distributions. MC has been widely used in radiotherapy [111, 112] and is considered the most accurate tool in determining the energy deposited to a medium by ionizing radiation. For problems with a large number of parameters, such as the dose received by patients due to radiotherapy treatments, MC is more advantageous than analytic algorithms, especially at regions with high heterogeneity [107, 122–124]. However, the quality of the output is strongly dependent on the accuracy of the MC model of the linear accelerator (linac), and discrepancies between the equipment and its MC model will translate into errors on the patient dose calculation [125]. A precise description of the equipment's geometry, and of the electron beam hitting the target, is therefore crucial [126]. The necessary detailed geometric information of the linac's components is often a commercial secret and is disclosed by vendors only upon non-disclosure agreements, making the faithful modeling of the equipment a complicated task.

The equipment used in this project was an Elekta Synergy®, installed at the Klinikum der Universität München (Klinikum Grosshadern). The linac is equipped with an Elekta Agility™ Multi-Leaf collimator (MLC), composed of 80 interdigitating leaf-pairs and two perpendicular jaws. A MC model of this linac was built using the Geant4 toolkit [113], initially with Geant4 version 10.01.p02 and then updated to version 10.05.p01.

In this project, the modeling of the linac head was divided into two parts: 1. The patient-independent part, or static part, which consists of all static elements in the linac head that are not dependent on a specific irradiation, i.e. which geometry remains unchanged, and 2. The patient-dependent part, or dynamic part, which comprises the MLC collimator (with all its leaves and diaphragms), that shapes the irradiation fields based on the treatment plan of the patient. Two individual Geant4 models were created for the static and dynamic parts of the linac head. The static model was used to produce Phase

Space (PhSp) files, which served as particle generators for further simulations performed with the dynamic model (subsection 3.1.1). The dynamic model was used to simulate the dose deposited to a relevant scoring volume, such as a virtual water phantom or a patient geometry reconstructed from computed tomography (CT) data.

For a proper representation of the equipment, the static part was modeled in details based on confidential information disclosed by the vendor through a non-disclosure agreement (NDA) between Elekta Limited, Klinikum Grosshadern and the Department of Medical Physics from the Ludwig-Maximilians-Universität München (LMU). As this geometric information is not always available to the users, a different approach for modeling the static linac part was developed as a side project (Chapter 4), based on the manipulation of validated PhSp files provided by the International Atomic Energy Agency (IAEA). The dynamic part was created as an analytical model of the MLC parameters, using only information freely available in the literature.

This chapter describes the Geant4 models of the static and dynamic parts of the linac head. The use of PhSp files was crucial, therefore the characteristics and importance of PhSps are also presented.

3.1 Methodology

3.1.1 Phase Space Files

A Phase Space (PhSp) is a file, created usually as a binary file, containing relevant information of particles crossing a reference surface (the PhSp plane). When a PhSp is created, different parameters of the particles of interest can be recorded, depending on the application and interest of the user.

In this work, the PhSp plane was located between the static part and the dynamic part of the linac head, as illustrated in Figure 3.1. The static part consists of the electron beam target, primary collimator, the flattening filter, the MU ionization chambers (monitor chambers), the backscatter plate and the mylar mirror. It can also be referred to as patient-independent part, as it does not change during patient treatment, except for irradiations with different energies and/or with electrons, in which case different flattening filters are required. As in this work only treatments with photon beams with a fixed nominal energy of 6 MV were used, all simulations were performed with the same flattening filter. The dynamic part, or patient-dependent part, encompasses the Multi-Leaf Collimator (MLC) and the two diaphragms (also called jaws). The dynamic part is linked to the treatment of the patient, as it shapes the fields to be irradiated, and is therefore patient-specific. Moreover, during the irradiation of a single patient, several fields might be involved, during which the MLC configuration is constantly changing.

Modeling the radiation transport through all the complex components of a linac head is very time consuming. Since the static part of the linac does not change, the output of

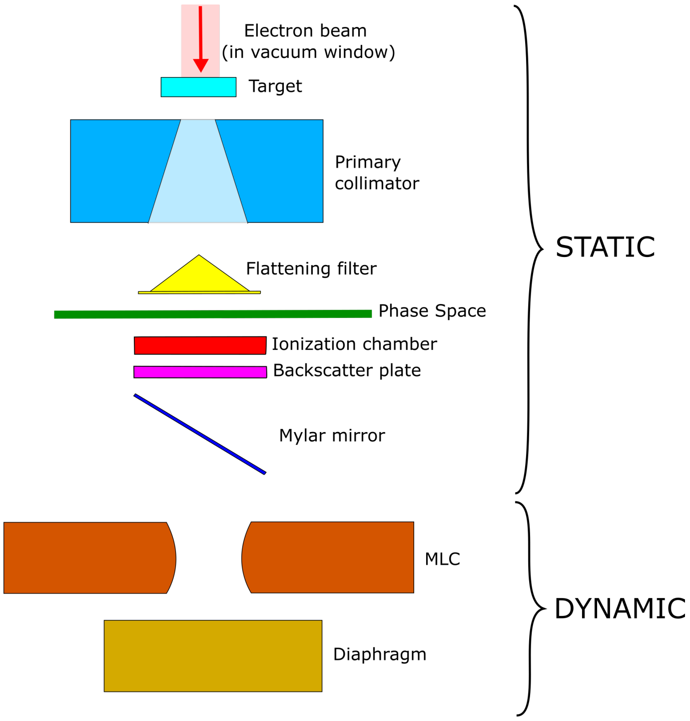

Figure 3.1: Schematic representation of the linac Monte Carlo (MC) model. The static and dynamic parts are depicted, together with the location of the phase space (PhSp) plane. Figure is not to scale.

particle simulation through this part can be considered stable, as long as a sufficiently large number of particles are simulated. To save computational time, a PhSp can be created between the static and dynamic parts (as seen in green in Figure 3.1), recording all relevant information of particles crossing the PhSp plane. This PhSp can be used as a replacement for the static components of the head, serving as the primary particle generator for future simulations. This practice improves simulation time, as the radiation transport through the static parts of the linac does not need to be repeatedly simulated [127]. Several papers have investigated the use of PhSp files for linac MC modeling [112, 128, 129]. Although full-linac simulations can also be potentially time-efficient if proper variance reduction techniques are adopted [130], in this work no variance-reduction techniques were employed. Instead, PhSp files were created for the specific linac of interest and used to replace the static part of the linac head.

Although the components of the static part of the linac head do not change, the incident electron beam characteristics, such as its energy and shape, can vary from their nominal values [131, 132]. It has been reported that the quality of the dose simulated with MC depends on the proper description of the incident electron beam [125, 126, 133]. Namely the mean energy $\overline{E}$, the energy spread σ_E and the shape (focal spot) of the beam incident on the target needed to be properly modeled and fine tuned for the equipment used in this work.

Starting from the nominal values, a range of values for the incident electron beam parameters was used and several combinations were investigated. For each combination, a PhSp file containing relevant information of the particles crossing the PhSp plane was generated and used to simulate percentage depth dose (PDD) and lateral profiles for several open squared fields in a water phantom. The simulated profiles were compared to measurements performed under the same conditions. The parameters of the electron beam were then iteratively tuned until simulations and measurements agreed. Therefore, the use of PhSp files was crucial in optimizing the initial electron beam parameters. The optimization process is presented in 3.1.7.

3.1.2 Geant4 Patient-Independent Model

The static part of the linac head was modeled following the geometric and material descriptions provided by the vendor through the NDA agreement. The geometric information of each component, as well as its corresponding material composition, were described in the user defined *DicomDetectorConstruction* class.

The world volume, representing the treatment room, was constructed as a 400 × 400 × 400 cm^3 air volume. The air material $G4_{AIR}$ was loaded from the database of the National Institute of Standards and Technology (NIST) [134] using the *G4NistManager* class. The isocenter of the treatment room corresponds to the origin of the Geant4 coordinate system, as illustrated in Fig. 3.2, where the world's coordinate system is denoted by (X_w, Y_w, Z_w) and the isocenter is at $(X_w, Y_w, Z_w) = (0, 0, 0)$. For a patient in the head-first supine position (HFS), the X_w axis corresponds to the left-right direction, the Y_w axis is the posterior-anterior direction and the Z_w axis is the cranio-caudal direction (Fig. 3.2). For the gantry positioned at 0 degrees, the beam travel direction is $-Y$ (red arrow in Fig. 3.2).

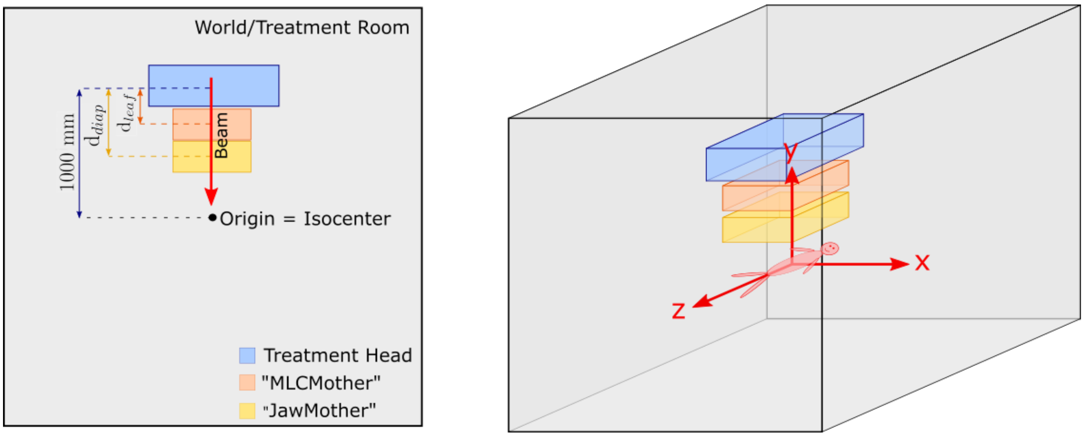

Figure 3.2: 2D representation of the treatment room (world, in gray), treatment head (in blue), MLC mother (in orange) and diaphragms mother (in yellow) volumes in the Geant4 model of the linac (left figure). The coordinate system relative to a patient positioned in a head-first supine (HFS) position is displayed. The images are not to scale.

Within the treatment room volume, a treatment head volume was created as a 100 × 110 × 100 cm^3 box, also filled with air. The treatment head volume was moved

100 cm in the positive Y_w axis direction, so that its origin was located exactly 100 cm from the isocenter of the treatment room (Fig. 3.2). Therefore, the treatment head coordinate system can be denoted by (X_h, Y_h, Z_h), with $X_h = X_w$, $Y_h = Y_w + 100$ cm and $Z_h = Z_w$. The configuration of the treatment head volume as the mother volume for all the components of the static part facilitates their construction in the coordinates given by the vendor.

As illustrated in Figure 3.1, the static model comprises the vacuum window, the electron target on its copper holder, the primary collimator, the flattening filter, the ionization chamber array, the backscatter plate and the mylar mirror. Each component was modeled as boolean solids using the basic Geant4 geometries ($G4Box$, $G4Trap$, $G4Tubs$ and $G4Cons$), following the geometric description provided by the vendor, and were placed inside the treatment head volume. The target was positioned at the center of the treatment head volume $(X_h, Y_h, Z_h) = (0, 0, 0)$, thus being 100 cm from the isocenter of the treatment room. With this configuration, the linac source, i.e. the target where the electron beam hits, is positioned at a source-to-axis distance (SAD) of 100 cm. This configuration corresponds to the physical treatment room. The remaining components were positioned with respect to the target following the vendor's description.

Finally, the static model also includes the PhSp plane, defining the place where the PhSp files were generated. The PhSp plane is a solid box created inside the *DicomDetectorConstruction* class, with 30×30 cm^2 area in the XZ plane, and 5 μm thickness. The PhSp plane (green in Fig. 3.1) is positioned at 16 cm from the center of the treatment head volume at $(X_h, Y_h, Z_h) = (0, 16, 0)$ cm, between the flattening filter and the IC chamber arrays. For absolute dose calibration, the dose deposited at the IC arrays (monitor chambers) might need to be considered (details in section 3.1.10), but since the IC chambers and the mylar mirror result in negligible attenuation of the photon beam [126], the PhSp plane, and the resultant PhSp files created there, were placed in this position. The purpose of the PhSp plane is to define an area where the characteristics of the particles crossing it should be scored; therefore, when the particles cross this solid, they must not interact in a way that modifies their characteristics. Hence, the PhSp plane is filled with vacuum: the code notices particles entering and exiting the solid, but no changes on the beam parameters occur in this interaction.

The electron beam was modeled using the *G4GeneralParticleSource* (GPS) class with the user-defined *DicomPrimaryGeneratorAction* class. The GPS allows the definition of several characteristics of the primary particle beam, such as the type of particle, its energy, its spatial and angular distribution, among others. The GPS can be conveniently configured via macro-based inputs commands, enabling parallel simulations with different configurations of the primary beam. For PhSp generation, all parameters of the initial electron beam hitting the target were defined inside macro files, and passed as input to the executable via the command line. More details of PhSp generation are presented in section 3.1.8.

The static model was used only for PhSp generation, so all simulations were performed with a fixed gantry position of 0 degrees. Irradiations at different gantry angles were performed for validation of the dynamic part, where the static part is replaced by a validated PhSp file and gantry rotations are properly accounted for.

3.1.3 Geant4 Patient-Dependent Model

The dynamic part of a linac, or patient-dependent part, corresponds to the beam shaping components, i.e. the Multi-Leaf collimator (MLC). The components of the static part of the linac were not included in the patient-dependent model. Instead, they were replaced by PhSp files generated between the flattening filter and the IC chamber arrays, using the patient-independent part. For consistency, the IC arrays and the mylar mirror, components of the static part, were included in the dynamic model as well; however their presence does not influence the photon beam [126]. The PhSp file acts as the primary particle source, defining the type of particles to be generated, their initial positions (starting at the PhSp plane position), their energy and angular distribution. The names of the PhSp files used were given as input inside the macro file via an user-defined command (section 3.1.6) and read inside the model by the *DicomPrimaryGeneratorAction* class.

3.1.3.1 The Agility MLC

The Elekta Synergy® linac used throughout this work is coupled to an Agility™ MLC system, composed of two banks of 80 interdigitating leaf pairs. Each leaf bank is mounted on a dynamic leaf guide (DLG) which can move simultaneously with the leaves at a maximum speed of 30 mm/s. Additionally, the DLGs can move up to 15 cm over the central axis (measured at isocenter). The 160 leaves can move independently and can travel up to 20 cm within the DLG (measured at isocenter), at a maximum speed of 35 mm/s [93, 135, 136]. Thus, in combination with the DLG, the leaves' range of motion is from -15 cm (across the central axis) to 20 cm, with a maximum speed of 65 mm/s. Table 3.1 lists the dynamic properties of the Agility™ MLC.

To describe the parameters of the leaves, it is convenient to use the definitions of the task group 50 of the American Association of Physicists in Medicine (AAPM) [92], described in Fig. 3.3. The leaf width is the smallest dimension of the leaf perpendicular to both the beam propagation direction and the leaf's direction of motion; the leaf length is the dimension parallel to the direction of leaf motion; the height of the leaf is the dimension of the leaf along the beam propagation direction. Additionally, the leaf sides are the surfaces in contact with the adjacent leaves, and the leaf end is the surface of the leaf which shapes the irradiation field and gets closer to the leaf on the opposite bank.

The leaves are 90 mm high, with a nominal projected width of 5 mm at the isocenter, and have a curved end design, following a curvature of radius R_{leaf} (Fig. 3.4). The leaves have a trapezoidal shape following the beam divergence, and adjacent leaves are positioned

Table 3.1: Dynamic parameters of the components of the Agility$^{\mathrm{TM}}$ MLC. All values are defined at isocenter [93, 135, 136].

Dynamic Property	Value
Maximum field size (cm^2)	40×40
DLG range with respect to central axis (cm)	5 to 20
Leaf range with respect to DLG (cm)	-20 to 0
Leaf range with respect to central axis (cm)	-15 to 20
Diaphragm range with respect to central axis (cm)	-12 to 20
DLG maximum speed (mm/s)	30
Leaf maximum speed (mm/s)	35
Leaf and DLG combined maximum speed (mm/s)	65
Diaphragm maximum speed (mm/s)	90

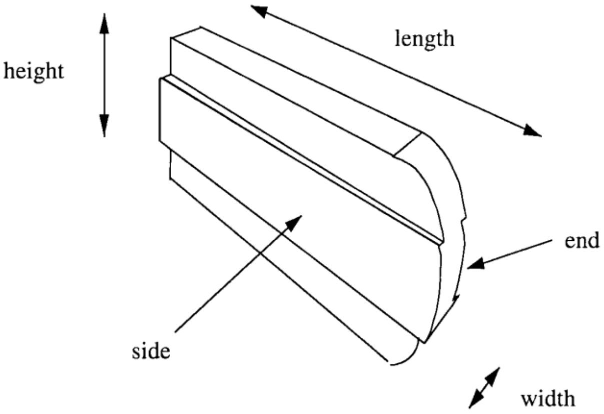

Figure 3.3: Schematic representation of a leaf with rounded ends and tongue and groove, illustrating the terminology defined by the AAPM Task Group 50 [92] and used in this work. Reprinted with permission from [92].

side by side with a 90 μm gap, represented in green in Fig. 3.5a. The interleaf leakage is minimized by the introduction of a small tongue and groove (T&G) on the leaves' sides, as shown in Fig. 3.5b. To further reduce the interleaf leakage, the leaf banks are rotated by a certain angle, such that the leaves are focused to a virtual focus at a distance α from the target and defocused with respect to the radiation source (Fig. 3.6) (called virtual tongue-and-groove, vT&G) [93, 137]. The combination of T&G and vT&G prevents the beam to pass directly through the leaf gaps. Two extra shielding blocks are positioned next to both sizes of the leaf banks (pink blocks in Fig. 3.6). The explanation and respective nominal values of the symbols in Figs. 3.4, 3.5 and 3.6 are listed in Table 3.2.

In addition to the two leaf banks and the DLGs, the Agility$^{\mathrm{TM}}$ MLC also includes two shaped collimators, called diaphragms or jaws, which move in the orthogonal direction of the leaves (Fig. 3.1). The diaphragms can move from -12 cm (across the central axis) to 20 cm, at a maximum speed of 90 mm/s (Table 3.1) [93, 135, 136]. Unlike previous MLC models, the Agility$^{\mathrm{TM}}$ system has no backup collimators moving in the direction of

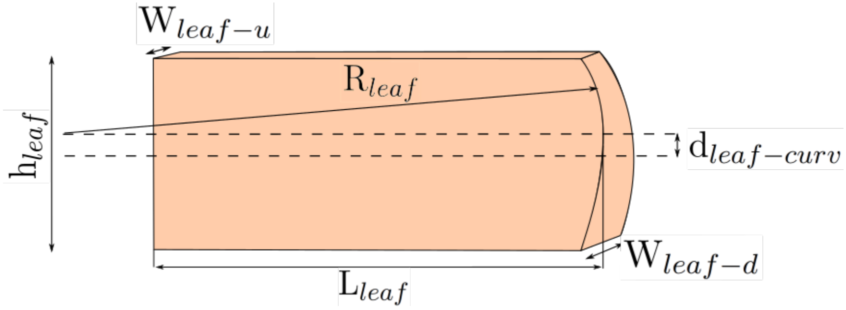

Figure 3.4: Schematic representation of the lateral view of one leaf, showing the geometry used for Monte Carlo (MC) simulations [136]. The representation is simplified and does not include the tongue and groove (described in detail in another figure). The description and respective values of W_{leaf-u}, W_{leaf-d}, $d_{leaf-curv}$, R_{leaf}, L_{leaf} and h_{leaf}, used in the MC model, are listed in Table 3.2. The figure is not to scale.

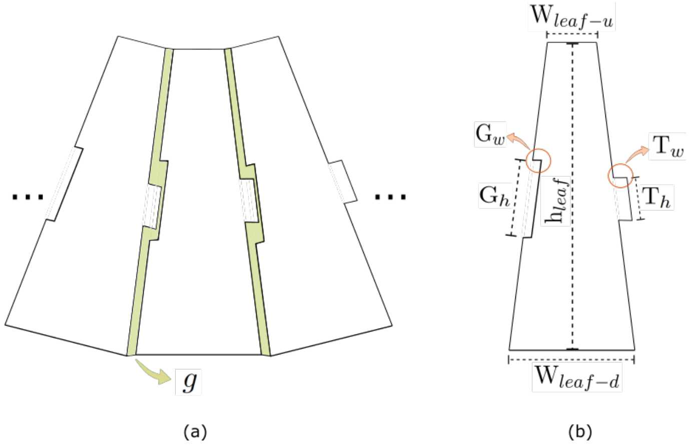

Figure 3.5: (a) Schematic representation of the constant gap g between adjacent leaves (in green). (b) Frontal view of one leaf, showing its trapezoid shape and the geometry of the tongue and groove (T&G). The description and respective values of W_{leaf-u}, W_{leaf-d}, h_{leaf}, G_w, G_h, T_w, T_h and g used in the MC model are listed in Table 3.2. The figures are not to scale.

leaf motion [135]. The diaphragms have a novel “Y-sculpted design”, with a full height of 77 mm in a Y-shaped zone along the field-defining rounded ends and along the central section (Fig. 3.7). Outside the thicker Y-shaped area is a region where the leaves always provide extra shielding, and therefore the diaphragm’s height is reduced to 30 mm. This novel design reduces the weight of the diaphragms and enables rapid movements and decreasing treatment times [93]. Similar to the leaves, the diaphragms have a rounded end following a curvature of radius R_{diap} as shown in Fig. 3.8. The explanation and respective nominal values of the symbols in Fig. 3.8 are listed in Table 3.2.

The 160 individual leaves are divided in two leaf banks, and inside each leaf bank, they are positioned side by side across the Z direction. The leaves move along the X direction, while the diaphragms move in the perpendicular direction, along the Z axis. The Agility^{TM} MLC has two additional static structures, positioned left and right of each leaf bank (pink, in Fig. 3.6). Fig. 3.9 illustrates the coordinate system from the beam’s eye view, as used in the Geant4 model. The crossline (patient’s left-right, for a patient

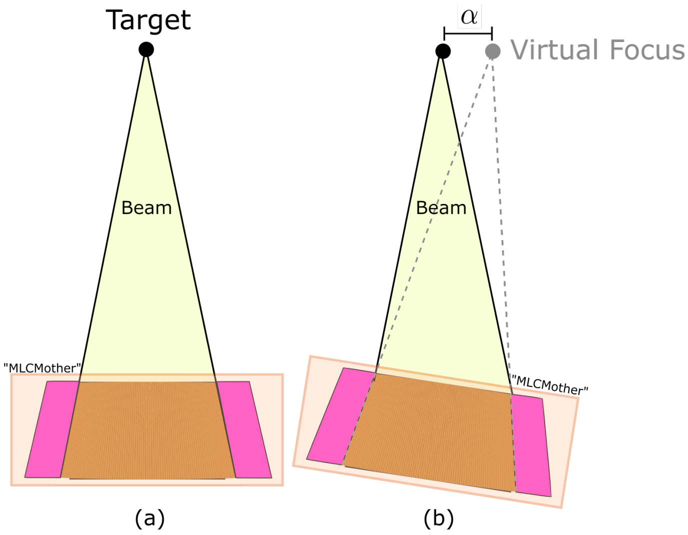

Figure 3.6: (a) Representation of the leaves positioned following the beam divergence; (b) The subsequent rotation of the leaf bank ("MLCMother" volume) towards a virtual focus, located at a distance α from the target, blocking the radiation traveling through the leaf gaps and creating a virtual tongue and groove (vT&G). The two extra blocks, positioned left and right of each leaf banks (pink), were also included in the MC model. The description and respective value of α used in the MC model is listed in Table 3.2. The figures are not to scale.

in HFS position) and the inline (gun-target) directions are also shown. In this coordinate system, the radiation beam travels towards the -Y axis (entering the page).

3.1.3.2 Leaf and Diaphragm Modeling

At the start of this project, the geometric information on the Elekta Agility$^{\mathrm{TM}}$ MLC was not available. Without detailed information of the geometry and material composition of the components, the quality of the MC model was compromised. As an alternative, an extensive literature review was performed in order to find and collect any available information regarding the Agility$^{\mathrm{TM}}$ MLC system.

The gathered information, with respective references, is listed in Table 3.2, and was the basis for a first model of the leaves in Geant4. These are mostly nominal values and can vary among different machines of the same model [141]. Therefore, some parameters need to be fine-tuned for the specific machine of interest, namely the height of the leaf, the density of the tungsten alloy used to model leaves and diaphragms, and the shift to the virtual focus (more details in 3.1.9). Following the descriptions in Figs. 3.4 and 3.5, the leaves were modeled as boolean solids in Geant4, using the basic Geant4 geometries (*G4Box*, *G4Trap* and *G4Tubs*). For the tongue and groove modeling, the T_w and G_w values were not found in the literature. Instead, they were chosen such that the tongue of one leaf is large enough to block the radiation through the gap with the adjacent leaf, but without them touching each other. Since the tongue has to cover the leaf gap, the value for T_w needs to be larger than 90 μm, and the value for G_w can be just enough to avoid friction between adjacent leaves. The values chosen and listed in Table 3.2 fulfill

Table 3.2: List of parameters' symbols, meaning and nominal values found in the literature and used for the construction of the Agility^{TM} MC model. The reference of each parameter is also given. Parameters marked with * are machine-specific and were optimized.

Parameter	Meaning	Nominal Value	Reference
F_{max}	Maximum field size	400×400 mm^2	[93, 135, 138]
F_{min}	Minimum field size	5×5 mm^2	[139]
N_{leaf}	Number of leaf pairs	80	[93, 138–141]
iso	Distance between target and isocenter	100 cm	[136, 140]
d_{leaf}	Distance from target to center of leaf bank	35.68 cm	[93, 135, 136, 142]
d_{diap}	Distance from target to center of diaphragms	47.05 cm	[93, 142]
$d_{leaf-curv}$	Position of leaf curvature with respect to leaf center	7.5 mm	[136, 140]
$d_{diap-curv}$	Position of diaphragm curvature with respect to diaphragm center	3.5 mm	[136]
R_{leaf}	Leaf's curvature radius	170 mm	[135, 136, 138, 140]
R_{diap}	Diaphragm's curvature radius	135 mm	[136]
L_{leaf}	Leaf's length	155 mm	[136, 140]
L_{diap}	Diaphragm's length	220 mm	Not found
W_{leaf-u}	Leaf's width up	1.47 mm	[136, 140]
W_{leaf-d}	Leaf's width down	1.91 mm	[136, 140]
W_{diap}	Diaphragm's width	235 mm	Calculated
$h_{leaf}*$	Leaf's height	90 mm	[93, 135, 136, 138–140]
h_{diap}	Diaphragm's height	77 mm	[93, 136, 143]
$\rho*$	Material density of the tungsten alloy for leaves and diaphragms	18.0 g/cm^3	[136, 144]
$\alpha*$	Shift to virtual focus (vT&G)	3.25 mm	[93, 138]
$W_{leaf-iso}$	Leaf width at isocenter	5 mm	[93, 135, 138–140, 143]
g	Interleaf gap	90 μm	[93, 140, 144, 145]
T_h	Tongue height	1 mm	[137]
G_h	Groove height	2 mm	[137]
T_w	Tongue width	0.098 mm	Not found
G_w	Groove width	0.03 mm	Not found
δ	Gap between closed opposing leaves (end)	1 mm	[136]
ϕ	Maximum angular Spread of the radiation beam	27.5 °	[136, 140]

both requirements, and a validation of the modeled tongue and groove geometries was performed against measurements (details in 3.1.12).

Once the leaf solids were modeled, they were placed inside a dedicated mother volume (MLCMother) filled with air (Fig. 3.6), with a position relative to the center of the MLCMother volume, side by side along the Z axis, with the constant gap of $g = 90$ μm between adjacent leaves (Fig. 3.5). The leaves move along the X axis (Fig. 3.9), and their X and Z positions, X_{Pos} and Z_{Pos} respectively, are determined by analytical expressions,

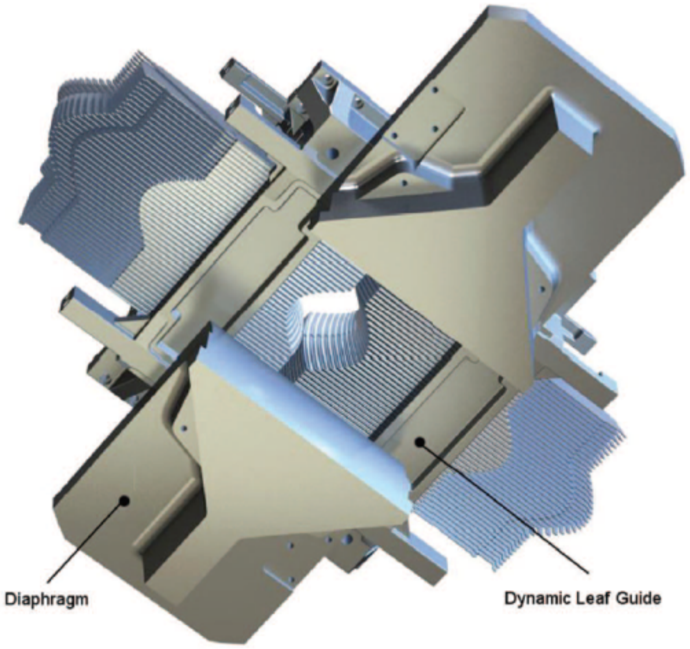

Figure 3.7: Schematic representation the Agility^{TM} MLC system from the patient's eye view, showing the leaves, the dynamic leaf guides and the sculpted Y-shaped diaphragms. Reprinted with permission from [93].

described in details in the following. By matching the information given in Table 3.2 and illustrated in Fig. 3.2, the distance from the target to the center of the MLC mother is set to d_{leaf} = 35.65 cm. Hence, the center of the MLCMother volume is positioned at a distance $iso - d_{leaf}$ = 64.32 cm from the isocenter of the treatment room (in the Y direction). The Y position Y_{Pos} for all leaves was set to zero, such that their centers are positioned at the distance d_{leaf} from the beam target. Finally, the vT&G can be modeled in Geant4 by simply rotating the MLCMother volume against the beam divergence around the X axis, by an angle determined by α (Fig. 3.6).

Similarly, the diaphragms were modeled as boolean solids, using *G4Box* and

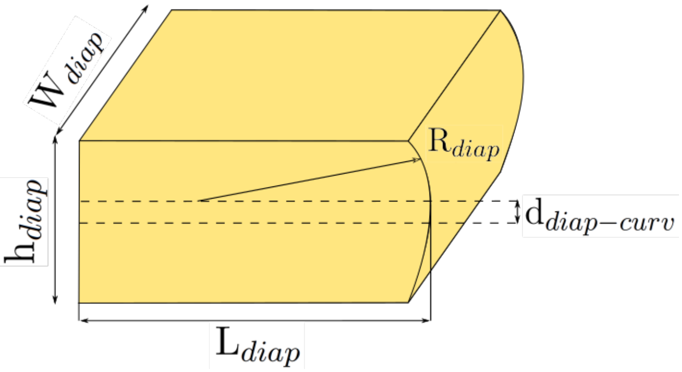

Figure 3.8: Schematic representation of the lateral view of one diaphragm, showing the geometry used for Monte Carlo (MC) simulations [136]. The description and respective values of W_{diap}, L_{diap}, h_{diap}, R_{diap}, and $d_{diap-curv}$ used in the MC model are listed in Table 3.2. The figure is not to scale.

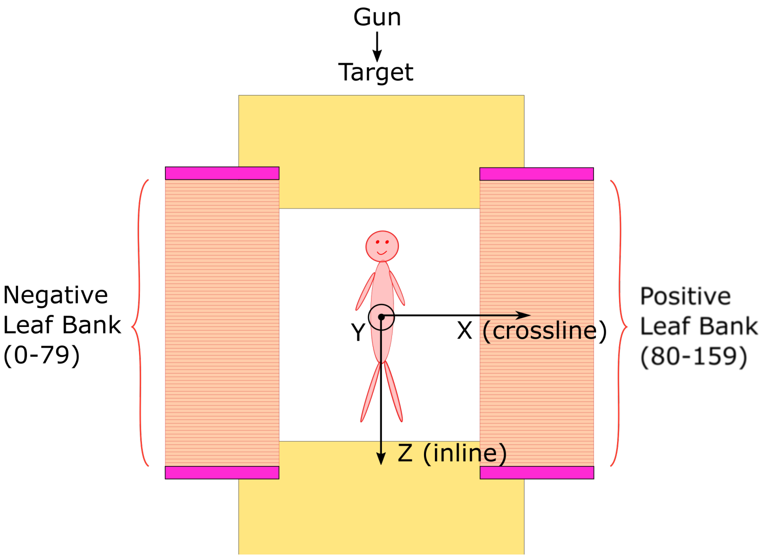

Figure 3.9: Representation of the coordinate system at the beam's eye view, used for MC modeling of the MLC system. The X axis (crossline direction) refers to the patient's left-right, while the Z axis (gun-target direction) matches the patient's cranio-caudal direction, for a patient in the HFS position. The Y axis (exiting the page) is the patient's posterior-anterior direction. The first 80 leaves (0 to 79) are on the negative side of the X axis (negative leaf bank), while the last 80 leaves (80 to 159) are located on the positive side (positive leaf bank). The figure is not to scale.

G4Tubs, following the description given in Fig. 3.8 and the values listed in Table 3.2. The diaphragms have a thicker Y-shaped region (Fig. 3.7). For MC simulation purposes, however, the diaphragm model can be simplified to have an uniform height of $h_{diap} = 77$ mm, as shown in Fig. 3.8 [136, 143]. Despite the extensive literature review, the values for W_{diap} and L_{diap} were not found. W_{diap} can be determined by the angular spread of the radiation beam ϕ, as illustrated in Fig. 3.10 and described by Eq. 3.1. As can be observed in Fig. 3.8, the length of the diaphragm L_{length} is larger than the radius of curvature R_{diap}. The diaphragm needs to be long enough to block the radiation when the diaphragm is on its maximum position over the central axis (120 mm), and still block the region behind it. Since the Geant4 model has no extra structure limiting the length of the diaphragms, L_{diap} was set to 220 mm. In reality, however, the length of the diaphragm is limited by other objects, as e.g. the cover of the linac head, and the value of L_{diap} might differ from the value used in this work. The diaphragms were placed inside a dedicated mother volume (JawMother) filled with air, so the position of each diaphragm is with respect to the isocenter of the JawMother volume. The X and Y positions of the diaphragms are fixed to zero to match the center of the JawMother volume ($X_{PosD} = 0$ and $Y_{PosD} = 0$). The position along the beam defining direction Z_{PosD} (Z axis, Fig. 3.9) was determined with analytic expressions, explained in the following. The JawMother volume is positioned at 52.95 cm from the isocenter of the treatment room volume (Y direction), and consequently at 47.05 cm from the beam target, as listed in Table 3.2.

$$W_{diap} = 2 \times d_{diap} \times tan\left(\frac{\phi}{2}\right) \tag{3.1}$$

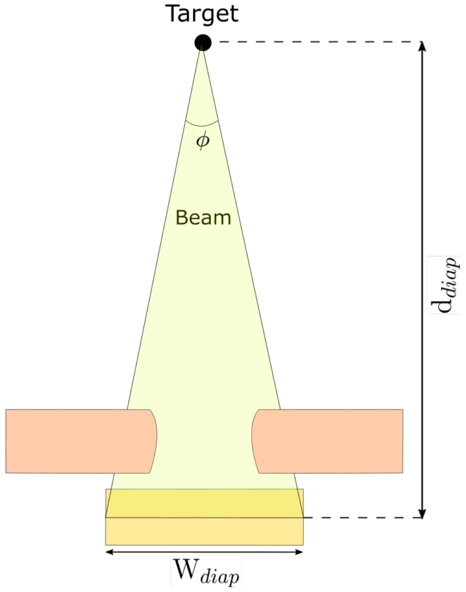

Figure 3.10: Schematic representation of the geometry used to calculate the width of the diaphragms W_{diap}, based on the maximum photon beam aperture ϕ. The figure is not to scale.

In this initial model, the material used for modeling the leaves, diaphragms and extra blocks was a tungsten alloy (tungsten: 75%, nickel: 3.75%, iron 1.25%) with density ρ of 18.0 g/cm3, as presented in Table 3.2 ([136, 144]). This value however is a nominal value, and variations might exist among equipment of the same model. The same applies to the values for leaf height h_{leaf} and the shift to virtual focus α (for vT&G) [141, 144]. Hence, h_{leaf}, ρ and α were set as adjustable parameters in the MC model, using user-defined commands (subsection 3.1.6), and iteratively fine-tuned until good agreement of simulations and measurements, performed in a water phantom under the same conditions, was achieved (details in 3.1.12).

3.1.3.3 Leaves: angular rotation θ_i

As previously described (subsubsection 3.1.3.2), the leaves are positioned side by side along the Z axis (Fig. 3.9). Due to the trapezoidal geometry and the constant gap between adjacent leaves, the leaves are rotated around the X axis by an angle θ_i, following the beam divergence as illustrated in Fig. 3.11. To define the Z_{Pos} of the leaves, their angular rotation θ_i with respect to the Y axis (CAX) needs to be determined first. For this purpose, it is considered that the outer sides of the first and last leaves in the leaf banks match perfectly the maximum field size F_{max} of the Agility^{TM} MLC. Thus, θ_{Fmax} in Fig. 3.11 is the angle of the outer side of the last leaf (Eq. 3.2), and θ_i is the angle of the middle of the last leaf, both with respect to the CAX. The trapezoid itself has an angular aperture of θ_t (Eq. 3.3).

$$\theta_{Fmax} = \arctan\left(\frac{0.5F_{max}}{iso}\right) \tag{3.2}$$

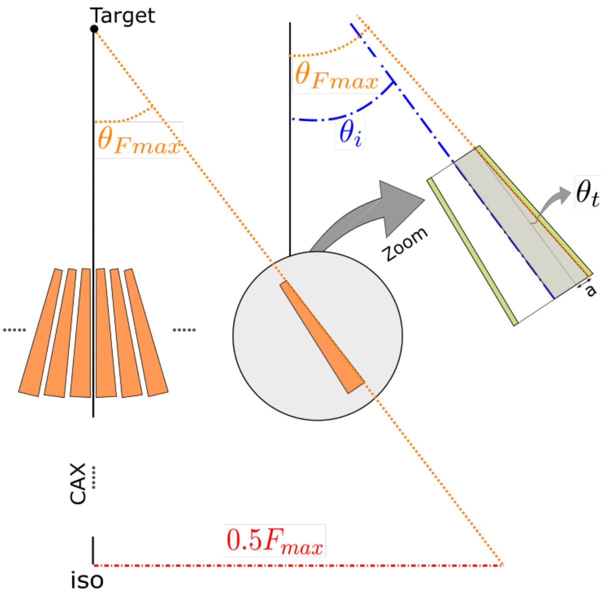

Figure 3.11: Schematic representation of the leaves' individual rotation angle θ_i with respect to the central axis. The leaves are rotated following the beam divergence. For illustration purposes, only half of the maximum field size is displayed ($0.5F_{max}$). The last leaf is shown in details, with its outer side perfectly matching the beam defined for the maximum field size. The angular aperture of the trapezoid θ_t is also represented. The figure is not to scale.

$$\theta_t = \arctan\left(\frac{0.5(W_{leaf-d} - W_{leaf-u})}{h_{leaf}}\right) \tag{3.3}$$

To determine the leaf rotation, a fundamental unit u is defined as the half of the leaf gap plus half of the leaf solid. This is to assure that the gaps between adjacent leaves are constant throughout the entire leaf bank. To illustrate the concept of u, Fig. 3.12 depicts three leaves rotated by their respective angles θ_1, θ_2 and θ_3 with respect to CAX. The six units u corresponding to the three leaves are represented by the pink and purple shadowed regions.

θ_{max} is the maximum angle in this example, defined by the outer side of the last leaf, taking into consideration the angle of the trapezoid θ_t (Eq. 3.4). However, θ_t is approximately 0.0023 radians (Eq. 3.3) and can be considered zero ($\theta_t \approx 0$), which leads to $\theta_{max} \approx \theta_3$. It is important to mention that only the θ_t of the last leaves were disregarded, since for all the other leaves θ_t is included in the units u and therefore considered in the angular rotation and subsequent Z_{Pos} positioning of the leaves.

$$\theta_{max} = \theta_3 + \theta_t \approx \theta_3 \tag{3.4}$$

Due to the leaves' geometry and the constant gap between leaves, it can be assumed that each unit u has its own rotation with respect to CAX, and the angular increment between adjacent u is constant. Hence, the angle of the last leaf θ_3 can be equally divided by the total number of units u which it encompasses. As seen in Fig. 3.12, θ_3 does not

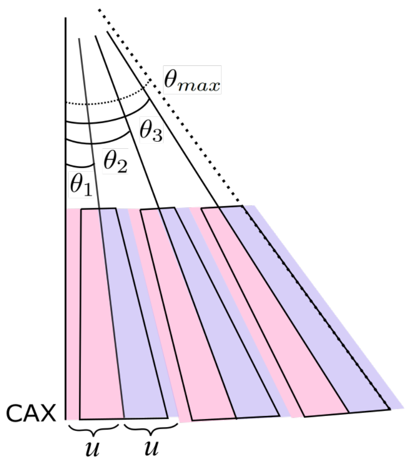

Figure 3.12: Schematic representation of three adjacent leaves rotated by θ_1, θ_2 and θ_3 with respect to the CAX. For each leaf, two fundamental units u , depicted in pink and purple, can be defined as half leaf gap plus half leaf solid. u is useful for the determination of the leaf's rotation. Figure is not to scale.

take into consideration the last unit u, encompassing five of the six u depicted. Therefore, the increment on the angular rotation of each unit u from the previous u, with respect to the CAX, is defined by γ in Eq. 3.5:

$$\gamma = \left(\frac{\theta_3}{5}\right) = \left(\frac{\theta_{max}}{5}\right) \tag{3.5}$$

Still referring to Fig. 3.12, the angular rotation of the first leaf θ_1 encompasses only one unit u, corresponding to an angle γ with respect to CAX (Eq. 3.6):

$$\theta_1 = \gamma \tag{3.6}$$

Between the centers of the first and the second leaves there are two units u. Hence, the angle of the second leave θ_2 is given by Eq.3.7:

$$\theta_2 = \theta_1 + 2\gamma \tag{3.7}$$

Similarly for the third leaf, θ_3 is the angle of the previous leaf, incremented by two units u (Eq. 3.8). Therefore, it can be inferred that the increment between the angular rotation of adjacent leaves is 2γ.

$$\theta_3 = \theta_2 + 2\gamma \tag{3.8}$$

To generalize this scenario to all the 80 leaves present in each leaf bank, it is important to describe how the leaves are numbered in the Geant4 model. The leaves are positioned symmetrically along the Z axis, i.e. 40 leaves at the positive side and 40 leaves at the negative side of the origin. In Geant4, the leaves are numbered from 0 to 79, as illustrated in blue in Fig. 3.13, with leaves 0 to 39 placed at the negative side of Z and leaves 40 to 79 at the positive size of Z.

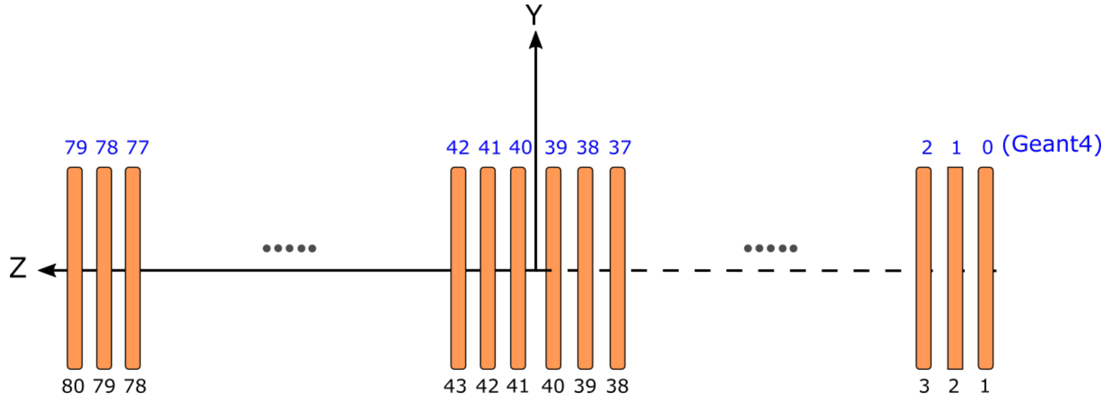

Figure 3.13: Simplified representation of the leaves' positioning along the Z axis to illustrate the numbering convention used in the Geant4 model. Leaves are not rotated and are depicted as rectangles instead of trapezoids; the tongue and groove are also not depicted. The figure is not to scale.

To define the angular increment between adjacent leaves, only half of the leaves need be considered, e.g. leaves 0 to 39. This assemble is composed by 80 units u. However, as shown in the example with three leaves, the last unit u can be ignored. Making an analogy to the example with three leaves, γ can be defined by Eq. 3.9, with θ_{max} replaced by θ_{Fmax}. θ_{Fmax} should also include the trapezoid angle θ_t. Again as in the example with three leaves, θ_t can be considered zero and θ_{Fmax} can be defined as in Eq. 3.2.

$$\gamma = \left(\frac{\theta_{Fmax}}{79}\right) \tag{3.9}$$

The outermost leaf in the considered group, in this case leaf 0, comprises 79 units u. Therefore its angular rotation θ_0 is simply given by Eq. 3.10:

$$\theta_0 = 79\gamma \tag{3.10}$$

The leaf 39, which is the first leaf next to the CAX, comprises only one unit u (similar to leaf 1 in Fig. 3.12), and its angular rotation θ_{39} is given by Eq. 3.11:

$$\theta_{39} = \gamma \tag{3.11}$$

The angular rotation θ_i of each leaf i, with $i = [0, 79]$, can be generalized by a linear equation (Eq. 3.12), where a and b are constant factors. By combining Eqs. 3.10, 3.11 and 3.12, the determination of a and b is straightforward, leading to the final analytic equation for determination of the angular rotation of the leaves inside one leaf bank (Eq. 3.13).

$$\theta_i = a + b \cdot i \tag{3.12}$$

$$\theta_i = 79\gamma - 2\gamma \cdot i \tag{3.13}$$

The leaves in the opposing bank are numbered from 80 to 159. Leaves 80 to 119 are positioned at the negative side of the Z axis while leaves 120 to 159 are at the positive side. Their angular rotations are mirrored to the angular rotation of leaves 0 to 79, and the same considerations presented in this method apply. Hence, the angular rotation for leaves 80 to 119 are simply given by Eq.3.14, with $i = [0, 79]$.

$$\theta_{i+80} = \theta_i \tag{3.14}$$

3.1.3.4 Leaves: Z position

Similar to the methodology used for determining the leaves' rotation, to calculate the leaves' position along the Z axis (Z_{Pos}) the gap g between adjacent leaves needs to be incorporated as part of the leaf's geometry. Additionally, the leaf width at the center of the leaf W_{leaf-c} plays an important role in the analytic expressions for Z_{Pos} and needs to be calculated. Using the variables already defined in Table 3.2 and a half of the trapezoid (as depicted in Fig. 3.14), the determination of W_{leaf-c} is straightforward and can be defined by Eq.3.15.

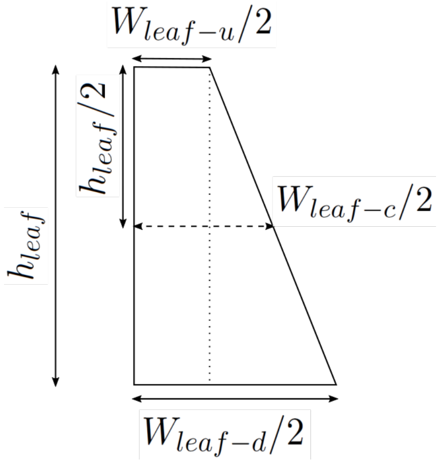

Figure 3.14: Half leaf representation, used to determine the leaf's central width W_{leaf-c} value, required for the calculation of the leaf's position along the Z axis Z_{Pos}. The description and respective values of W_{leaf-u}, W_{leaf-d} and h_{leaf} used in the MC model are listed in Table 3.2. The figure is not to scale.

$$W_{leaf-c} = \frac{W_{leaf-u} + W_{leaf-d}}{2} \tag{3.15}$$

Considering the leaf as an "extended" trapezoid with an extra width of $g/2$ on each side, three new variables are introduced: d_1, d_2 and d_3, corresponding to the trapezoid's width in the upper, central and bottom parts, respectively (Eqs. 3.16, 3.17 and 3.18), as illustrated in Fig. 3.15a.

$$d_1 = W_{leaf-u} + g \tag{3.16}$$

$$d_2 = W_{leaf-c} + g \tag{3.17}$$

$$d_3 = W_{leaf-d} + g \tag{3.18}$$

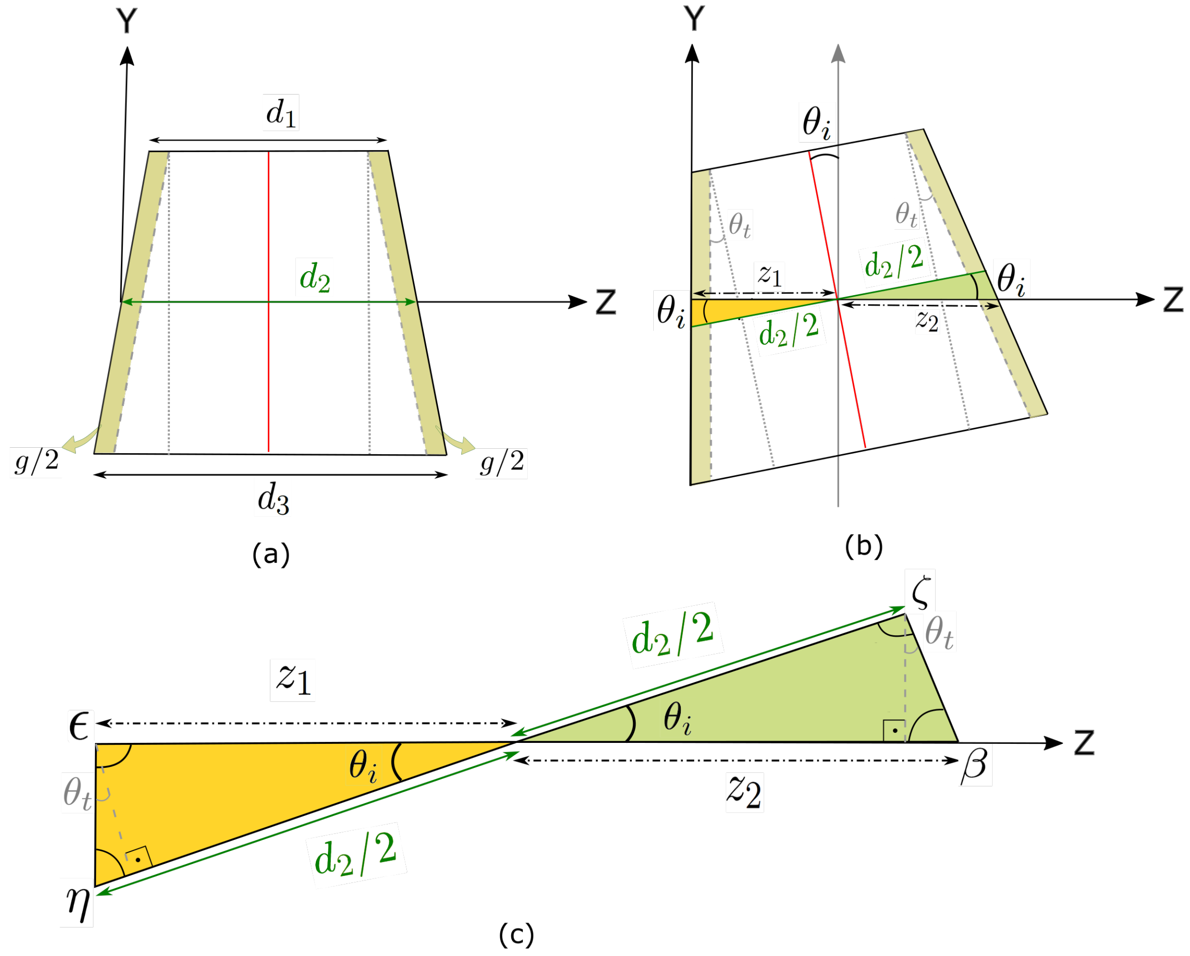

Figure 3.15: (a) Extended leaf widths in the upper, central and bottom parts, d_1, d_2 and d_3, respectively. (b) Leaf rotated by θ_i, and the respective Z position z_1 defined as the projection of $d_2/2$ along the Z axis. (c) The two highlighted triangles (orange and green in detail), together with the angles ϵ, η, β and ζ, are needed for calculation of z_1 and z_2, as described in the text. Figures are not to scale.

In Geant4, the position of an object is given relative to its central point. For one leaf, this central position corresponds to the intersection of the green and red lines depicted in Fig. 3.15a, with the green line and the Z axis superimposed. The position of the leaf in the Z axis is defined simply by $d_2/2$, considering the origin of Z at the intersection of the Y and Z axis. When the leaf is rotated by an angle θ_i, the leaf's position on Z is then defined by the projection of the first half of d_2 over the Z axis, depicted as z_1 in Fig. 3.15b. The variable z_2 is the projection of the other half of d_2 over the Z axis and will be useful to determine the position of the subsequent leaves.

The values of z_1 and z_2 can be calculated based on the two opposing triangles highlighted in green and orange in Fig. 3.15b. Fig. 3.15c shows in detail the geometry which allows the definition of the following trigonometric relationships given by Eqs. 3.19, 3.20, 3.21 and 3.22:

$$\epsilon = 180^\circ - \theta_i - \eta \tag{3.19}$$

$$\eta = 90^\circ - \theta_t \tag{3.20}$$

$$\beta = 180^\circ - \theta_i - \zeta \tag{3.21}$$

$$\zeta = 90^\circ + \theta_t \tag{3.22}$$

Replacing Eq. 3.20 in Eq. 3.19 and Eq. 3.22 in Eq. 3.21 results in:

$$\epsilon = 180^\circ - \theta_i - (90^\circ - \theta_t) = 90^\circ - (\theta_i - \theta_t) \tag{3.23}$$

$$\beta = 180^\circ - \theta_i - (90^\circ + \theta_t) = 90^\circ - (\theta_i + \theta_t) \tag{3.24}$$

Using the law of sines on the orange triangle results in Eq. 3.25 :

$$\frac{z_1}{\sin\eta} = \frac{d_2/2}{\sin\epsilon} \tag{3.25}$$

Replacing Eqs. 3.20 and 3.23 in 3.25 and isolating z_1 leads to:

$$z_1 = \frac{d_2}{2}\left(\frac{\sin(90^\circ - \theta_t)}{\sin(90^\circ - (\theta_i - \theta_t))}\right) \tag{3.26}$$

Using the relationship between sine and cosine of complementary angles (Eq. 3.27)

$$\begin{aligned} \sin\kappa &= cos(90^\circ - \kappa), \qquad & \sin(\kappa + 90^\circ) &= cos\kappa, \\ \cos\kappa &= sin(90^\circ - \kappa), \qquad & \cos(\kappa + 90^\circ) &= -sin\kappa, \end{aligned} \tag{3.27}$$

and applying Eq. 3.27 to Eq. 3.26, the value of z_1 is finally given by Eq. 3.28:

$$z_1 = \frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_i - \theta_t)}\right) \tag{3.28}$$

In the same way for the green triangle, applying the law of sines results in Eq. 3.29

$$\frac{z_2}{\sin\zeta} = \frac{d_2/2}{\sin\beta} \tag{3.29}$$

Replacing Eqs.3.22 and 3.24 in 3.29 and isolating z_2 gives Eq. 3.30

$$z_2 = \frac{d_2}{2}\left(\frac{\sin(90^\circ + \theta_t)}{\sin(90^\circ - (\theta_i + \theta_t))}\right) \tag{3.30}$$

Finally, with the relationships in Eqs. 3.27, z_2 is defined by Eq. 3.31

$$z_2 = \frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_i + \theta_t)}\right) \tag{3.31}$$

z_1 and z_2 are the two quantities used to define the position of the leaves side by side over the Z axis: they are the projections of the first and second half of the extended leaf's horizontal midline (d_2) into the Z axis, respectively. Moreover, as z_1 and z_2 are functions of the leaf's rotation θ_i, each leaf has its own specific values of z_1 and z_2, which can hence be written as $z_1(i)$ and $z_2(i)$, with $i = [0, 79]$. For the central leaves, i.e. those closest to the CAX (leaves 39 and 40 from Fig. 3.13), only z_1 is needed. To demonstrate it, let's focus on leaves $i = [0, 39]$ (leaves on the negative side of the Z axis), more specifically on leaves 39 and 38, as illustrated by Fig.3.16. For the purpose of this demonstration, the absolute values of $Z_{Pos}(i)$ will be considered, but the correct signal will be properly accounted for at the final equations. The position of the central leaf 39 $Z_{Pos}(39)$ is given by Eq.3.32:

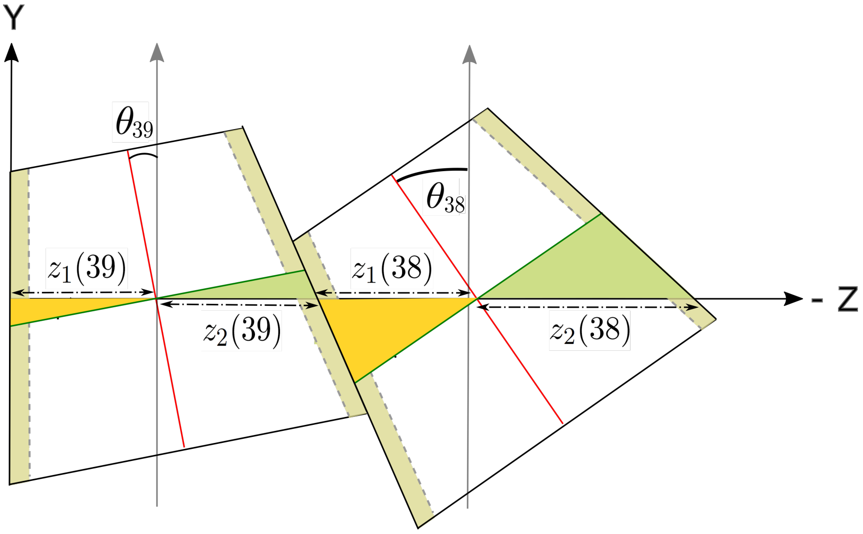

Figure 3.16: Simplified representation of leaves 39 and 38, the two leaves closest to the CAX on the negative side of the Z axis, with respective rotation angles θ_i, z_1 and z_2. The figure is not to scale.

$$Z_{central} = Z_{Pos}(39) = z_1(39) = \frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_{39} - \theta_t)}\right) \tag{3.32}$$

The position of leaf 38, $Z_{Pos}(38)$, will be the position of the former leaf 39, plus the distances defined by $z_2(39)$ and $z_1(38)$:

$$\begin{aligned} Z_{Pos}(38) &= Z_{central} + z_2(39) + z_1(38) \\ &= Z_{central} + \frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_{39} + \theta_t)} + \frac{\cos\theta_t}{\cos(\theta_{38} - \theta_t)}\right) \\ &= Z_{central} + \frac{d_2}{2}\cos\theta_t\left(\frac{1}{\cos(\theta_{39} + \theta_t)} + \frac{1}{\cos(\theta_{38} - \theta_t)}\right) \end{aligned} \tag{3.33}$$

Similarly, the position of leaf 37 will be the position of leaf 38 plus $z_2(38)$ and $z_1(37)$:

$$\begin{aligned}
Z_{Pos}(37) &= Z_{central} + Z_{Pos}(38) + z_2(38) + z_1(37) \\
&= Z_{central} + \frac{d_2}{2}\cos\theta_t \left(\frac{1}{\cos(\theta_{39}+\theta_t)} + \frac{1}{\cos(\theta_{38}-\theta_t)} + \frac{1}{\cos(\theta_{38}+\theta_t)} + \frac{1}{\cos(\theta_{37}-\theta_t)} \right) \\
&= Z_{central} + \frac{d_2}{2}\cos\theta_t \left(\frac{1}{\cos(\theta_{39}+\theta_t)} + \frac{1}{\cos(\theta_{38}+\theta_t)} + \frac{1}{\cos(\theta_{38}-\theta_t)} + \frac{1}{\cos(\theta_{37}-\theta_t)} \right)
\end{aligned} \tag{3.34}$$

This process goes on iteratively until the last leaf $i = 0$. For leaves $i = [0, 39]$, $Z_{Pos}(i) \leq 0$. From Eqs. 3.33 and 3.34, a generalized formula for the positions of the subsequent leaves can be inferred:

$$Z_{Pos}(i) = -\left[Z_{central} + \frac{d_2}{2}\cos\theta_t \sum_{j=i}^{38} \left(\frac{1}{\cos(\theta_{j+1}+\theta_t)} + \frac{1}{\cos(\theta_j-\theta_t)} \right) \right], \text{ for } \quad 0 \leq i \leq 38 \tag{3.35}$$

with $Z_{central}$ given by Eq. 3.32. For leaf 39 ($i = 39$) the Z position is:

$$Z_{Pos}(39) = -Z_{central} = -\frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_{39}-\theta_t)} \right) \tag{3.36}$$

Leaves $i = [40, 79]$ are positioned symmetrically to leaves $i = [0, 39]$ with respect to the CAX, over the positive side of the Z axis, such that $Z_{Pos}(i) \geq 0$. Hence, their positions can be simply calculated by Eq. 3.37. The position of leaf 40 (the first one on the positive side) is given by Eq. 3.38.

$$Z_{Pos}(i) = Z_{central} + \frac{d_2}{2}\cos\theta_t \sum_{j=41}^{i} \left(\frac{1}{\cos(\theta_{j-1}+\theta_t)} + \frac{1}{\cos(\theta_j-\theta_t)} \right), \text{ for } \quad 41 \leq i \leq 79 \tag{3.37}$$

$$Z_{Pos}(40) = Z_{central} = \frac{d_2}{2}\left(\frac{\cos\theta_t}{\cos(\theta_{40}-\theta_t)} \right) \tag{3.38}$$

3.1.3.5 Leaves: X position

The leaves move along the X direction (crossline, Fig. 3.9) to shape the radiation field, according to positions determined by the treatment planning system (TPS). Leaves $i = [0, 79]$ are located in the negative side of X axis "negative leaf bank" in Fig. 3.9) and can move up to 15 cm across the positive side; leaves $i = [80, 159]$ are located in the positive side of the X axis ("positive leaf ban" in Fig. 3.9), and can move up to 15 cm across the negative side. Similar to Z_{Pos}, the position of leaves over the X axis X_{Pos} is also determined by analytic equations. However, unlike Z_{Pos} which is constant, X_{Pos} is dynamic and will be unique for each control point of a treatment plan, defined by values extracted from TPS dicom planning files (RTPlan.dcm). A MATLAB script (Mathworks Inc., Natick, MA) was used to extract all the necessary information from the RTPlan.dcm files and is presented in Appendix A.

To derive the analytic equation for X_{Pos}, Fig. 3.17 will be used. It illustrates an open leaf positioned at the positive side of the X axis, assuming that the leaf position at isocenter is F_X (in green), as given by the RTPlan.dcm file. F_X is the aperture of the leaf at the isocenter level and defines the field size. F_X corresponds to an aperture p with respect to point P (both in green), with P being the point where the radiation is tangent to the leaf end, beyond which the leaf blocks the radiation beam. Point C is the center of the leaf solid. A is the angular aperture of the radiation field corresponding to F_X (Eq. 3.39), q is the distance from the target to the leaf's center of curvature (Eq. 3.40), s (in blue) is the distance from the center of curvature of the leaf to point P, and l and r (in pink and purple, respectively) are supporting variables that will be used in the following. The remaining variables are defined in Table 3.2. All equations are expressed in millimeters.

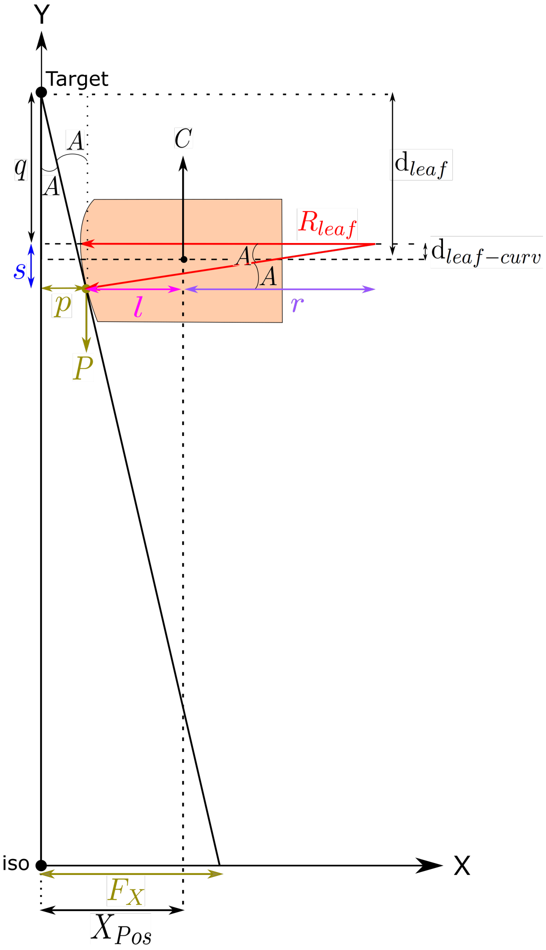

Figure 3.17: Simplified representation of an open leaf, defining a field of size F_X at isocenter level. For determination of X_{Pos}, several parameters are needed, all illustrated in the figure and explained in the text. The figure is not to scale.

$$A = \arctan\left(\frac{F_X}{iso}\right) \tag{3.39}$$

$$q = d_{leaf} - d_{leaf-curve} = 356.8 - 7.5 = 349.3\ mm \tag{3.40}$$

From Fig. 3.17 the X position of the leaf X_{Pos} for an aperture F_X at isocenter is given by Eq. 3.41

$$X_{Pos} = p + l \tag{3.41}$$

Hence, to obtain X_{Pos}, p and l need to be determined. From Fig. 3.17, p is given by Eq. 3.42

$$p = (q + s) \cdot \frac{F_X}{iso} \tag{3.42}$$

and s given by Eq. 3.43

$$s = R_{leaf} \cdot \sin A \tag{3.43}$$

Replacing Eqs. 3.40 and 3.43 in Eq. 3.42, p can be rewritten as:

$$p = (349.3 + R_{leaf} \cdot \sin A) \cdot \frac{F_X}{iso} \tag{3.44}$$

To find l, the sum $l + r$ is considered:

$$l + r = R_{leaf} \cdot \cos A \tag{3.45}$$

From Fig. 3.17, r is half of the leaf length ($L_{leaf}/2$) plus the difference between L_{leaf} and the leaf's curvature radius R_{leaf}. From the values listed in Table 3.2, r can be derived as Eq. 3.46. Replacing r in Eq. 3.45 leads to l being described by Eq. 3.47.

$$r = \frac{L_{leaf}}{2} + (R_{leaf} - L_{leaf}) = 92.5\ mm \tag{3.46}$$

$$l = R_{leaf} \cdot \cos A - r = R_{leaf} \cdot \cos A - 92.5\ mm \tag{3.47}$$

Finally, replacing Eqs. 3.44 and 3.47 in Eq. 3.41, X_{Pos} is given by Eq. 3.48 as

$$X_{Pos} = (349.3 + R_{leaf} \cdot \sin A) \cdot \frac{F_X}{iso} + (R_{leaf} \cdot \cos A - 92.5)\ mm \tag{3.48}$$

with A given by Eq. 3.39, R_{leaf} and iso listed in Table 3.2 and F_X retrieved from the RTPlan.dcm file. Eq. 3.48 is valid for the entire range of movement of the leaves at the positive bank (leaves $i = [80, 159]$), i.e. for open leaves and leaves crossing the CAX. Moreover, each leaf will have a position $F_X(i)$ defined by RTPlan.dcm, such that $X_{Pos}(i)$

is a function of leaf i. Hence, $X_{Pos}(i)$ is described by Eq. 3.49, with $A(i)$ given by Eq. 3.50. All values are given in mm.

$$X_{Pos}(i) = (349.3 + R_{leaf} \cdot \sin A(i)) \cdot \frac{F_X(i)}{iso} + (R_{leaf} \cdot \cos A(i) - 92.5)\ mm, \qquad \text{for} \quad 80 \leq i \leq 159 \tag{3.49}$$

$$A(i) = \arctan\left(\frac{F_X(i)}{iso}\right) \tag{3.50}$$

The $X_{Pos}(i)$ for leaves on the negative bank (leaves $i = [0, 79]$) is derived in a similar way, and is described by Eq. 3.51, with $A(i)$ given by Eq. 3.50. Eq.3.51 is valid for the entire range of movement of the leaves at the negative bank (leaves $i = [80, 159]$). All values are given in mm. Is it important to mention that $F_X(i)$ values are retrieved from the RTPlan.dcm as negative or positive values: negative values correspond to open leaves from the negative bank or closed leaves from the positive bank (leaves crossing the CAX), while positive values correspond to open leaves from the positive bank or closed leaves from the negative bank (leaves crossing the CAX).

$$X_{Pos}(i) = (349.3 - R_{leaf} \cdot \sin A(i)) \cdot \frac{F_X(i)}{iso} - (R_{leaf} \cdot \cos A(i) - 92.5)\ mm, \qquad \text{for} \quad 0 \leq i \leq 79 \tag{3.51}$$

$X_{Pos}(i)$ as defined in Eqs. 3.49 and 3.51 use the light field edge to position the leaves. The light field edge is determined by the tangent of the radiation field with the leaf end, indicated by point P in Fig. 3.17. However, the values of F_X retrieved from the RTPlan.dcm correspond to radiation field edges, which is defined as the position where the intensity of the radiation drops to 50% of the maximum value. Fig. 3.18 illustrates the radiation and light field edges (in red and green, respectively), where the offset Δ between light and radiation field edges can be observed. Therefore, the values of F_X need to be corrected for this offset before being fed into Eqs. 3.49 and 3.51. In this work, values provided by the vendor for TPS calibration were used [136], however calibration methods are available in literature [143, 146, 147]. For leaves $i = [0, 79]$, the offset is $\Delta = 0.26\ mm$; for leaves $i = [80, 159]$, $\Delta = -0.26\ mm$.

Finally, the gap δ between closed opposing leaves needs to be accounted for in $X_{Pos}(i)$. This is achieved by introducing a condition of minimum distance between leaves: if $F_X(i) = F_X(i + 80)$, $X_{Pos}(i)$ and $X_{Pos}(i + 80)$ are incremented by shift of 0.5δ each in the opposing directions, specified at leaf position, not at isocenter.

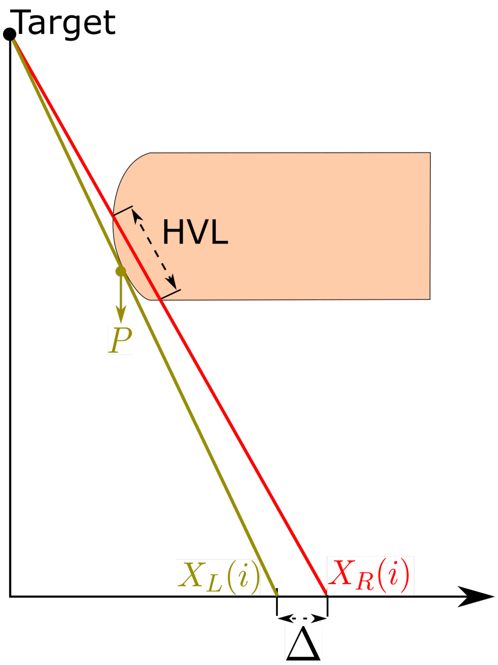

Figure 3.18: Radiation field edge X_R and light field edge X_L, depicted in red and green, respectively. Δ is the offset between X_R and X_L. The figure is not to scale.

3.1.3.6 Diaphragms: X and Z positions

The diaphragms move perpendicular to the leaves' motion, i.e. along the Z axis. While their position in the X axis (X_{PosD}) is constant and centralized inside the JawMother volume, their position along the Z axis (Z_{PosD}) depends on the values retrieved from the RTPlan.dcm file. Hence, for the diaphragms $X_{PosD} = 0$. To determine Z_{PosD}, Fig. 3.19 is used following the same methodology adopted for determining X_{Pos} for the leaves. For an aperture F_Z at isocenter (in green, retrieved from RTPlan.dcm file), p corresponds to the aperture with respect to point P (both in green), where P is the point where the radiation is tangent to the diaphragm end, beyond which the diaphragm blocks the radiation beam. C is the center of the diaphragm solid, A is the angular aperture of the radiation field correspondent to F_Z (Eq. 3.52), q is the distance from the target to the diaphragm's center of curvature (Eq. 3.53), s (in blue) is the distance from the center of curvature of the diaphragms to point P, and l and r (in pink and purple, respectively) are supporting variables that will be used in the following. The remaining variables are defined in Table 3.2.

$$A = \arctan\left(\frac{F_Z}{iso}\right) \tag{3.52}$$

$$q = d_{diap} - d_{diap-curve} = 470.5 - 3.5 = 467\ mm \tag{3.53}$$

From Fig. 3.19, Z_{Posd} is given by Eq. 3.54, with p and s defined as Eqs. 3.55 and 3.56,

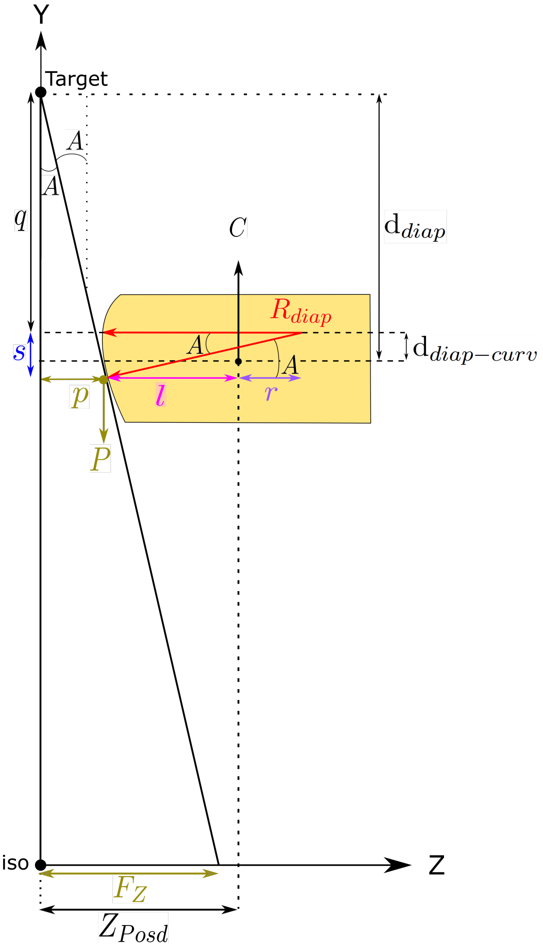

Figure 3.19: Simplified representation of an open diaphragm, defining a field of size F_Z at isocenter level. For determination of Z_{Pos}, several helping parameters are needed, all illustrated in the figure and explained in the text. The figure is not to scale.

respectively.

$$Z_{Posd} = p + l \tag{3.54}$$

$$p = (q + s) \cdot \frac{F_Z}{iso} \tag{3.55}$$

$$s = R_{diap} \cdot \sin A \tag{3.56}$$

From Fig. 3.19, l can be written as Eq. 3.57:

$$l = R_{diap} \cdot \cos A - r \tag{3.57}$$

Using the values listed in Table 3.2, r is calculated with Eq. 3.58:

$$r = \frac{L_{diap}}{2} - (L_{diap} - R_{diap}) = 25 \tag{3.58}$$

Finally, by replacing Eqs. 3.53 and 3.56 in 3.55, Eq. 3.58 in 3.57, and the resultant l and p in Eq. 3.54, Z_{Posd} is defined by Eq. 3.59, with A given in Eq. 3.52.

$$Z_{PosD1} = (467 + R_{diad} \cdot \sin A) \cdot \frac{F_Z}{iso} + (R_{diap} \cdot \cos A - 25) \tag{3.59}$$

Eq. 3.59 describes the position Z for the diaphragm located at the positive side of the Z axis, hereon referred to as D1. The position Z for the other diaphragm, located at the negative side of the Z axis (D2), is obtained following the same method and is given by Eq. 3.60. Eqs. 3.59 and 3.60 are valid for the entire range of the diaphragm. All values are given in mm.

$$Z_{PosD2} = (467 - R_{diad} \cdot \sin A) \cdot \frac{F_Z}{iso} - (R_{diap} \cdot \cos A - 25) \tag{3.60}$$

Z_{PosD1} and Z_{PosD2}, as defined in Eqs. 3.59 and 3.60, use the light field edge to position the diaphragms. Similar to F_X for the leaves' positions, the values of F_Z retrieved from the RTPlan.dcm refer to radiation field edges. Therefore, the values of F_Z also need to be corrected for the offset between light and radiation field edges. Using values provided by the vendor for TPS calibration [136], the offsets were calculated as $\Delta = 0.24\ mm$ for the D1 and $\Delta = -0.24\ mm$ for D2. There is no constraint for closed diaphragms, as they are never closed during irradiation (otherwise the beam is completely blocked).

3.1.3.7 Extra Blocks

The extra blocks were modeled as trapezoids with the same height as the leaves (h_{leaf}) and with thickness correspondent to ten leaves. The length of the block is $2.5 \times L_{leaf}$ such that one block extends along both leaf banks. They are placed along the Z axis inside the MotherMLC next to the first and last leaf of both leaf banks, as seen in Fig. 3.6, following the methodology for Z_{pos}.

3.1.3.8 Water Phantom

The patient-dependent model also included a virtual water phantom, modeled as a water-filled box with 67.5×56×64.5 cm^3, matching the dimension of the physical water phantom used for measurements performed throughout this work (Blue Phantom 2, IBA Dosimetry, Schwarzenbruck, Germany). The virtual water phantom was constructed using an individual function inside the *DicomDetectorConstruction* class and could be activated or deactivated separately, depending on the goal of the simulation performed.

3.1.4 EPID Model

The Geant4 model of the patient-dependent part of the linac includes a model of the EPID panel. The EPID is composed of several different structures, as described in Chapter 2. Despite its complex architecture, for Monte Carlo purposes the EPID can be modeled as a water box of 41×18×41 cm^3. It is placed in the central axis at $Y = -57\ cm$, such that the source-to-surface distance (SSD) is 157 cm. The scoring layer used for simulation is placed inside the EPID box at 3 cm depth, with 160 cm source-to-detector distance (SDD). This 3 cm water thickness above the scoring layer accounts for the build-up effect caused by the copper plate located at the top of the physical panel [32, 81, 148], while the 15 cm water column behind the scoring plane mimics the backscatter observed on the real detector [148].

3.1.5 Patient Model

A virtual model of the patient can be created inside the patient-dependent Geant4 model by importing CT images of patients into the code. The images were imported using the *DicomHandler* class, which reads and stores the information contained in the CT image, such as pixel size, slice thickness, slice position, CT number of each voxel, among others. The CT images are saved in the DICOM format (Digital Imaging and Communications in Medicine, .dcm extension), and one image corresponds to one slice of the patient CT. Once the images were imported, the patient was modeled as a voxelized grid by the *DicomDetectorConstruction* and *DicomPhantomZSlicerHead* classes, taking into account each voxel's positions, size, mass density and material composition. The CT number of each voxel was used to obtain its density and material composition, as described in the following.

To create the patient geometry inside the Geant4 model, the user needs to provide an input file containing relevant information needed by the *DicomHandler* class. This class was created based on the "Medical DICOM" example provided by the Geant4 repository, and to be consistent with the documentation available for this example, this input file will be hereon referred to as Data.dat. The user should include the following information in the following order:

1. Compression value;
2. Number of CT images to be used;
3. Name (with path) of a file containing the shift to isocenter;
4. One line for each CT image used, containing the name (with path) of the file without extension;
5. Number of materials used to construct the patient model;
6. One line for each material containing its name and the upper density boundary.

The compression value determined in the first line defines if the patient will be modeled in Geant4 with the same resolution as in the CT image (compression value 1), or if the resolution will be decreased. The compression value should be given as powers of 2

(2^n). Hence, for a CT image with 512×512 voxels, a compression value of 2 will create a volume with 256×256 voxels, a compression value of 4 will result in a image with 128×128 voxels, etc. The decrease in resolution is given by simply "merging" several voxels into one single voxel. In case compression is used, the CT number of the resultant voxel is the average of the CT number of the individual voxels merged. The second line determines the number of CT images to be used for creating the patient volume inside the Geant4 model.

In this work, the patient geometry was constructed centered at the isocenter of the treatment room volume, such that the patient volume was symmetrically distributed in all directions. In some cases, the isocenter of the CT image is not necessarily defined at position $(X, Y, Z) = (0, 0, 0)$, as is the isocenter of the treatment room. The third line in Data.dat contains a text file (.txt) with the appropriate shift in all directions to correct for this discrepancy, if applicable. The file should have 3 lines, where the shifts in X, Y and Z directions are given in mm in the first, second and third lines, respectively. If no shift is necessary, a text file with three lines (zero at each line) should be provided anyway, otherwise the simulation will crash. This file was not originally required in the DICOM example and is an adaptation to this work.

Next, one line for each CT DICOM image to be reconstructed is needed, containing the name (and path) of the file without the .dcm extension. The number of lines and files should match the value defined in item 2 above. Each CT image used corresponds to one patient slice along the Z axis (axial plane, Fig. 3.9). The slices will be reconstructed separately, but will be merged at runtime to form a unique patient volume. Therefore, the name of the slices should be listed according to their Z position in a continuous way, with decreasing value of Z.

The subsequent line in Data.dat should refer to the number of different materials which can be used to reconstruct the patient voxels. Finally, one line for each material used should be given, with the material's name and its density value, as created in the *DicomDetectorConstruction* class. The number of materials listed should match the value defined in item 5, and should be listed in increasing order of density. These values given for density in Data.dat are considered the upper density boundary for the respective material: the voxels with density between 0 and the first value will be constructed with the first material; the voxels with density between the first and the second values will be assigned to the second material, etc.

After reading the Data.dat files and determining which CT images will be reconstructed, the Geant4 code needs to convert the CT number of the voxels of every image to density values, and thereafter determine which material will be used to construct them. The *DicomHandler* reads the DICOM files, stores relevant information and defines the densities of the voxels. To do so, an extra file named "CT2Density.dat" is used, set as an environmental variable. This file contains the calibration points to convert CT numbers into density values, and was retrieved from the commissioned TPS used with the Elekta

Synergy® in Klinikum Grosshadern. For voxels with CT number between two points, the voxel density is obtained by linear interpolation. This calibration is center-specific and is usually determined using a CT scan from a characterization phantom with tissue like-inserts, where the material composition and density of each insert is known. For the TPS used with the Elekta Synergy® at Klinikum Grosshadern, the calibration points are presented in Table 3.3. The determination of the values in Table 3.3 is described in detail in the master thesis of S. Schmid and A. Resch [149, 150].

Table 3.3: Values of CT number and density, used for calibration. CT values lower than -1000 and bigger than 4000 are assigned to the lowest and highest density values, respectively (0.001 g/cm^3 and 3.708 g/cm^3) [149, 150].

CT number (HU)	Density (g/cm^3)
-1000	0.001
-711	0.290
-555	0.440
-104	0.943
-53	0.980
-8	1000
18	1.051
63	1.094
209	1.141
218	1.149
451	1.332
821	1.559
1237	1.822
4000	3.708

The material composition of the voxels was determined based on their densities: the density range considered in this work (Table 3.3) was divided into 86 values, and each value was correlated to a different material. Therefore, 86 different materials were created and used to reconstruct the patient voxels inside the Geant4 model. The 86 new materials were obtained with a method that correlates CT numbers, given in Hounsfield Unit (HU), to material composition and density (g/cm^3) [151, 152], described in details in the master thesis of S. Schmid and A. Resch [149, 150].

To summarize: the CT image files listed in Data.dat are read by *DicomHandler*, which converts the CT values of the voxels into density based on the conversion points given in “CT2Density.dat”. Once the density of the voxel is determined, the voxel is assigned to one of the 86 materials listed in Data.dat (as previously described), and a patient voxel is created by *DicomDetectorConstruction* inside the Geant4 model.

3.1.6 User-defined commands

For different simulations with different geometric configurations, the parameters can be changed inside the Geant4 classes. This requires however that the user compiles the code for each different configuration, creating an executable file for this specific scenario, preventing parallel simulation with different parameters, as for example simulation of different fields at different gantry angles.

One alternative is to adopt user-defined commands, which allows setting the values of some variables via macro files (.mac). Macro files can be given as input to the simulation: the executable reads the value of the variable given in the macro file, and assigns this value to the proper parameter defined inside the Geant4 classes. Different simulations can be run in parallel by simply changing the parameters given in the macro files. Geant4 provides a set of built-in user commands to control the number of particles to be simulated, verbose levels, several visualization parameters, among others. A description of the Geant4 built-in command can be found in the Geant4 documentation [153].

For the electron beam optimization, the build-in user-defined commands from the GPS (general particle source) were used to set the parameters for the electron beam hitting the target. For generation of PhSp files, a new user-defined command was created at *DetectorLinac* class to define the name of the PhSp binary file. More details will be given on subsection 3.1.7.

For the patient-dependent model, user-defined commands were created inside the *DicomLinac* class to accommodate specific needs, as listed below:

1. /define/leafParameters: sets the density of the tungsten alloy (ρ, in g/cm^3) and height (h_{leaf}, in mm), both tunable leaf parameters. The parameters should be specified in this order. They are passed to the *DicomDetectorConstruction* class and used to construct the leaves, as well as the diaphragms and the extra blocks (all these components are modeled with the same density);
2. /define/virtualFocus: sets the shift to virtual focus α of the leaf bank, also a tunable parameter. This parameter is passed to the *DicomDetectorConstruction* class and used to rotate the “MLCMother” volume around the X axis towards the virtual focus (Fig. 3.6);
3. /define/leafGap: defines the gap g between adjacent leaves;
4. /constructor/Parameters: sets the parameters of the field to be simulated, as extracted from the RTPlan.dcm files (Appendix A), in the following order: gantry rotation (in degrees), a text file containing the position of each leaf (at isocenter, in mm - $F_X(i)$ in Eqs. 3.49 and 3.51), and a text file containing the position of each diaphragm (at isocenter, in mm - F_Z in Eqs. 3.59 and 3.60). The *DicomDetectorLinac* class takes these input parameters and positions the leaves and diaphragms inside their respective mother volumes accordingly;

5. /PSF/name: sets the name of the PhSp binary file to be used as primary particle source.

The components of the patient-dependent model could be rotated around the Z axis (i.e. around the patient model) to reproduce irradiation from different gantry angles. A clockwise rotation direction was assumed, in agreement with the linac used. The value of the gantry angle, set with the /constructor/Parameters in the macro file, is used to rotate the "MLCMother" and the "JawMother" volumes around the Z axis. All PhSp files were created using the patient-independent model with fixed gantry position at 0 degrees. To accommodate gantry rotation for particle generation inside the patient-dependent model, the value retrieved for the gantry rotation angle is also passed to the *DicomPrimaryGeneratorAction* such that the initial particles emerging from the PhSp file were rotated accordingly. The patient-independent model was used only for PhSp generation and did not include gantry rotation.

Moreover, the dose scoring throughout this work was performed using command-based scoring. This Geant4 feature allows the user to define the characteristics of the scoring volumes inside macro files, such as the dimension and position of the scoring volume and the quantity to be scored.

3.1.7 Incident electron beam optimization

The energy spectrum of the electron beam, which exits the vacuum window and hits the target producing the photon beam of the linac, strongly influences the dosimetric output of the equipment. The shape and size of the focal spot, i.e. the spatial distribution of the electron beam incident on the target, also affects the output of the linac. Consequently, the quality of dose simulations from a linac MC model strongly depends on the proper modeling of the initial electron beam, whose characteristics can vary from the nominal values provided by the vendors, being machine-specific [125, 131, 132, 144, 154–157]. Therefore, the initial electron beam properties needed to be optimized in order to match the equipment used and was a crucial step in the development of this work. Several methods for optimization of the initial electron beam parameters are proposed in the literature [125, 126, 132, 144, 154, 157–159]. The majority of them iteratively varies the electron beam parameters, namely its mean energy ($\overline{E}$) and energy spread (σ_E), its spatial distribution and its angular spread, to find the best set of values that results in a good match between simulations and measurements.

Sheikh-Bagheri et al. [154] investigated the influence of the electron beam energy and spatial distribution on the off-axis factor and percentage depth dose (PDD) profile for a 10×10 cm^2 field. They concluded that the off-axis factor is strongly influenced by the mean energy $\overline{E}$ and spatial distribution of the electron beam, and insensitive to the full width of half maximum of its energy distribution ($FWHM_E$) and its angular spread. The PPD, on the other hand, is insensitive to the spatial distribution of the beam, but affected by $\overline{E}$ and $FWHM_E$.

Aljarrah et al. [125] used PDDs and lateral dose profiles at different depths in water from 4×4 cm^2, 10×10 cm^2 and 20× 20cm^2 fields to evaluate their sensitivity to the beam $\overline{E}$ and spatial distribution. They also investigated different metrics for a quantitative evaluation of the agreement between simulations and measurements. Moreover, they proposed a minimum required dataset for the electron beam optimization, with the intention to prevent unnecessary simulations. They concluded that the PDDs are not sensitive to the evaluated parameters, and therefore the minimum set for beam optimization comprises lateral dose profiles for several field sizes at several depths in water, along with only one PDD profile for a large field. Finally, a definitive conclusion for the best metric providing a quantitative analysis of the agreement between simulations and measurements could not be achieved. The proposed metrics might lead to several optimal parameters combinations, but the ultimate decision on the best set relies on visual inspection of the dose distributions.

In the work of J. Fleckenstein [132], it is stated that the flattening filter is strongly responsible for shaping the photon beam of the linac, and therefore the electron beam parameters should be optimized using quadratic fields of at least 30×30 cm^2, which entirely cover the flattening filter in the lateral dimension.

Chibani et al. [159] investigated the influence of $\overline{E}$, $FWHM_E$, radial distribution and angular divergence of the primary electrons on the PDDs and lateral profiles of fields ranging from 2×2 cm^2 to 35×35 cm^2, showing that very small and very large fields are more sensitive to the $\overline{E}$, $FWHM_E$ and angular distribution than the 10×10 cm^2 field. Therefore, the optimization of the initial electron beam parameters should include a broad range of field sizes.

Despite the several approaches proposed, modeling the electron beam hitting the target remains a challenge when using MC methods to simulate dose distributions in radiotherapy, and a consensus on the best methodology is yet missing. For a deeper overview of other methods, please refer to Chetty et al. [160] and Ma et al. [161].

As the ultimate goal of the MC model in this work is to simulate dose distributions for clinical IMRT treatment plans, which comprise several fields in a broad range of sizes, it was important to assure that the model properly describes smaller and larger fields. Therefore, the optimization process was based on PDDs and lateral dose profiles at two different depths in water (15 mm and 100 mm) for fields ranging from 2×2 cm^2 to 30×30 cm^2. The optimized parameters were the mean electron beam energy $\overline{E}$, the energy spread around its mean value σ_E, the beam's spatial distribution and its angular spread.

Measurements performed at the Klinikum Grosshadern with the Elekta Synergy® machine coupled with the Agility$^{\text{TM}}$ MLC for the 6 MV photon beam were used as benchmark. The PDDs and lateral dose profiles, both in inline and crossline directions, were measured using a micro diamond detector (microDiamond 60019, PTW, Freiburg, Germany) in a water phantom (Blue Phantom 2, IBA Dosimetry, Schwarzenbruck, Germany) at source-to-surface distance (SSD) of 90 cm. The micro diamond detector has a sensitive volume of 0.004 mm^3, offers sub-millimetric spatial resolution and is indicated for fields sizes ranging from 1×1 cm^2 to 40×40 cm^2 [162]. PDD profiles were measured from the surface of the water to 300 mm depth, with 1-mm steps in the buildup region, 2-mm up to the 50th measuring point and 4-mm for the rest of the measurement. Inline and crossline profiles were measured at 15 mm and 100 mm depth, with 1-mm steps in the penumbra region and 4-mm elsewhere,

It is important to mention that the electron beam optimization and the leaf optimization were performed separately. For the electron beam optimization, the patient-independent model (3.1.2) was used to generate PhSp files. For each configuration of the initial electron beam, i.e. for each combination of initial parameters, one PhSp was generated. This PhSp was then used as primary beam generator on the patient-dependent model (3.1.3) to simulate the PDD and lateral dose distributions on a virtual water phantom, for the fields considered and under the same conditions of the benchmark measurements. For both PDD and lateral profiles, the dose was recorded using the command-based scoring feature from Geant4 [163], hereon referred to as scoring meshes. The simulations were compared to the reference measurements and this process was repeated iteratively until the best set of electron beam parameters was found, based on the visual inspection of the distributions. As the electron optimization process involves open squared fields, the PDDs and lateral profiles in this case are less sensitive to the tunable leaf parameters (leaf height, density of the tungsten alloy and shift to virtual focus), which affect mainly the beam transmission through the leaves and diaphragms. Hence, the MLC model could be constructed with the values listed in Table 3.2. Once the optimal parameters of the incident electron beam were determined, the correspondent PhSp files were used as initial particle generator on the patient-dependent model and the MLC was optimized. More details are given in section 3.1.9.

First, the energy spectrum of the incident electron beam was manipulated, while the angular spread Φ was kept constant using the nominal values provided by the vendor [132]. The initial source was constructed as a point source with a Gaussian energy distribution, positioned at 3 cm from the exit of the vacuum window. The mean energy $\overline{E}$ was varied from 6.0 to 7.0 MeV in steps of 0.2 MeV, each time combined with $FWHM_E$ values of 0.01, 0.5 or 1.0 MeV (Table 3.4). The electron source was created using the *G4GeneralParticleSouce* class, which does not accept the spread of the energy distribution in terms of $FWHM$. The respective values of the standard deviation σ_E were calculated based on the relationship between $FWHM$ and σ of a Gaussian distribution (Eq. 3.61) and fed into the particle generator. As PDD profiles in the central axis are primarily determined by photons incident on the phantom around the central axis itself, they are less sensitive to variations of the spatial distribution of the electron beam [154], and are suited for the energy optimization step. For each set of $\overline{E}$ and σ_E, a PhSp file was created. This PhSp was used as primary beam generator on the patient-dependent model to simulate the PDD distributions on a virtual water phantom for 2×2 cm^2, 5×5 cm^2, 10×10 cm^2 and 30×30 cm^2 fields. The dose was scored with a 1-mm resolution along the central axis, using column-like scoring meshes. For both simulated and measured PDDs, the dose distribution was normalized to the dose scored at the maximum dose depth (d_{Dmax}). Based on visual inspection, optimal values for $\overline{E}$ and σ_E were selected and used for subsequent simulations.

$$FWHM = 2\sigma\sqrt{2ln(2)} \tag{3.61}$$

The optimization of the focal spot and the angular spread of the beam was performed using measured lateral profiles as benchmark. The lateral profiles are highly sensitive to changes in the spatial distribution of the focal spot, particularly at the penumbra region [126, 156, 158, 164]. The focal spot depends on the type of source used, on the size of the source and on the angular spread of the electrons emerging from it. J. Fleckenstein [132] has modeled an Elekta Synergy® using a point source, while other authors have proposed circular [165] or elliptical focal spots [126, 166] as a better match. Therefore the three types of source were investigated. The source size was changed by changing the angular spread, if point, the angular spread and radial distribution, if circular, and by changing the angular spread and spatial distribution along the minor and major axis, if elliptical.

For circular sources, the radial distribution is defined by the radius of the beam and its standard deviation σ_r. The radius was kept constant as 0.05 mm, while σ_r was set to 0.5 mm or 1.0 mm. Similarly, the elliptical sources were modeled with a fixed value of 0.1 mm for both minor and major axis. The elliptical characteristic of the source was achieved by changing the standard deviation of the beam in the inline and crossline directions. For the inline and crossline directions, the values investigated were $\sigma_{in} = [0.0001, 0.3, 0.5]$ mm and $\sigma_{cr} = [0.0001, 0.8, 1.0]$ mm, respectively (Table 3.5).

Table 3.4: Mean energy $\overline{E}$ and $FWHM_E$ values used for the energy optimization of the initial electron beam. The value for angular spread (radial) was kept constant at the nominal value provided by the vendor [132].

$\overline{E}$ (MeV)	$FWHM_E$ (MeV)	Φ_e (rad)
6.0	0.01 0.5 1.0	0.02
6.2	0.01 0.5 1.0	0.02
6.4	0.01 0.5 1.0	0.02
6.6	0.01 0.5 1.0	0.02
6.8	0.01 0.5 1.0	0.02
7.0	0.01 0.5 1.0	0.02

Table 3.5: Source type, size and standard deviation values of the beam's spatial distribution, used for focal spot optimization of the initial electron beam.

Source type	Radius (mm)	σ_r (mm)	Minor/major axis (mm)	σ_{in} (mm)	σ_{cr} (mm
Point	-	-	-	-	
Circular	0.5	0.0 0.5 1.0	-	-	-
Elliptical	-	-	1.0	0.001 0.3 0.5	0.001 0.8 1.0

For a point source, the angular distribution is radially symmetric (Φ_r). For circular and elliptical sources, the angular spread can be set individually for each direction. Different values for the angular spread in the inline and crossline directions were evaluated, Φ_{in} and Φ_{cr}, respectively (Table 3.6).

Table 3.6: Source type and respective angular distribution values, used for focal spot optimization of the initial electron beam.

Source type	Φ_r (rad)	Φ_{in} (rad)	Φ_{cr} (rad)
Point	0.020 0.022, 0.025 0.035	0.011, 0.017	0.011, 0.014, 0.022
Circular Elliptical	-	0.011, 0.020 0.021, 0.022 0.035	0.011, 0.014, 0.017 0.020, 0.022, 0.023 0.024, 0.025, 0.035

Similar to the energy optimization, one PhSp was generated for each different parameter set, using the patient-independent model. The PhSps were given as initial particle generator to the patient-dependent model to simulate the lateral profiles inside a virtual water phantom, using scoring meshes with 1 mm resolution. The values of each parameter investigated were based on results published by others [126, 132, 165, 166] and by the iterative comparison of the simulated profiles to the benchmark measurements. Both measured and simulated profiles were normalized to their respective dose at CAX.

The optimal parameters for the focal spot were selected based on visual comparison of the simulated lateral profiles and benchmark measurements. It is important to mention that the energy distribution strongly influences the lateral profiles of larger fields [160]. Therefore, after the focal spot optimization, it might be necessary to reevaluate the values for the electron beam energy before a final set of initial beam parameters can be determined.

Finally, once a set of optimal parameters was found, 500 unique PhSp files were generated, each using 10^8 statistically independent initial electron histories, yielding a total of 5×10^{10} initial electron histories that could be used for further simulations.

3.1.8 Phase Space generation and application

As described in 3.1.2, the PhSp plane was created in the patient-independent model as a vacuum-filled solid with 5 μm thickness, placed 16 cm away from the target (i.e. 84 cm away from the isocenter). For creation of the PhSp files, the PhSp solid was set as a Sensitive Detector (SD) [163]. Particles entering the upper boundary of this SD (the boundary closer to the target) were detected by the $DicomSourceSurfaceSD$ class, which uses a filter to select only particles entering the SD from above. The *DicomSourcePhspData* class creates a binary file (.bin), named with the string passed to the /PSF/name command in the macro file (3.1.6), and stores seven parameters in the following order:

1. Particle type, following the PDG encoding (22 corresponds to photons, ±11 for

electrons and positrons, respectively [163, 167]);

2. Particle energy (in MeV);
3. Position along the X axis (in mm), where the particle enters the PhSp solid;
4. Position along the Z axis (in mm), where the particle enters the PhSp solid;
5. Momentum direction along the X axis, in terms of direction cosine (P_x);
6. Momentum direction along the Y axis, in terms of direction cosine (P_y);
7. Momentum direction along the Z axis, in terms of direction cosine (P_z).

All variables were set as $G4float$, corresponding to 4 bytes each. Therefore, 28 bytes were allocated for each particle stored in the .bin file. For the PhSp generated with the optimal electron beam parameters (3.1.7), with 10^8 initial electron histories, the size of each file was around 200 MB. As each particle corresponds to 28 bytes, each PhSp file stored approximately 7.5×10^6 particles.

The PhSp binary files were later used as particle primary generators to simulations performed with the patient-dependent model. This means that every new particle created in the patient-dependent model carries the parameter values stored inside the PhSp.

The *DicomPrimaryGenerator* class of the patient-dependent model opens the .bin file. A pointer is set to the beginning of the file and starts reading the information stored inside of it. The pointer reads the first 4 bytes and assigns this value to a variable called *particleType*. The pointer than moves on to the 5^{th} byte, reads the next 4 bytes and assigns the value to the variable *energy*. This process goes on until 7 parameters are assigned to 7 variables, namely *particleType*, *energy*, $PosX$, $PosZ$, P_x, P_y and P_z. It is important that the parameters read from the pointer are assigned in the same order used for PhSp generation, otherwise the values would be assigned to the wrong parameter.

A variable named *nbytesper_vertex* is used to define the size of each particle, and is set to 28. After going through 7 steps of 4 bytes, the pointer is now at the 28^{th} byte. The *nbytes_pervertex* sets the pointer "on hold" and one particle is created using the assigned values for particle type, energy, position and momentum. The position along the Y axis was not scored by the PhSp, and is set to 840.005 mm upstream the isocenter, corresponding to the position of the upper boundary of the PhSp plane from the patient-independent model, where the filter for particle scoring was set. The pointer can now move to the next bytes and the process repeats until the end of the binary file.

It is important to notice the maximum number of particles stored inside a PhSp: if the user tries to run a simulation using more particles than available, the simulation will crash. The number of particles inside a PhSp can be determined by dividing the binary file size (in bytes) by the value of bytes per particles (28 in this study).

3.1.9 Leaf optimization

Some parameters of the leaves can vary from the nominal values given by the vendors, such as the density of the tungsten alloy (ρ), the leaf height (h_{leaf}) and the shift of the MLC towards a virtual focus (α) to create the virtual tongue and groove. All of these parameters highly influence the intra and interleaf transmission of the MLC system [144]. Intra and interleaf transmission is the percentage dose deposited in a certain point, due to radiation passing through and between adjacent leaves respectively, with respect to the dose deposited at CAX of an open 10×10 cm^2 field at the same depth [93]. In order to properly model the Elekta Agility$^{\text{TM}}$ MLC used in this work, such parameters needed to be tuned until an agreement between simulated and measured transmission curves was achieved.

The transmission measurement was performed using a water phantom (Blue Phantom 2, IBA Dosimetry, Schwarzenbruck, Germany), and a silicon diode detector (Diode P T60016, PTW, Freiburg, Germany), at 100 cm SSD. The diaphragms were positioned in the maximum open position (± 20 cm at isocenter level) and the leaves were totally closed at -15 cm from isocenter. A lead block was positioned under the gap between closed leaves to minimize leakage. A dose profile was measured at 1.6 cm depth along the inline direction (perpendicular to leaves movement direction) with 1 mm resolution (Fig. 3.20). To determine the transmission, the measured profile was normalized to the dose measured at CAX of an open 10×10 cm^2 at the same depth. Following the same configuration, the simulations were performed inside a virtual water box positioned at 100 cm SSD. The dose for the transmission profile was recorded at 1.6 cm depth in water. The resolution of the scoring grid was set to 1 mm in the inline direction (Z axis), 10 mm along the CAX (Y axis) and 100 mm along the crossline direction (Z axis, one voxel only). The larger dimensions along Y and Z were chosen to improve the signal-to-noise ratio of the simulations [168, 169]. The dose at CAX for the 10×10 cm^2 field, used for normalization of the transmission profile, was simulated at the same depth with one voxel of 4×4 mm^2 in the X-Z plane, and 10 mm along the Y plane. This voxel size was chosen to improve statistics, and assuming the dose is relatively uniform in this small volume.

First, the leaf height and density were modeled using the values listed in Table 3.2 ($h_{leaf} = 90\ mm$ and $\rho = 18.0\ g/cm^3$). These values were kept constant, while different values for the shift to the virtual focus α were investigated ($\alpha = [0.0, 1.5, 2.5, 3.0, 3.5]\ mm$). For each value of α, one pair of transmission profile and dose at CAX was simulated and visually compared to the benchmark measurement. The values of 3.0 mm and 3.5 mm yielded the best results and were further used for the height and density optimization.

The shape of the simulated curves obtained with α values of 3.0 mm and 3.5 mm was in agreement with the measurement. However, the transmission values of both simulations were higher than the measurements, i.e. the simulated profile was “shifted” upwards (more details later on in the 3.2, indicating that the modeled leaves were less attenuating than the physical leaves. Therefore, 18 different combinations of h_{leaf} and ρ, which lead to an

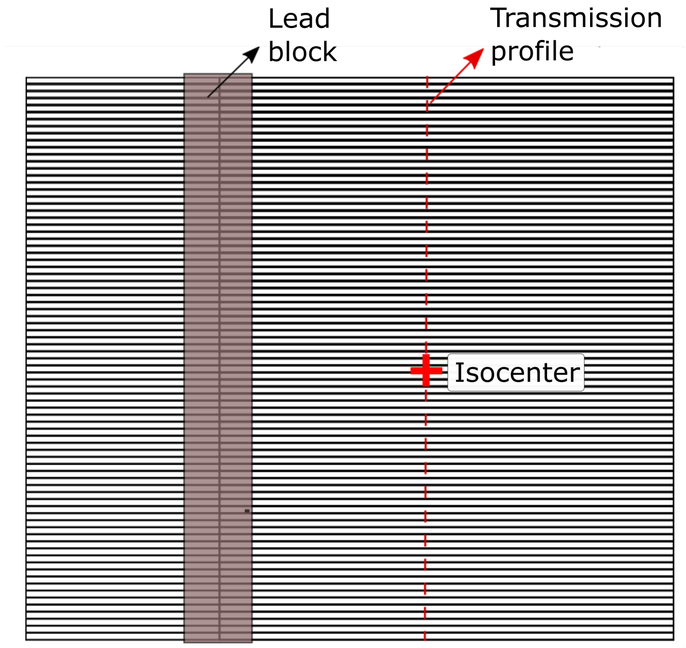

Figure 3.20: Representation of the setup for penumbra measurements and simulations. The leaves were totally closed at -15 cm from isocenter. A lead block, depicted in brown, was positioned under the closed leaves' tips to minimize leakage. The measured/simulated transmission profile direction and location is represented by the red dashed line, while the red cross represents the isocenter. Figure is not to scale.

increased attenuation (i.e. higher density and/or higher leaves), were used for transmission simulations, with each one the pre-selected α values, and fine-tuned until a good agreement with the measurements was achieved. ρ values were varied between 18.0 and 18.4 g/cm^3, and h_{leaf} was varied between 90.0 and 95.0 mm.

3.1.10 Absolute Dose Calibration

The simulations performed for the electron beam and leaf parameters optimization were all used as relative dose values, i.e. normalized to a certain reference point. For simulations of dose deposited inside a patient geometry, the scored dose needs to be properly quantified in terms of absolute values.

As pointed out by J. Fleckenstein [132], the amount of particles necessary to explicitly simulate a relevant dose value in a water voxel (e.g. 2 Gy deposited at a 8 mm^3 voxel) would exceed the calculation possibilities. The mean dose deposited at the same voxel could be obtained with relatively low uncertainty using a reasonable number of primary incident electrons. Moreover, IMRT plans are composed of several fields with different weights, quantified in terms of monitor units (MUs). This field weighting needs to be accommodated in the simulation results, and this can be done using a calibration factor between the MU values and the dose deposited in the scoring volume of a simulation per incident electron in the target (in Gy/electron).

The calibration factor was obtained following the method proposed by Popescu et al. [170], and it relies on the definition of Monitor Units. Monitor Units (MU) are used to quantify the output of clinical linacs, and are measured by the ionization chambers (IC) positioned between the flattening filter and the backscatter plate (Fig. 3.1), hereon called monitor chambers. The monitor chambers are usually calibrated such that 100 MUs correspond to an absorbed dose of 1 Gy at the depth of maximum dose (D_{max}) for a 10×10 cm^2 field in a water phantom positioned at 100 cm SSD [85]. This will be hereon referred to as the calibration condition. The absolute dose $D_{i,abs}$ deposited in a voxel i can be described by Eq. 3.62:

$$D_{i,abs} = D_{i,e} \frac{D_{ch(10\times10)}}{D_{ch}} \frac{D_{abs}^{cal}}{D_{e}^{cal}} U \tag{3.62}$$

$D_{i,e}$ is the dose deposited in voxel i per incident electron in the target, $D_{ch(10\times10)}$ and D_{ch} are the doses deposited in the monitor chambers for a 10×10 cm^2 and one arbitrary fields, respectively. D_{abs}^{cal} is the absolute dose deposited in a voxel for the calibration condition (assigned to 1 Gy for 100 MUs, or 1 cGy/MU), while D_{e}^{cal} is the dose deposited per incident electron at calibration condition. Finally, U is the number of monitor units.

The term $D_{ch(10\times10)}/D_{ch}$ accounts for variable dose deposition in the monitor chambers for variable field sizes, due to backscattered particles from the MLC. If the dose deposited at the monitor chambers does not vary with the field size, this term can be set to 1 [170, 171]. Due to the presence of the backscatter plate in Elekta linacs, the amount of particles backscatter to the monitor chamber is less than 0.35% [172, 173]. Moreover, the dose at the monitor chambers was scored for simulations performed with 50 PhSp (i.e. $N_e = 5 \times 10^9$), for 2×2 cm^2, 5×5 cm^2, 10×10 cm^2, 15×15 cm^2, 20×20 cm^2 and 30×30 cm^2 fields, to evaluate the dependence of the deposited dose with the field size, and has shown negligible fluctuation. Therefore, in this work, the term $D_{ch(10\times10)}/D_{ch}$ has been set to 1, and Eq. 3.62 becomes Eq. 3.63, with C_{abs} being the calibration factor defined in Eq. 3.64.

$$D_{i,abs} = D_{i,e} \frac{D_{abs}^{cal}}{D_{e}^{cal}} U = D_{i,e} C_{abs} U \tag{3.63}$$

$$C_{abs} = \frac{D_{abs}^{cal}}{D_{e}^{cal}} \tag{3.64}$$

Eq. 3.63 yields the absolute simulated dose deposited at a voxel i at any position inside a water phantom or patient geometry [170], taking as input the dose deposited at the same voxel per incident electron and the number of MUs of the field being simulated. The calibration factor C_{abs} is a constant value and can be determined by performing one simulation at the calibration conditions. Hence, the dose deposited at D_{max} in a water phantom, positioned at SSD of 100 cm, due to an irradiation with a 10×10 cm^2 field, was simulated using 150 different PhSp files as input, each generated with 10^8 incident electrons (N_e). The scoring volume (mesh) was constructed to have the same size as the sensitive volume of the Farmer chamber (PTW 30013, PTW, Freiburg, Germany, 0.6 cm^3 sensitive volume), used for the calibration measurement of the physical linac, and centered at 1.6 cm depth in the water phantom, corresponding to the D_{max} of the 6 MV photon beam. The result of each simulation was added and then divided by the total number of incident electrons combined ($N_e = 1.5 \times 10^{10}$) to obtain D_e^{cal}; a value of 1 cGy is attributed to D_{abs}^{cal}, and C_{abs} was easily obtained. The uncertainty of the calibration factor was estimated using the history-by-history method described in the following. If applicable, particle recycling needs to be accounted for in the absolute dose conversion [170]. This is done by dividing Eq. 3.63 by the number of times each particle is used in a simulation ($rec + 1$), where rec is the recycling factor (Eq. 3.65).

$$D_{i,abs} = \frac{D_{i,e} C_{abs} U}{(rec + 1)} \tag{3.65}$$

For IMRT plan simulations, each field was simulated separately and the absorbed dose was converted into absolute dose using the MU values of each individual field, obtained from the RTPlan.dcm files (Appendix A). The sum of the individual fields yields the dose delivered from the entire treatment plan.

3.1.11 Uncertainty Estimation

The accuracy of MC methods relies not only on the proper description of the modeled geometry, but also on the amount of particle histories used in the simulations. Since MC is a statistical method, which relies on random numbers to predict the interactions of particles with the geometries, the resultant simulation comes with an intrinsic statistical uncertainty. Based on the Strong Law of Large Numbers, the simulated dose deposited in one voxel $D_{i,e}$, taken as the average over the total number of incident electrons N_e, converges to the expected value D_i if an infinite number of incident electrons are simulated. However, for a sufficiently large number of initial electrons simulated ($N_e \to \infty$), the resultant value represents the expected value within a certain statistical uncertainty.

In this work, the statistical uncertainty of the dose simulated to any voxel at a virtual phantom or patient was estimated using the history-by-history approach [174, 175]. The dose deposited in a certain voxel i per incident electron $D_{i,e}$ can be described as the arithmetic mean $\overline{D_i}$ of the statistically independent dose deposits from all the incident electrons N_e in one simulation (Eq. 3.66):

$$D_{i,e} = \overline{D_i} = \sum_{j=1}^{N_e} \frac{D_j}{N_e} \tag{3.66}$$

According to the central limit theorem, when $N_e \to \infty$, the probability distribution of $\overline{D_i}$ follows a Gaussian distribution. Therefore, the standard deviation of the mean $\sigma_{\overline{D_i}}$

is given by Eq. 3.67.

$$\sigma_{\overline{D_i}} = \sqrt{\frac{1}{N_e - 1}\left(\sum_{j=1}^{N_e}\frac{D_j^2}{N_e} - \left(\sum_{j=1}^{N_e}\frac{D_j}{N_e}\right)^2\right)} \tag{3.67}$$

The statistical uncertainty throughout this work was estimated using Eq. 3.67, including the statistical uncertainty of the absolute calibration factor σ_{Cabs}. For absolute dose results, calculated using Eq. 3.63, the total uncertainty was estimated as the combination of $\sigma_{\overline{D_i}}$ and σ_{Cabs}, using the rules of uncertainty propagation. Hence, for any voxel, the absolute dose uncertainty is given by Eq.3.68:

$$\sigma_{D_{i,abs}} = \left(\frac{U}{rec+1}\right)\sqrt{(D_{i,e}^2(\sigma_{Cabs})^2 + (C_{abs})^2(\sigma_{D_{i,e}})^2} \tag{3.68}$$

3.1.12 Validation

The process of modeling the Elekta Synergy® linac head coupled with the Agility™ MLC involved many challenging optimization steps, namely the electron beam source and the tunable MLC geometric parameters. Moreover, the position of the leaves and diaphragms was modeled following analytic equations derived mainly from data found in the literature. In order to properly simulate the dose delivered to a patient geometry, it is important to verify that the model is able to reproduce different sets of benchmarks. The benchmarks can be either measurements or calculations performed by a commissioned calculation engine, such as a commissioned TPS.

First, the Geant4 model of the linac was validated against a set of profiles measured inside a water phantom for 2×2 cm^2, 5×5 cm^2, 10×10 cm^2 and 30×30 cm^2 fields. PDD and lateral profiles at different depth inside the virtual water (both along inline and crossline directions), were simulated with 200 PhSp files with particle recycling ($rec = 19$), following the descriptions in 3.1.7. Both the measured and simulated data were normalized to the dose at maximum depth D_{max} (for PDDs) and to the dose at the central axis (for lateral profiles). This procedure verifies if the incident electron beam has been properly modeled. For this step, the validation is performed using visual inspection.

Next, the inter and intraleaf transmission was simulated following the descriptions in 3.1.9, using 150 PhSp files and a recycling factor of 6, and visually compared to measurements performed under the same conditions. This validation verifies if the tunable MLC parameters have been properly optimized.

To verify if both the physical and virtual T&G have been properly modeled, the picked-fence test was used. The picked-fence test is composed of twelve different fields, with several dents where interdigitating leaves from opposite sides intertwine, as illustrated in Fig. 3.21 [176]. Due to its high spatial resolution, the EPID attached to the linac (ELEKTA iViewGT®) has been used to measure this irradiation pattern. The fields of the picked-fence test were irradiated directly to the EPID. The picked-fence test fields were simulated individually, each one using 150 PhSp files with particle recycling factor of 6. The dose at the virtual EPID was recorded using a scoring mesh positioned at 3.0 cm depth in the virtual EPID (3.1.4), with 0.4-mm resolution, matching the spatial resolution of the physical EPID [148]. Both measurements and simulations have been normalized to the dose received by the voxels at CAX. Finally, the simulated fields were summed up and compared to the measurements.

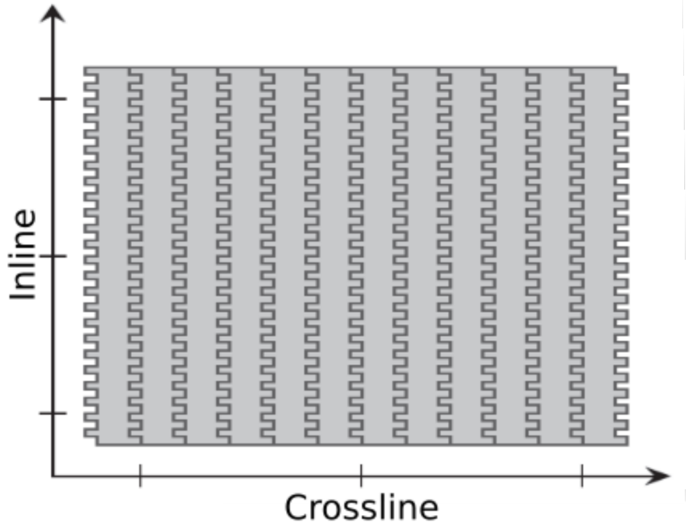

Figure 3.21: Representation of the twelve different fields comprising the picked-fence test, used for investigation of the tongue-and-groove effect. Figure is not to scale. Reprinted with permission from S. Neppl [176].

To verify the absolute dose calibration, the water phantom (Blue Phantom 2, IBA Dosimetry, Schwarzenbruck, Germany) was irradiated with 300 MUs for different square fields. Point dose values, measured with a Farmer chamber (PTW 30013, PTW, Freiburg, Germany) at D_{max} in water were used as benchmarks. Simulations were performed under the same conditions with the scoring volume matching the sensitive volume of the chamber used. The resultant dose from simulations were converted into absolute dose values following the methodology described in section 3.1.10 and compared to the measurements. The relative difference was used to quantify the agreement between benchmark measurements and simulations.

Subsequently, validations for 3D dose distributions were performed. In this step, a cylindrical water equivalent phantom (Octavius 4D, PTW, Freiburg, Germany) and a respective virtual model were used. The virtual model was constructed using a CT image from the phantom. For simulations, the scoring mesh was created as a 1.8×1.8×1.8 mm^3 grid, perfectly matching the voxel size of the CT image. Moreover, the benchmark dose distributions were calculations performed using a commissioned TPS [177, 178], with 2.5×2.5×2.5 mm^3 calculation grid size. First, different squared fields were analyzed (2×2, 5×5, 10×10 and 26×26 cm^2) and both dose profiles (PDDs and lateral profiles) as well as the 3D dose distributions, in absolute dose values, were used for validation. For the profiles, the agreement was evaluated with visual inspection; for the 3D dose distributions, the gamma evaluation method (introduced in section 3.1.13) was used. The simulations of all the squared fields were performed with particle recycling factor of 14 and 50 PhSp files.

Finally, 3D dose distributions inside the Octavius phantom from two clinical step-and-shoot IMRT, one head and neck (H&N) and one prostate plan, were calculated using the TPS. Both plans were calculated with a 6 MV photon beam nominal energy. The gantry rotation, leaves and diaphragms positions and MU values of each field from both plans were extracted from the RTPlan.dcm files and used as input to the Geant4 model. Each field was simulated with 30 PhSp files and particle recycling factor of 24. The resultant simulations were summed to obtain the 3D dose distribution for the total treatment plan, and converted into absolute dose values. Tables 3.7 and 3.8 list the details of both IMRT fields. The agreement to the TPS calculations, also in absolute values, was evaluated in terms of the gamma evaluation (3.1.13). Table 3.9 summarizes the validation process.

Table 3.7: Characteristics of the head and neck (H&N) step-and-shoot IMRT plan used for 3D dose validation. The plan was calculated with a 6 MV photon beam nominal energy.

H&N IMRT Beam	1	2	3	4	5	6	7	8	9
Gantry Angle (Deg)	160	85	55	25	0	335	305	275	200
Control Points (CP)	10	9	13	10	12	8	11	9	10
Total MU/Beam	80.9	58.0	71.9	47.1	56.8	40.2	62.1	75.8	77.4

Table 3.8: Characteristics of the prostate step-and-shoot IMRT plan used for 3D dose validation. The plan was calculated with a 6 MV photon beam nominal energy.

Prostate IMRT Beam	1	2	3	4	5	6	7	8	9
Gantry Angle (Deg)	180	140	100	60	20	340	300	260	220
Control Points (CP)	7	7	7	6	8	7	7	5	7
Total MU/Beam	74.9	85.1	101.4	76.4	77.9	82.8	103.3	70.1	92.2

Table 3.9: Summary of the dose distributions used for validation of the entire MC model.

Dose Distribution	Irradiation geometry	Dose	Reference Distribution	Evaluation
PDD	Squared fields at water phantom	Normalized	Measurements with IC	Visual inspection
Lateral profiles	Squared fields at water phantom	Normalized	Measurements with IC	Visual inspection
T&G effect	Picked-fence at EPID	Normalized	Measurements with EPID	Visual inspection
PDD	Squared fields at Octavius	Normalized	TPS calculations	Visual inspection
Lateral profiles	Squared fields at Octavius	Normalized	TPS calculations	Visual inspection
3D	Squared fields at Octavius	Absolute value	TPS calculations	Gamma index
3D	Step-and-shoot IMRT	Absolute value	TPS calculations	Gamma index

3.1.13 Gamma Index Evaluation

The gamma index evaluation, or simply gamma evaluation, introduced by Low et al. [179, 180] is a quantitative method to compare two different dose distributions, widely used in radiotherapy. The gamma evaluation takes into consideration not only differences in the dose values, based on a maximum dose difference (DD) criterion, but also geometric discrepancies, based on a maximum distance-to-agreement (DTA) criterion. The DD and DTA criteria define a region of interest (ROI). In the gamma evaluation, each voxel of a reference dose distribution (RD, index ref) is compared to all the voxels in an evaluated distribution (ED, index ev) located inside the ROI, in terms of the DD and DTA, yielding a gamma value Γ defined by Eq. 3.69,

$$\Gamma(\vec{r}_{ref}, \vec{r}_{ev}) = \sqrt{\frac{|\vec{r}_{ev} - \vec{r}_{ref}|}{DTA^2} + \frac{D_{ev}(\vec{r}_{ev}) - D_{ref}(\vec{r}_{ref})}{DD^2}} \tag{3.69}$$

where $\vec{r}_{ref}$ and $\vec{r}_{ev}$ are the position of the reference and evaluated voxels inside the ROI, respectively, and $D_{ev}(\vec{r}_{ev})$ and $D_{ref}(\vec{r}_{ref})$ are the dose values of the reference and evaluated voxels, respectively. The first term in Eq. 3.69 refers to the distance to agreement between the two voxels, while the second term to dose difference. The point in the ED with the lowest Γ value is considered to be the best match: the resulting γ value for every element in the RD is the minimum of all Γ entries, as defined by Eq. 3.70:

$$\gamma(\vec{r}_{ref}) = min\{\Gamma(\vec{r}_{ref}, \vec{r}_{ev})\}\forall\{\vec{r}_{ev}\} \tag{3.70}$$

A fail/pass criteria, based on the DD and DTA combined analysis, can then be defined for every evaluated voxel in the ROI as Eq. 3.71:

$$\begin{aligned} &\gamma(\vec{r}_{ref}) \leq 1, \text{ evaluated voxel passes,} \\ &\gamma(\vec{r}_{ref}) > 1, \text{ evaluated voxel fails.} \end{aligned} \tag{3.71}$$

The gamma evaluation was implemented using MATLAB (Mathworks Inc., Natick, MA). In this work, the resultant γ value given by Eq. 3.70 is multiplied by the sign of the dose difference between the reference and evaluated voxels [96]; the γ value thus indicates if an undedosage ($\gamma < 0$) or an overdosage ($\gamma > 0$) is found in the evaluated voxel. Moreover, the implementation script takes into account the different voxel sizes of the reference and comparison dose distributions. The reference and comparison dose distributions were the benchmark (i.e. TPS dose calculations) and the simulations, respectively.

For all validations performed with the gamma evaluation, DD and DTA values of 3% and 3 mm were used, respectively, and will be denoted as (3%,3 mm). The DD and DTA values are limited by the level of uncertainty in the MC dose simulations (2.8% on average) and by the grid size from both MC and TPS dose distributions (1.8 mm and 2.5 mm, respectively).

Finally, the passing rate, i.e. the proportional amount of voxels passing the gamma evaluation, can be used as a measure of the agreement between evaluated and reference dose distributions. Only voxels located inside the outer borders of the Octavius phantom were considered for calculating the passing rates. Low dose voxels can lead to an overestimation of the gamma passing rates depending on the DD and DTA criteria chosen. To avoid this bias, a dose threshold can be used to select only voxels receiving at least a certain fraction of the maximum dose. A threshold of 25% was used.

3.2 Results

3.2.1 Incident electron beam optimization

The incident electron beam energy optimization process relied mostly on PDD curves of several squared fields. The visual inspection of the agreement between simulated and measured PDDs led to several different sets of sub-optimal values for $\overline{E}$ and $FWHM_E$. The focal spot optimization process followed considering all the sub-optimal sets. It was rapidly noticed, however, that for lower energy values, simulated lateral profiles for bigger fields, especially the 30×30 cm^2, presented very pronounced shoulders at the offset region of the profile, clearly not matching the reference measurement (Fig. 3.22). The discrepancy decreased with increasing mean energy values, confirming that the lateral profiles are sensitive to the beam mean energy [132, 159], especially for larger fields. Therefore, during the focal spot optimization, it might be necessary to reevaluate the values for the mean beam energy before a final set of initial beam parameters can be determined, and it is recommended to include lateral profiles of bigger fields in the beam energy optimization process.

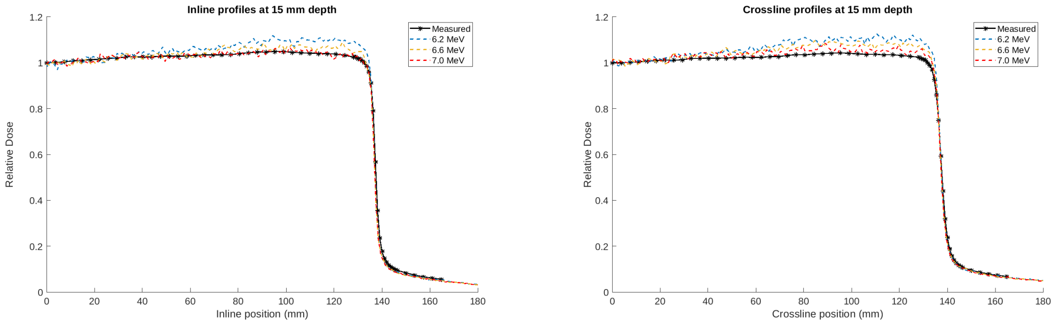

(a) Inline **(b)** Crossline

Figure 3.22: Measured and simulated inline (a) and crossline (b) profiles at 15 mm depth in water. The simulations were performed with three different PhSp files generated with initial electron beam mean energy $\overline{E}$ of 6.2, 6.8 and 7.0 MeV. For lower energies, the simulated profiles have higher shoulders than the measured profiles at the off axis position.

The energy parameters of $\overline{E} = 7.0$ MeV and $FWHM_E = 0.5$ MeV resulted in the best agreement of PDDs, for several fields, and shoulder regions, for lateral profiles of the 30×30 cm^2 field. In Fig. 3.22 it can be observed that, despite becoming better, the agreement between the measured and simulated profiles is still bad, especially along the crossline direction. This is because the shape of the shoulders depends not only on the energy of the beam, but also on the focal spot shape. Therefore, the focal spot optimization continued using $\overline{E} = 7.0$ MeV and $FWHM_E = 0.5$ MeV. This optimization process relied on lateral profiles at two different depths in water for different field sizes. Both the shoulders and the penumbra regions were used as guidance to iteratively find the best set of beam focal spot size and angular spread. Figs. 3.23 and 3.24 show how the shoulders and the penumbra region varies for the different beam parameters listed in Table 3.10, the crossline profiles being more sensitive to the focal spot parameters.

Finally, it was found that an elliptical beam, with an asymmetric spatial distribution and angular spread, resulted in the best agreement between measured and simulated profiles for all the fields and depths considered (Beam N. 12 in Table 3.10). These values were considered the optimal electron beam parameters for modeling the 6 MV photon beam of the Elekta Synergy® used in this study. 500 PhSp files with 10^8 initial electrons were generated with the final optimized parameters and used for validation.

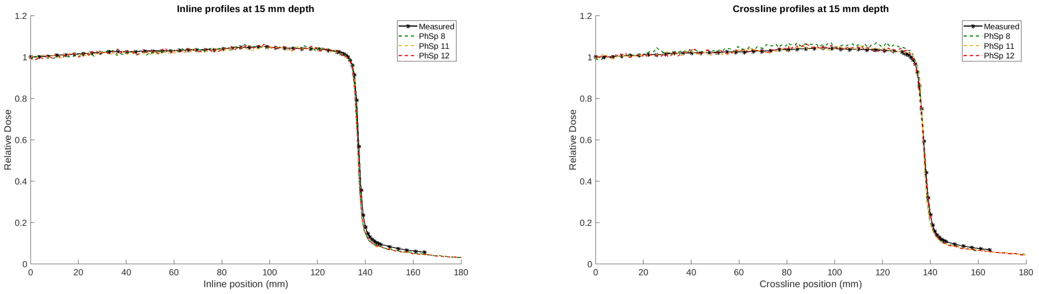

(a) Inline **(b)** Crossline

Figure 3.23: Measured and simulated inline (a) and crossline (b) profiles at 15 mm depth in water. The simulations were performed with three different PhSp files generated with different initial electron beam source types (Table 3.10).

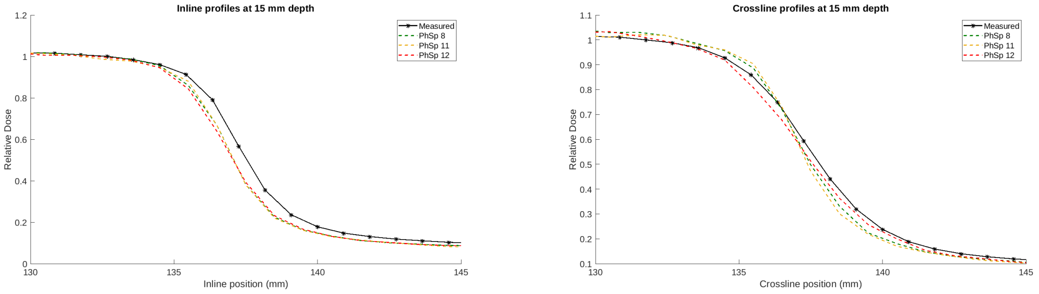

(a) Inline **(b)** Crossline

Figure 3.24: Measured and simulated inline (a) and crossline (b) profiles at 15 mm depth in water, with emphasis at the penumbra region. The simulations were performed with three different PhSp files generated with different initial electron beam source types (Table 3.10).

Table 3.10: Three different beam parameter sets investigated during focal spot optimization, together with the nominal values provided by the vendor [132]. Beam N. 12, marked with *, is the final optimized parameter set.

Beam N.	Shape	$\overline{E}$ (MeV)	$FWHM_E$ (MeV)	σ_{in} (mm)	σ_{cr} (mm)	Φ_{in} (rad)	Φ_{cr} (rad)
8	Point	7.0	0.5	-	-	0.011	0.020
11	Point	7.0	0.5	-	-	0.011	0.011
12*	Elliptical	7.0	0.5	0.5	1.0	0.011	0.014
Nominal	Circular	6.5	0.5	0.5	0.5	-	-

3.2.2 Leaf optimization

Using the PhSp files produced with the optimal incident electron beam as particle source generator in the patient-independent model, the tunable leaf parameters were optimized (3.1.9), starting with the shift to virtual focus α. Fig. 3.25a shows the measured transmission profile, together with the simulations obtained using all values of α investigated. For these simulations, the other tunable parameters h_{leaf} and ρ were modelled using the nominal values from Table 3.2. An $\alpha = 0\ mm$ represents the scenario where the leaf bank is not rotated, and the leaves are positioned following the beam divergence, justifying the high transmission values. Once the leaf bank is rotated, the transmission drops significantly, even for a small shift of $\alpha = 1.5\ mm$. For better visualization, Fig. 3.25b displays the transmission curves without the unrotated case.

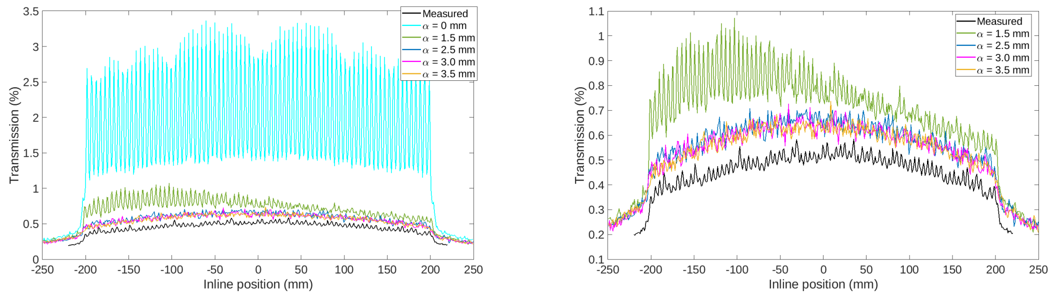

(a) **(b)**

Figure 3.25: Transmission curves simulated using all investigated α values (a), together with the measured curve; figure (b) excludes the non-shift case ($\alpha = 0\ mm$). h_{leaf} and ρ were modelled using the nominal values (Table 3.2).

From Fig.3.25, it is clear that the simulated transmission curves with $\alpha = [2.5, 3.0, 3.5]\ mm$ come closer to the measured reference. However, it is possible to observe that the shape of the simulated curves with $\alpha = [3.0, 3.5]\ mm$ are more similar to the measured curve, while the curve with $\alpha = 2.5\ mm$ has slightly higher values towards the negative inline positions (Fig. 3.26).

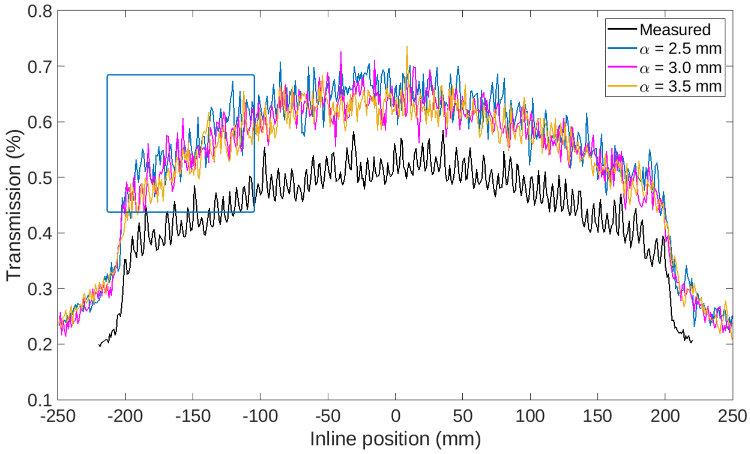

Figure 3.26: Simulated transmission curves with α values of 2.5 mm, 3.0 mm and 3.5 mm. With respect to the measured reference (in black) all three simulated curves present higher transmission, however, their shape resemble the measured curve. For the 2.5 mm shift (blue curve) the left part of the curve was higher than the other two simulated curves (region inside the blue box). h_{leaf} and ρ were modelled using the nominal values (Table 3.2).

Despite the good shapes, the simulated transmission values are above the measured values, indicating that the modeled leaves are not as attenuating as the physical leaves (3.1.9). The optimization of h_{leaf} and ρ should be towards more attenuating conditions, i.e. higher leaves and/or higher densities. By individually analyzing the curves with $\alpha = 3.0\ mm$ and $\alpha = 3.5\ mm$, it is not possible to identify one best match (Fig. 3.27). Hence, the leaf optimization continues using both values of $\alpha = 3.0\ mm$ and $\alpha = 3.5\ mm$, combined to the different values of h_{leaf} and ρ.

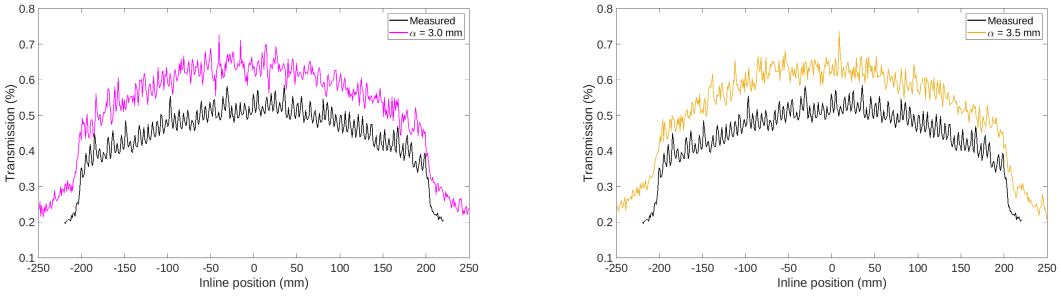

(a) $\alpha = 3.0\ mm$

(b) $\alpha = 3.5\ mm$

Figure 3.27: Individual comparison of the measured transmission curve against simulations performed with $\alpha = 3.0\ mm$ (a) and $\alpha = 3.5\ mm$ (b). h_{leaf} and ρ were modelled using the nominal values (Table 3.2).

For both $\alpha = 3.0\ mm$ and $\alpha = 3.5\ mm$, simulations performed with $h_{leaf} = 95\ mm$ and $\rho = 18.4\ g/cm^3$ resulted in good agreement between measured and simulated transmission curves (Fig. 3.28).

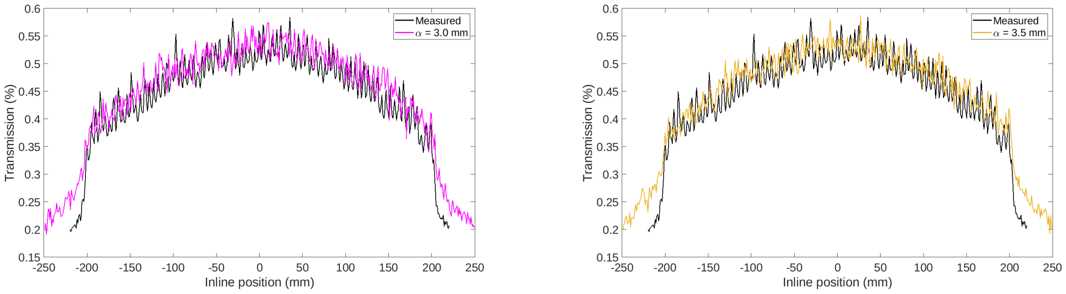

(a) $\alpha = 3.0\ mm$ **(b)** $\alpha = 3.5\ mm$

Figure 3.28: Individual comparison of the measured transmission curve against simulations performed with $\alpha = 3.0\ mm$ (a) and $\alpha = 3.5\ mm$ (b), with the optimal values for h_{leaf} and ρ values of 95 mm and 18.4 g/cm^3, respectively.

By the time the leaf optimization started, no information was found in the literature regarding the value of α. However, during the optimization, Onizuka et al. [138] reported a leaf bank rotation value of $\approx 0.9\ mrad$, which corresponds to a shift value of $\alpha = 3.25\ mm$. The transmission was then simulated using the best ρ and h_{leaf} combination with $\alpha = 3.25\ mm$, yielding good agreement with the measurements (Fig. 3.29). Hence, $\alpha = 3.25\ mm$ was taken as the final optimal value to be consistent with the published data. Finally, the optimal set of leaf parameters for modeling the Agility$^{\mathrm{TM}}$ MLC used in this work were $\alpha = 3.25\ mm$, $h_{leaf} = 95\ mm$ and $\rho = 18.4\ g/cm^3$ (Table 3.11). Since the leaves and diaphragms have the same density, according to the vendor-provided information [136], the diaphragms were also modeled with $\rho = 18.4\ g/cm^3$.

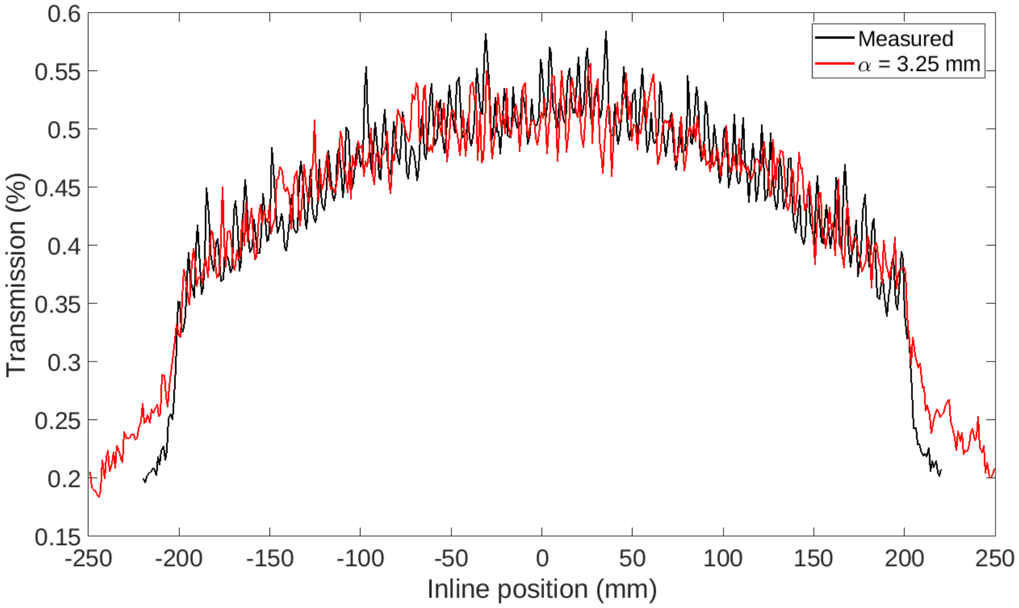

Figure 3.29: Measured transmission curve and simulations performed with the final optimal set of leaf parameters: $\alpha = 3.25\ mm$, $h_{leaf} = 95\ mm$ and $\rho = 18.4\ g/cm^3$.

Table 3.11: Final optimal set of leaf parameters used for modeling the Agility$^{\text{TM}}$ MLC.

α (mm)	$h_{leaf}(mm)$	ρ (g/cm^3)
3.25	95.0	18.4

3.2.3 Absolute Dose Calibration

The resultant absolute dose calibration factor C_{abs} obtained with 150 PhSp files, simulated in the calibration condition (as described in section 3.1.10), is equal to:

$$C_{abs} = (7.245 \pm 0.073) \times 10^{13} = 7.245 \times 10^{13} \pm 1.005\% \tag{3.72}$$

The fluctuation of dose deposited in the monitor chamber due to backscatter radiation from the MLC was tested for different field sizes. Fig. 3.30 shows the dose of all squared fields simulated, normalized to the dose value for the 10×10 cm^2 field, while table 3.12 lists the relative difference of the doses with respect to the value for the 10×10 cm^2 field. The maximum absolute relative difference observed is for the 30×30 cm^2, and amounts to 0.24%. Therefore, for field sizes ranging from 2×2 cm^2 to 30×30 cm^2, the dose at the monitor chamber changed by less than 0.3%, in agreement with results reported in the literature [172, 173]. This result confirms that no extra correction for backscatter radiation in the monitor chamber is needed for absolute dose conversion.

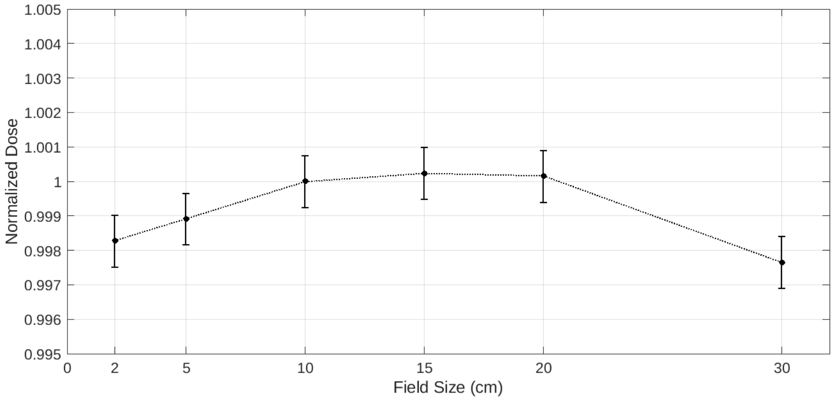

Figure 3.30: Simulated dose deposited at the monitor chamber for different field sizes, normalized by the value of the 10×10 cm^2 field. A maximum absolute relative difference of 0.24% was observed for the 30× 30 cm^2.

Table 3.12: Relative dose difference (%) of simulated dose deposited at the monitor chamber for different field sizes, with respect to the value of the 10×10 cm^2 field.

Field Size (cm^2)	Relative Difference (%)
2×2	-0.17
5×5	-0.11
10×10	0
15×15	0.02
20×20	0.02
30×30	-0.24

3.2.4 Validation

After the incident electron beam and leaf parameters were properly optimized, several validation simulations were performed. Simulations were performed using the final PhSp files generated with the optimal electron beam parameters (Table 3.10), in combination with the patient-dependent model with leaves and diaphragms modeled using the optimal parameters found (Table 3.11). The simulated data were compared to reference benchmarks as described in 3.1.12.

Fig. 3.31 shows the simulated PDDs profiles (red) for the 2×2 cm^2, 5×5 cm^2, 10×10 cm^2 and 30×30 cm^2 fields, together with the reference measurements (blue). The PDD curves are analyzed in terms of normalized dose with respect to the dose at D_{max}, and the curves have been shifted vertically for better visualization. Based on visual inspection, the simulated curves agree well to the reference measurements. Simulated lateral profiles along inline and crossline directions, at both 15 mm and 100 mm depth in water, are also in good agreement with the measured references for all fields analyzed. Fig. 3.32 and 3.33 show the simulated (red) and reference measured (blue) curves for all fields at 15 mm and 100 mm depth, respectively. The curves are normalized by the value at the CAX (0 mm) and the plots are shifted vertically for visualization. The good agreement between the benchmark measurements and the simulations, performed with the final PhSp files, indicates that the incident electron beam has been successfully optimized and the values listed in Table 3.10 properly represent the incident beam of the physical linac. This can also be an indication of the quality of the leaves and diaphragms positioning, serving as an initial validation of the analytic equations from sections 3.1.3.3, 3.1.3.4, 3.1.3.5 and 3.1.3.6. The tunable parameters of the leaves have been validated using the inter and intraleaf leakage (Fig. 3.29).

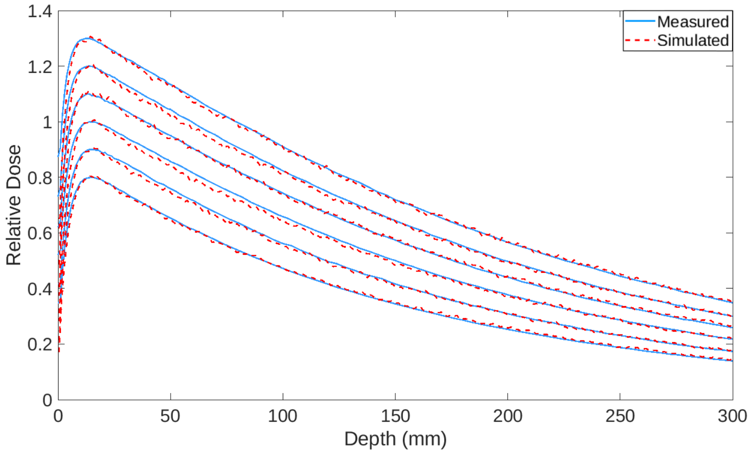

Figure 3.31: Measured (blue) and simulated (red) PDD profiles in the water phantom for all fields used for validation. Both measured and simulated profiles are normalized to the value at the depth of maximum dose (D_{max}). A good agreement can be observed between calculated and simulated profiles throughout the entire curve. The curves have been shifted vertically for better visualization.

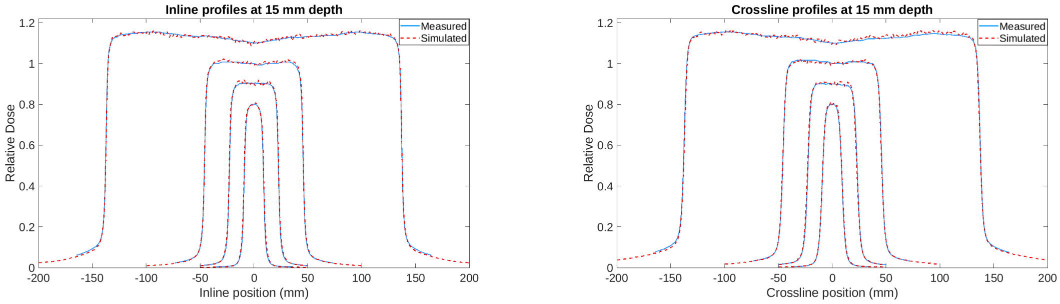

(a) Inline

(b) Crossline

Figure 3.32: Measured (blue) and simulated (red) inline (a) and crossline (b) profiles for the different fields used for validation, at 15 mm depth in the water phantom. Both measured and simulated profiles are normalized to the value at CAX. The simulated profiles are in good agreement with the reference measurements. The curves have been shifted vertically for better visualization.

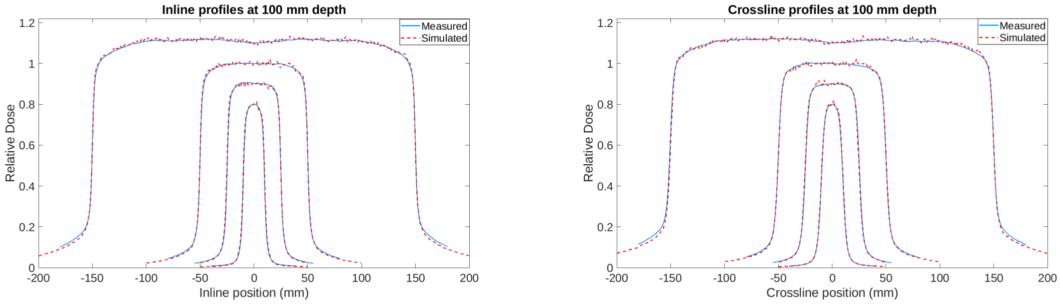

(a) Inline

(b) Crossline

Figure 3.33: Measured (blue) and simulated (red) inline (a) and crossline (b) profiles for the different fields used for validation, at 100 mm depth in the water phantom. Both measured and simulated profiles are normalized to the value at CAX. The simulated profiles are in good agreement with the reference measurements. The curves have been shifted vertically for better visualization.

The T&G model of the leaves has been validated using the picked-fence test. The measured and simulated doses (normalized to the dose at CAX) are shown in Fig. 3.34a and 3.34a, respectively. To evaluate the magnitude of the T&G effect, a dose profile along the inline direction was used from both dose distributions, illustrated in Fig. 3.35; for better visualization, Fig.3.36 shows only the central region around the CAX (100 mm wide). The simulated profile is reasonably in good agreement with the measurement, indicating that the T&G of the MLC, both virtual and physical, have been properly modeled. The T&G effect results in an underdosage of about 20% on both measurements and simulations, in accordance to values reported by other studies [93, 181].

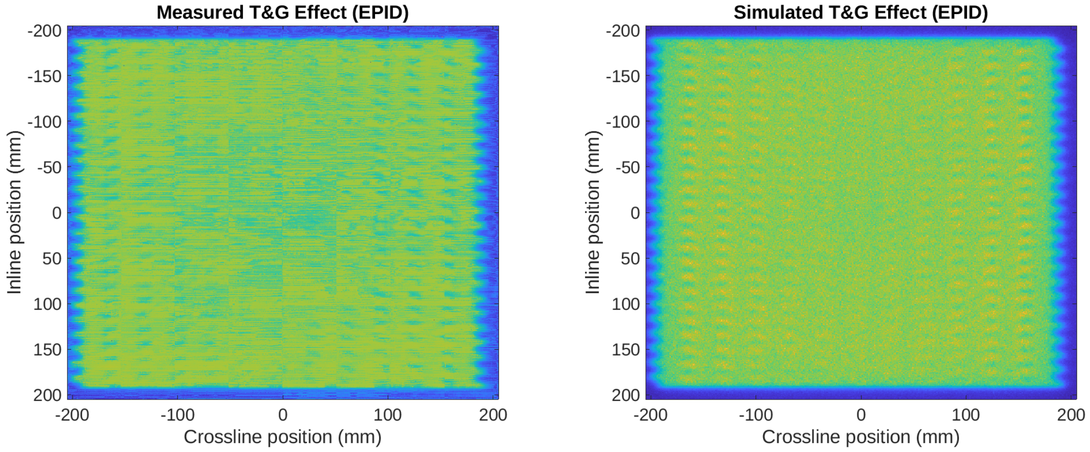

(a)

(b)

Figure 3.34: The picked-fence test measured with the EPID (a) and simulated with the virtual EPID (b). The systematic variation of the signal is caused by the tongue-and-groove effect.

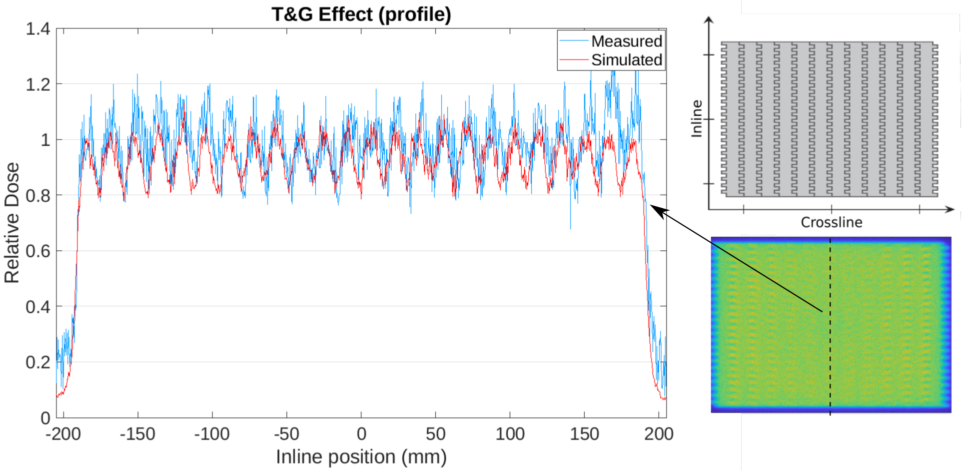

Figure 3.35: Dose profile (figure left) along the inline direction (dotted line at figure right, bottom) used to evaluate the T&G effect. The picked-fence test used for this validation step is also displayed (figure right, top).

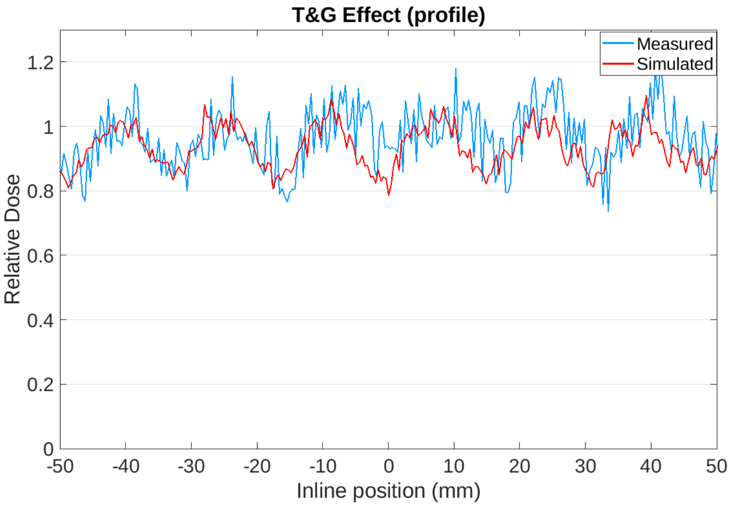

Figure 3.36: The T&G effect at the region around the CAX (100 mm wide). An underdosage of about 20% is observed for both measurement and simulation. The simulated profile is in good agreement with the measurement, indicating that the T&G of the Agility MLC, both virtual and physical, have been properly modeled in the Monte Carlo model. Reprinted with permission from [182].

Point dose measurements for different fields irradiated with 300 MU were used to evaluate the absolute dose calibration factor (3.2.3). Fig. 3.37 shows the measured and simulated dose values for the 5×5 cm^2, 15×15 cm^2, 20×20 cm^2 and 30×30 cm^2 fields. The measured values match the simulated values within the simulation uncertainty. The measured and simulated absolute dose values for all fields are listed in Table 3.13, together with their relative difference. The relative differences for all fields are under 1% (absolute values), with the maximum of 0.7% observed for the 15×15 cm^2 field.

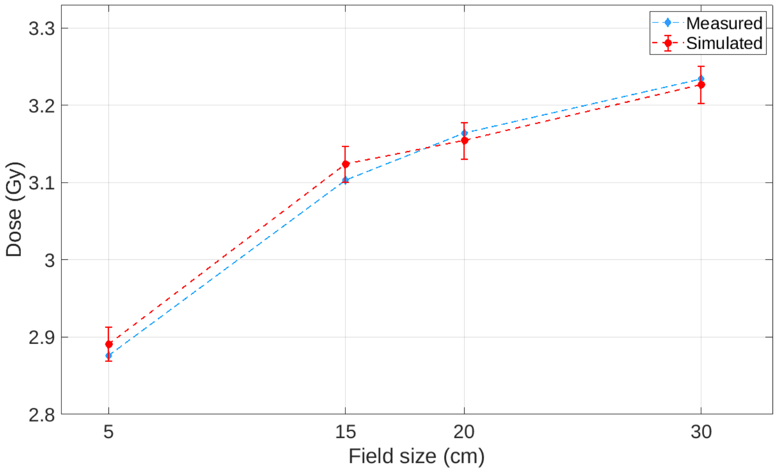

Figure 3.37: Point dose values at the depth of maximum dose D_{max}, measured (blue) and simulated (red), of different squared fields for an irradiation with 300 MUs. The values are given in absolute dose values to validate the absolute dose calibration factor. The measured values match the simulated values within the simulation uncertainty.

The dose distribution inside the Octavius phantom was simulated and compared

Table 3.13: Measured and simulated absolute point dose values of different squared fields, for an irradiation with 300 MUs, together with the relative difference between them.

Field Size (cm^2)	Measured Dose (Gy)	Simulated Dose (Gy)	Relative Difference (%)
5×5	2.876	2.891	-0.5
15×15	3.103	3.123	-0.7
20×20	3.164	3.154	-0.3
30×30	3.234	3.227	-0.22

to calculations performed with the TPS, for the 2×2 cm^2, 5×5 cm^2, 10×10 cm^2 and 26×26 cm^2 squared fields. The PDDs and lateral profiles (at 15 mm depth) are shown in Fig. 3.38 and 3.39, respectively; the curves have been shifted vertically for better visualization. A fairly good agreement can be observed between calculated and simulated profiles.

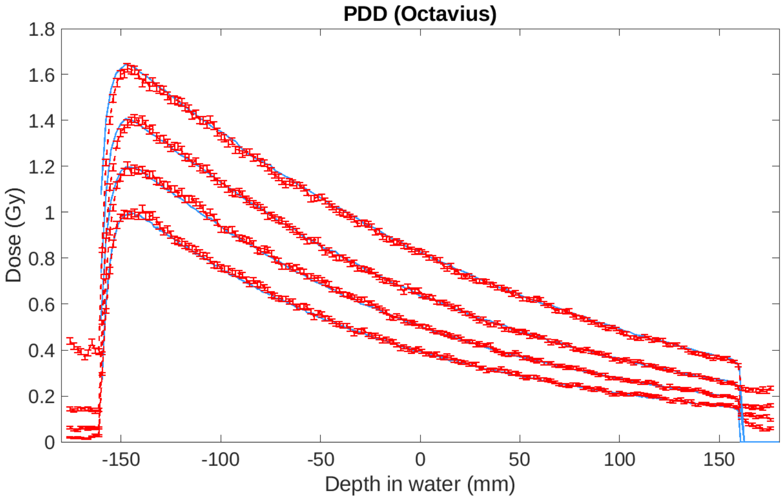

Figure 3.38: TPS (blue) and MC (red) PDD profiles in the Octavius phantom for all fields used for validation. Both measured and simulated profiles are given in absolute dose values. A good agreement can be observed between calculated and simulated profiles inside the phantom. The curves have been shifted vertically for better visualization.

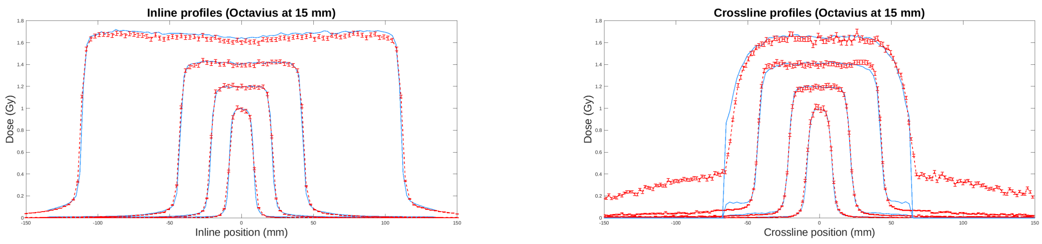

(a) **(b)**

Figure 3.39: TPS (blue) and MC (red) inline (a) and crossline (b) profiles in the Octavius phantom for all fields used for validation. Both measured and simulated profiles are given in absolute dose values. A good agreement can be observed between calculated and simulated profiles inside the phantom. The curves have been shifted vertically for better visualization.

The 3D dose distributions for all the squared fields were evaluated using the gamma evaluation. Using the 26×26 cm^2 field as an example, Fig. 3.40 shows the TPS calculated and simulated dose distribution in one slice of the axial plane, together with the gamma value distribution. In the gamma distribution, red voxels represent an overderdosage of the simulations with respect to the reference calculations, while blue voxels (inside the phantom) represent an underdosage; green voxels pass the (3%, 3 mm) criteria. Overall, a good agreement can be observed between the calculated and simulated dose distributions, with a passing rate of 96.7%.

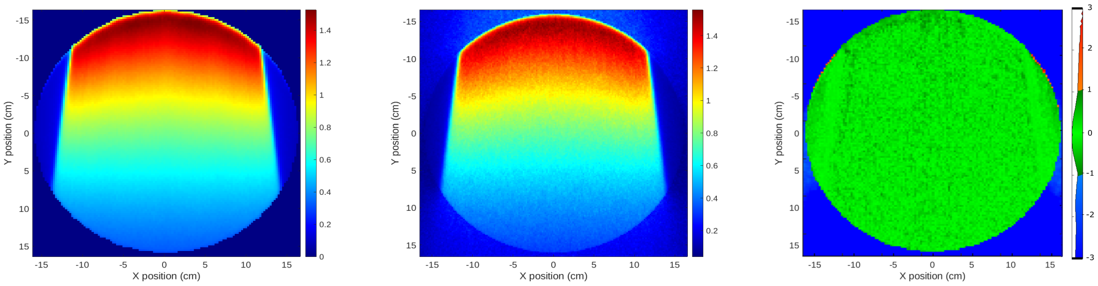

(a) TPS **(b)** Geant4 **(c)** Gamma

Figure 3.40: (a)TPS calculation, (b) MC simulation and (c) respective gamma evaluation of the axial plane of the Octavius phantom, irradiated with the 26×26 cm^2 field. In the gamma evaluation, voxels in green correspond to voxels passing the (3%/3mm) criteria, red voxels represent an overderdosage of the simulations with respect to the reference calculations, and blue voxels (inside the phantom) represent an underdosage.

For the H&N and prostate IMRT plans, the TPS calculations and MC simulations are presented in Fig. 3.41 and 3.42, respectively, together with their respective gamma distributions. The gamma evaluation was performed with (3%, 3 mm) criteria. For both IMRT plans, a reasonably good agreement was observed between the dose distributions. Finally, the gamma passing rates for all the fields evaluated using the Octavius phantom are presented in Table 3.14. For all fields, a satisfactory agreement was observed between the compared distributions, with passing rates higher than 96.6%.

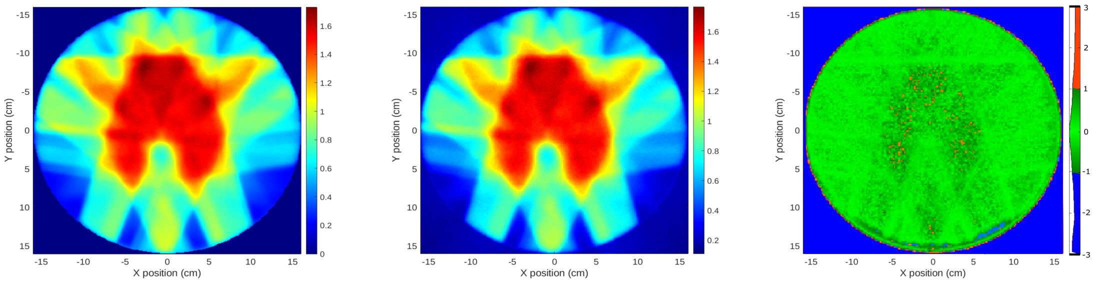

(a) TPS **(b)** Geant4 **(c)** Gamma

Figure 3.41: (a) TPS calculation, (b) MC simulation and (c) respective gamma evaluation of the axial plane of the Octavius phantom, irradiated with the H&N IMRT plan. In the gamma evaluation, voxels in green correspond to voxels passing the (3%/3mm) criteria, red voxels represent an overderdosage of the simulations with respect to the reference calculations, and blue voxels (inside the phantom) represent an underdosage.

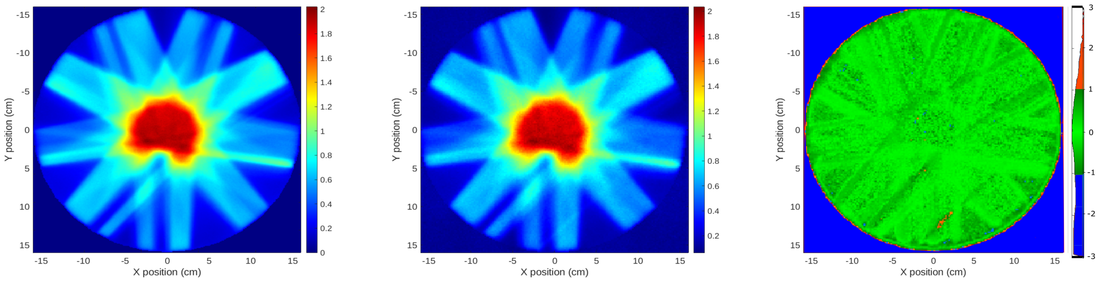

(a) TPS **(b)** Geant4 **(c)** Gamma

Figure 3.42: (a) TPS calculation, (b) MC simulation and (c) respective gamma evaluation of the axial plane of the Octavius phantom, irradiated with the prostate IMRT plan. In the gamma evaluation, voxels in green correspond to voxels passing the (3%/3mm) criteria, red voxels represent an overderdosage of the simulations with respect to the reference calculations, and blue voxels (inside the phantom) represent an underdosage.

Table 3.14: Gamma passing rates for the different irradiation investigated with the Octavius phantom, according to the (3%/3mm) criteria. Only voxels positioned inside the phantom and scoring more than 25% of the maximum dose value were considered.

Geometry	Irradiation	Passing Rate (%)
Octavius	2×2	99.6
Octavius	5×5	96.6
Octavius	10×10	97.2
Octavius	26×26	96.7
Octavius	IMRT H&N	98.1
Octavius	IMRT Prostate	98.9

3.3 Discussions

The construction of a MC model for a clinical linac was a time consuming iterative process. At the start of this project, the geometric parameters of the static part of the linac were not available, and as an alternative, a methodology for manipulation of existing validated PhSp files was developed as a side project (Chapter 4). Once a NDA between Elekta Limited, Klinikum Grosshadern and the Department of Medical Physics from LMU was signed, the linac head could be modeled following detailed geometric information provided by the vendor.

Even with the geometric information available, the incident electron beam parameters can vary from the nominal data provided by the vendor and among linacs of the same model, being therefore machine-specific [125, 131, 132, 144, 154–157]. Since there is no final consense in the literature regarding the best procedure for the electron beam optimization, the initial idea was to use PDD curves of several squared fields to tune the beam energy distribution ($\overline{E}$ and $FWHM_E$) and the lateral profiles for focal spot optimization.

The energy optimization using only PDD curves resulted in different sub-optimal combinations of $\overline{E}$ and $FWHM_E$, and their influence on the lateral profiles was only observed once the focal spot optimization process started. Therefore, it was necessary to reevaluate the values of $\overline{E}$ and $FWHM_E$ before a final optimal set of energy parameters could be determined. Although very useful, the PDDs should not be used alone for energy optimization, as the energy distribution of the electron beam, and consequently of the resultant photon beam, strongly influences the lateral profiles of larger fields, as reported by Sheikh-Bagheri et al. [154] and Chibani et al. [159] and observed in Fig. 3.22. Hence, it is recommended to perform the energy optimization using not only PDDs of several field sizes, but also at least one set of inline and crossline profiles for a larger field (e.g. 30×30 cm^2).

Different beam types and sizes have been investigated for the optimization of the focal spot, with the best results obtained with elliptical focal spots, in agreement with results reported elsewhere [126, 166]. An elliptical focal spot can be created in different ways, for example using a point source with asymmetric angular spreads in the inline and crossline directions, or using elliptical sources. PhSp 8 and 11 were created with point sources (Beam N. 8 and 11 in Table 3.10, respectively), with different angular spreads in the crossline direction Φ_{cr}. The influence of the different Φ_{cr} in the penumbra region was minimal, and both simulated crossline profiles were steeper than the reference measurements (Fig. 3.24). When using an elliptical source, however, the simulated penumbra region became less steep, resulting in a better agreement with the measurements (Beam N. 12 from Table 3.10). The penumbra of lateral profiles is therefore sensitive not only to the angular spread of the beam, but to its spatial distribution as well.

The differences between nominal and optimal values were not negligible. Consequently, even with access to confidential information from the vendor, the electron beam parameters are machine-dependent and need to be obtained for each linac individually through an iterative optimization process. With respect to the results obtained in this work, some steps could be spared, as for example using only elliptical beams rather than investigating different beam shapes. Finally, lateral profiles of different squared fields and at two different depths in water were sufficient for the focal spot optimization.

With the final set of electron beam parameters, 500 PhSp files were generated each using 10^8 initial electron histories. The files were created between the flattening filter and the monitor chambers (Fig.3.1), such that the dose deposited at the monitor chambers could be recorded during a simulation. This was necessary to evaluate the amount of backscatter radiation coming from the MLC system into the chambers, and consequently the dependence of the dose recorded by them on the simulated field size: if a correlation was observed, an extra correction factor should be included for the absolute dose calibration [132, 183]. The backscatter plate present in the Synergy head efficiently blocks the scattered radiation, and the observed variation due to different field sizes was less than 0.3% (Fig. 3.30 and Table 3.12), similar to results reported elsewhere [172, 173]. Therefore, in this study, the dose at the monitor chambers was considered independent of the field size. Moreover, as the monitor chambers and the mylar mirror do not influence the transmission of the photon beam [126], they could be introduced in the subsequent simulations as part of the patient-dependent model, without relevant discrepancy in the dose results and simulation time. Hence, a correction factor for backscatter radiation could be introduced in the future, if the user wants to make an even more precise absolute dose calibration.

Even though the geometric parameters of the Agility^{TM} MLC have become available on the course of this project, the patient-dependent model could be entirely constructed using information available in the literature (Table 3.2), which is a great advantage in case the information is not disclosured by the vendors. Some parameters, however, are machine-specific and needed to be tuned to match measured data [141], and the choice of which parameters to tune was made once transmission simulations started. It was clear that the MLC should be shifted towards a virtual focus [93, 137], but at the time the MLC started to be modeled, no information was found about the magnitude of this shift (α). After evaluating several α values, the simulated transmissions with α of 3.0 mm and 3.5 mm resembled the measurements in shape, however not in magnitude (Fig. 3.27): the simulated transmission was clearly higher than the measured one. Since the transmission is measured with closed leaves, the modeled leaves should become more attenuating, in order to reduce the simulated transmission to the correct level. The attenuating parameters that could be tuned were the density of the tungsten alloy ρ, used to construct the leaves, and the height of the leaves h_{leaf}. Hence, to reduce the simulated transmission, such parameters were optimized towards higher ρ and/or h_{leaf} combinations, until the best results were obtained with $\rho = 18.4\ gm/cm^3$ and $h_{leaf} = 95.0\ mm$. Both values of α resulted in a good agreement between simulation and measurements, but ultimately a shift value of $\alpha = 3.25\ mm$ was investigated and selected as the optimal value, to be consistent with more recently published results [138]. The optimal values for the leaf model varied significantly from the nominal values, confirming that these parameters need to be adjusted to the machine being modeled, even when geometric information has been provided by the vendor.

The optimization for both the incident electron beam and leaf parameters are time consuming iterative processes, involving the generation of multiple PhSp files and the simulation of many dose distributions. Even more iterations could be performed resulting in potentially better values. In this work, once a satisfactory visual agreement between simulations and measurements (PDDs and lateral profiles for squared fields) was achieved, the optimization process stopped. Ultimately, the MC user wants to describe a specific machine. Therefore, it is up to the user to tune such parameters accordingly and to decide how much time it can be dedicated into this iterative process, based on the desired accuracy.

Once both the electron beam and the leaf parameters were optimized, several validation simulations were performed and compared to benchmark dose distributions. First, PDDs and lateral profiles at 15 mm and 100 mm depth in water were simulated and compared to measurements. This analysis was performed by visual inspection using relative normalized dose values. For the PDDs, the dose was normalized to the dose at D_{max} (1.6 cm for the 6 MV photon beam modeled); for the lateral profiles, the normalization value was the dose at the CAX point of each curve. As can be observed in Fig 3.31, the simulated PDD curves are in reasonably good agreement with the respective measurements throughout the entire depth evaluated. For the lateral profiles, at both 15 mm and 100 mm depth, the simulated curves match the measurements both at the penumbra region and at the off-axis positions (shoulders), as observed in Figs. 3.32 and 3.33 respectively. This is an important result, as it is a strong indication that both the energy of the electron beam and its spatial and angular distributions (focal spot) have been properly modeled. The good agreement between simulated and measured lateral profiles also indicates that the field size was properly modeled. In other words, the field-size related positioning of each component of the MLC was successfully described by the analytic equations Eqs. 3.51, 3.49, 3.59 and 3.60. These equations are general and could be used to model any Agility$^{\text{TM}}$ MLC. Moreover, the method used to obtain the analytic equations is also general and could be applied to any MLC system upon system-specific adaptation.

Up to this point, the validation of the model was performed only by visual inspection between normalized simulations and reference measurements. Subsequent validations of 3D dose distributions inside the Octavius phantom were performed with absolute dose values, and the agreement between the simulations and the reference TPS calculations was quantified using the gamma evaluation method, with (3%, 3 mm) criteria. Four different squared fields and two clinical IMRT plans were analyzed, with promising gamma passing rates for all cases (Table 3.14). The H&N plan had a passing rate of 98.1%, while the value for the prostate plan was 98.9%, both above the recommended passing rate of 95% for clinical IMRT plans with (3%, 3 mm) criteria [184, 185]

The H&N and prostate IMRT plans were composed of 92 and 61 different control points (i.e. different segments), respectively, distributed over several different gantry angles (Tables 3.7 and 3.8). As expected for this type of treatment plan, they were comprised of several small fields with complex shapes, as opposite to the squared fields. Since the gamma evaluation method takes into consideration both the absolute difference in dose and discrepancies in its spatial distribution, the wrong positioning of one single leaf, a wrong photon beam energy spectrum and/or the wrong conversion into absolute dose values could lead to reduced gamma passing rates. The high passing rates observed for both IMRT plans are therefore a strong validation of both MC models of the Elekta Synergy linac and Agility MLC used. Ultimately, the validations performed with the IMRT plans can be considered the most important and comprehensive validation step: it proves that the combined MC models of the linac properly and faithfully represents the equipment used, including the electron beam parameters, the static components of the linac head, the components of the MLC system, the analytic equations used for leaves and diaphragms positioning and the absolute dose calibration. Moreover, the model properly uses the retrieved data from the RTPlan.dcm files for dose simulation (gantry angle, leaves and diaphragms position) and for absolute dose conversion (MU values of each control point).

3.4 Conclusions

A validated MC model of the Elekta Synergy linac, coupled with an Elekta Agility MLC system was developed. The model is divided in two parts: the patient-independent part, comprising the static components of the linac head, and the patient-dependent part, comprising the beam shaping components, i.e. the MLC collimator system. The first was produced based on geometric information disclosed by the vendor, while the latter was modeled using only information available in the literature. Despite the information available, either provided by the vendor or found in literature, both models had parameters that are machine-specific and needed to be fine-tuned to represent the equipment used. An improved description of the incident electron beam in the patient-independent part enabled the generation of PhSp files, which more correctly represented the characteristics of the 6 MV photon beam produced by the static components of the Elekta Synergy linac used in this work. Such PhSp files served as primary particle sources to the patient-dependent model. Ultimately, the parameters gathered from literature together with the analytical models for leaves and diaphragms positioning, can be used to produce a reliable equipment-specific geometry of MLCs for MC simulations. In combination with the PhSp manipulation method described in Chapter 4, an entire MC model of a linac could potentially be produced without depending on vendors' confidential information, which may not be easily accessible.

The MC models were extensively validated using several 2D and 3D dose distributions, both in a water phantom and in a water-equivalent cylindrical phantom. Even though for several validation steps the agreement between simulations and references was based only on visual inspection, for the final validations using 3D dose distributions of complex IMRT plans the agreement was quantified in terms of gamma passing rates. The high gamma passing rates obtained are a strong indicative of the reliability of the produced PhSp files and the MLC MC model.

The final purpose of the patient-dependent model was to simulate dose distributions inside patient-like geometries, constructed using patient CT data. Hence, the Geant4 model is able to read CT dicom images and use the stored information therein to construct each patient voxel inside the model. Together with a MATLAB script, which extracts the relevant information of RTPlan.dcm files from any TPS-calculated plan, the parameters of IMRT fields are passed to the model via macro files, thanks to user-defined commands, enabling parallel simulation of different fields at patient-like geometries. This was essential for the MC simulations required on Chapter 5.

Finally, the model was created only for the 6 MV photon beam, however, the same methodology can be extended to other beam energies and other linacs, provided the necessary changes are applied.

CHAPTER 4

Phase Space Optimization

At the start of this project, the geometric parameters the Elekta Synergy® linac were not available, and as an alternative, a methodology for manipulation of existing validated PhSp files was developed. The IAEA offers a database of fully validated PhSp files, available for general use [186], which can be implemented with different MC codes like EGSnrc [187], PENELOPE [188] and Geant4 [127]. However, currently the IAEA database provides PhSp validated only for a limited number of linac models. A validated PhSp for the Elekta Synergy® is not available at the IAEA database. As a first attempt, a PhSp validated for a similar equipment, the Elekta Precise®, was used as a surrogate for the static part of the linac head. Discrepancies as high as 11% and 26% were observed for Percentage Depth Dose (PDD) and lateral profiles in water, respectively, when compared to measurements [189]. These discrepancies can be attributed to differences in the linac's geometries and initial electron focal spots. The Elekta Precise® PhSp was generated using a circular-shaped electron beam [190–192], while better results for the Elekta Synergy® have been reported when using an elliptical electron beam [126, 166]. Moreover, the output of a linac simulation is very sensitive to other components of the head, such as the primary collimator and flattening filter [133].

To try overcome the lack of geometric information, the approach presented in this chapter consists in manipulating some information stored in the Elekta Precise IAEA PhSp for the 6 MV photon beam, in order to adapt it to better represent the 6 MV photon beam from an Elekta Synergy®. The IAEA PhSps are binary files containing information regarding particles' position, energy and direction, along with supplementary parameters that are not considered here. In this approach, the energy and directions of the particles were manipulated, as these parameters are directly correlated to the electron beam and focal spot. This approach consisted in a master thesis project supervised by the author of this work [193], and parts of this chapter were previously published in Physica Medica [194].

4.1 Methodology

The IAEA PhSp is a binary file, with information stored in an unknown sequence, making it troublesome to extract only selected parameters. To overcome this issue, a new PhSp plane was created at 27.9 mm below the original IAEA PhSp and scored only the relevant parameters of particles emerging from the IAEA PhSp, in a specific order, into a new PhSp here referred to as LMU PhSp.

The optimization was an iterative process, in which the energy and direction of particles stored in the LMU PhSp were randomly changed. For each iteration, a new PhSp was created and used to simulate PDD or lateral profiles in a virtual water phantom for different squared fields. The simulated profiles were then compared to measurements performed under the same conditions through cost functions. The next iteration would then change the energy or direction in the direction of the decreasing cost. The definition of cost functions was essential for the optimization process, serving as a guidance for the manipulation steps, as a great number of iterations was needed and qualitative or visual evaluations would have been unfeasible.

4.1.1 Measurements and simulations

All measurements were performed in the Elekta Synergy® linac equipped with an Elekta Agility® MLC installed at the Klinikum Großhadern. PDD and lateral profiles (at 15 and 100 mm depth) for 2×2, 5×5, 10×10 and 20×20 cm^2 fields were measured in the water phantom with a microDiamond detector, at SSD of 900 mm for the 6 MV photon beam. Further details on the measurements can be found in Chapter 3. The lateral profiles at 15 mm depth were used for optimization, while validation was performed for both depths.

Simulations for PDD and lateral profiles were performed under the same conditions in a virtual water phantom, using an in-house developed MC model of the Agility MLC (Chapter 3). The static part of the linac head was replaced by the perturbed PhSps. Particle recycling was used to decrease the statistical noise, with recycling factors ranging from 14 to 24 times, depending on the simulation configuration: for cropped PhSp (section 4.1.2) particles were recycled 24 times due to the reduced number of particles inside the PhSp files; for simulations performed with full PhSp, particles were recycled 14 and 19 times for smaller and larger fields, respectively.

For both PDD and lateral profiles, the dose was recorded using the command-based scoring feature from Geant4, hereon referred to as scoring meshes. Column-like scoring meshes were used for PDD simulations, with 1-mm resolution in the central axis direction, to account for the dose deposits at different depths in water. Two-dimensional (2D) scoring meshes were used for the lateral profiles simulations, with 1 mm resolution in the inline (gun-target) and crossline (lateral) directions. The size of the 2D mesh was determined such that it scores the entire penumbra region, and is therefore dependent on the field size.

Simulation uncertainties were calculated using the batch-approach [175]. Different from the history-by-history method used in Chapter 3, in the batch-approach the total number of incident electron histories is not used on the calculation of the mean and standard deviation values of the dose deposited in each voxel. Instead, these values are calculated using the results obtained from independent batches (or instances) of simulations performed using the same configuration. The mean dose deposited in a voxel i ($\overline{D_i}$) is simply the average over all instances (Eq. 4.1), where N is the number of instances, and the standard deviation $\sigma_{\overline{D_i}}$ is given by Eq. 4.2. Only voxels scoring at least 10% of the maximum dose were considered for the uncertainty estimation, as in low dose areas (e.g. profile tails) fewer interactions take place and a non-representative increase in the uncertainty is observed. For all simulations performed in this chapter, the average uncertainty was 2%, with a maximum value of 3%. All simulations were performed using Geant4 v.10.01.p01.

$$\overline{D_i} = \sum_{j=1}^{N} \frac{D_j}{N} \tag{4.1}$$

$$\sigma_{\overline{D_i}} = \sqrt{\sum_{j=1}^{N} \frac{(D_j - \overline{D_i})^2}{N(N-1)}} \tag{4.2}$$

4.1.2 Cropping the Phase Space

To reduce simulation time and improve the efficiency of the iteration steps, the manipulated LMU PhSps were radially cropped into smaller PhSps. In this process, a new PhSp was created, containing the information of particles within a specific radial distance from the center of the original PhSp, here called cropping radius (Fig. 4.1). To ensure that no relevant information was disregarded by cropping, the size of the smallest appropriated cropping radius was determined by the position of the particles inside the original LMU PhSp, the dimension and spatial distribution of the scoring mesh and the field size. Moreover, the influence of the cropping on resultant simulations was assessed by comparing simulations performed with the full LMU PhSp versus cropped LMU PhSps.

For PDD simulations, the column-like scoring mesh was positioned in the center of the radiation field, and is therefore mostly, though not exclusively, influenced by the particles emerging from the central region of the PhSp. Moreover, the energy optimization was based on PDD profiles from a 10×10 cm^2 field, therefore determining the appropriate cropping radius for this configuration was essential.

New perturbed PhSps were created, with energy perturbation of 3 MeV applied separately to the particles within different radially symmetric regions, namely 0-5mm, 5-10 mm, 10-15 mm, 15-20 mm, 20-40 mm and 40-60 mm. These PhSps were used to simulate PDD profiles for the 10×10 cm^2 field, and the results were compared to the curve simulated with the original IAEA PhSp via the cost values (section 4.1.3). The

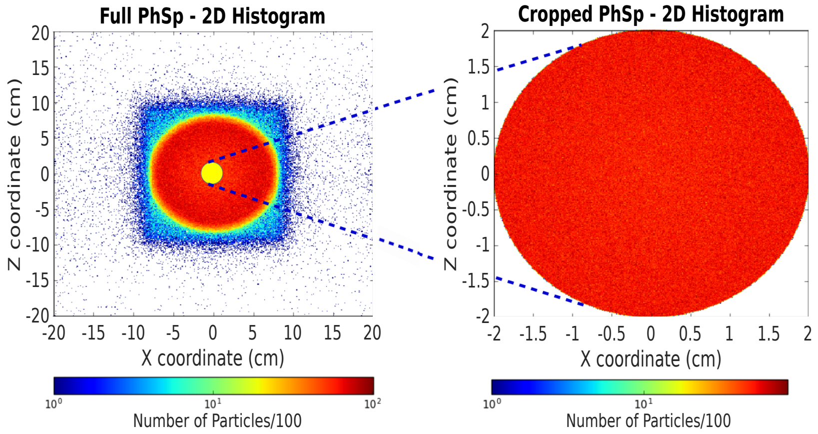

Figure 4.1: Cropping of the original PhSp (left) into a radially smaller PhSp (right). In this example, the cropping radius is 20 mm, which reduces the simulation time by a factor of 20. Reprinted with permission from [194].

lower the cost value, the weaker the influence of the particles in the perturbed region to the outputted PDD profile. Therefore, the cropping radius for PDD simulations was chosen as the value beyond which the particles' perturbation did not significantly influence the simulation results.

To determine the effective cropping radii for lateral profiles simulations, a simple geometric magnification was used (Fig. 4.2). The IAEA PhSp was recorded at 272.10 mm below the top of the Bremsstrahlung target (X-ray source), and the LMU PhSp was generated on a plane 27.9 mm below it, thus 300.0 mm away from the source. The aperture of the collimator (BC), beyond which the radiation is blocked for a certain field size (OD) at isocenter, is given by Eq. 4.3. The distance from the source to the isocenter AO is 1000 mm, while AB is the distance from the target to the LMU PhSp. BC marks the radial region within which the particles strongly influence the lateral profiles, for an $OD \times OD$ cm^2 field, providing a good first guess for the minimum acceptable cropping radius for this field size. The impact of particles emerging beyong this region drops substantially, however the effective cropping radius was always taken as a higher value, often twice the minimum cropping radius value, as a conservative measure.

$$BC = OD.\frac{AB}{AO} = OD.\frac{300\ mm}{1000\ mm} \tag{4.3}$$

4.1.3 PDD cost and energy optimization

As the PDD profile is highly correlated to the photon beam energy [85], the cost between simulated and measured PDDs was used as a benchmark for energy optimization. To account for the different sampling spacing, the simulated and measured PDD curves were first normalized to their respective maximum values, resulting in relative dose values. The normalized PDD curves were then fitted to double exponential functions (Eq. 4.4), where

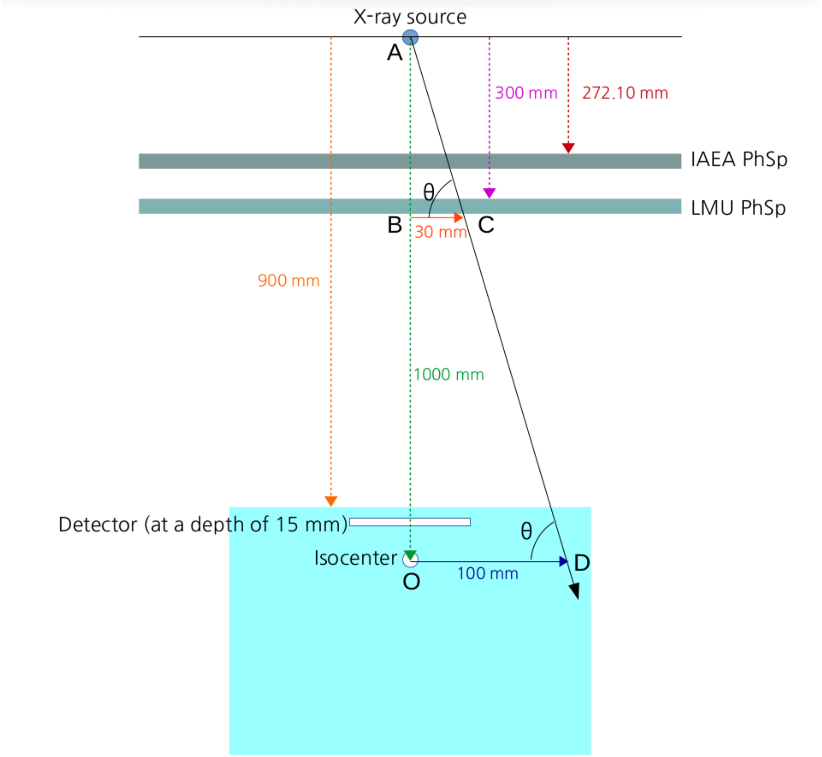

Figure 4.2: Schematic representation of the setup for lateral profile simulations and measurements. The distance from the X-ray source to isocenter (AO) is 1000 mm, with source-to-surface distance of 900 mm. The LMU PhSp is placed at 300 mm away from the source (AB). The 2D scoring mesh for lateral profiles (in case of simulations) or the microDiamond detector (for measurements) are placed at 15 mm or 100 mm depth in the water phantom. (Not drawn to scale). Reprinted with permission from [194].

$D(d)$ is the dose at depth d and a_1, a_2, b_1, b_2 and c are the fitting parameters. The domain $d = [0, 300]$ mm was used to fit both curves. For the 10×10 cm^2 field, figure 4.3 shows the fitted curves for measured and simulated data. The adjusted R^2 values were greater than 0.996 and the Root Mean Square Errors (RMSEs) were lower than 0.016 for PDD fitted from all field sizes.

$$D(d) = a_1 e^{-b_1 d} - a_2 e^{-b_2 d} + cd \tag{4.4}$$

The difference between the measured and simulated fitted curves, referred to as cost value, can be quantified and reflects the disagreement between the data. The cost value was defined by a χ^2 function (Eq. 4.5), where n is the number of points used in the calculation of the dose, D_i^m and D_i^s are the doses at point i, for the fitted curves of the measured and simulated PDD, respectively. The χ^2 function is simple, reasonably insensitive to statistical uncertainty [125] and able to reflect the disagreement between measured and simulated PDD, being a good metric choice

$$\chi^2 = \frac{1}{n} \sum_{i=1}^{n} (D_i^m - D_i^s)^2 \tag{4.5}$$

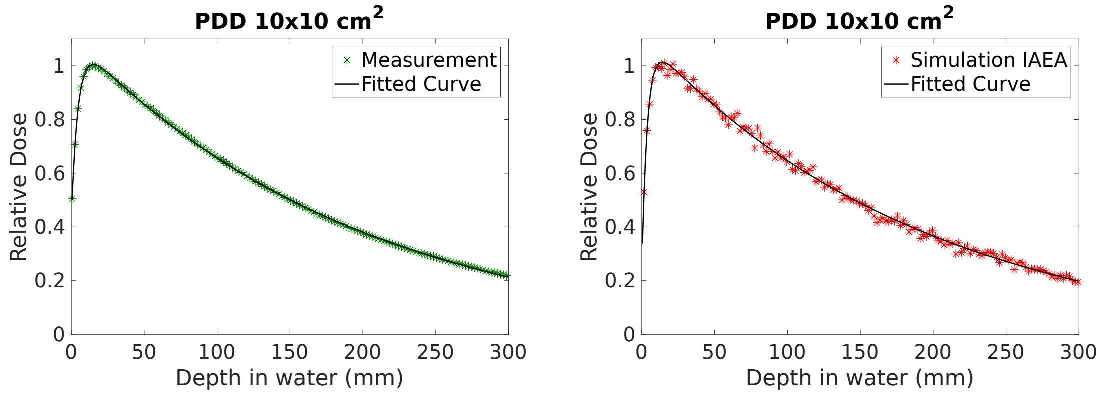

(a) Measurements

(b) Simulations (IAEA PhSp)

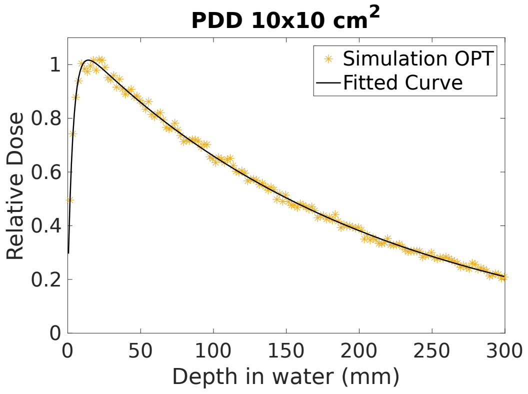

(c) Simulations (optimized PhSp)

Figure 4.3: PDD profiles with fitted curves for the 10×10 cm^2 field. (a) Measurements, (b) simulations performed with the original IAEA PhSp and (c) simulations performed with optimized PhSp. For all three data sets and all fields analyzed, the adjusted R^2 values are higher than 0.996; RMSEs are smaller than 0.015. Reprinted with permission from [194].

In the PDDs simulated using the Elekta Precise® IAEA PhSp, the dose values obtained in the region of Transient Charged Particle Equilibrium were lower than the measured values [189], suggesting that the energy of electron beam used to generate the IAEA PhSp (5.75 MeV [190]) was lower than the electron beam energy of the Elekta Synergy® used for measurements. Additionally, several studies have reported higher electron beam energies when modeling the Elekta Synergy® linac [166][195]. Therefore, the optimization of the energy of particles stored in the LMU PhSp should involve an increase in value.

The manipulation of energy was performed as an increase of the energies of all particles stored inside the LMU PhSp, except for particles with energy of approximately 0.511 MeV (annihilation peak), which should not be affected. As an example, a perturbation of 0.2 MeV to the PhSp means that 0.2 MeV is added to the energy of all suited particles (Fig. 4.4). Such manipulation shifts the mean and median of the energy spectrum by the same amount.

To determine the optimal increasing factor needed for the energies, several cropped PhSp were created, with cropping radius of 20 mm. The energies of particles inside each cropped PhSp were increased by 0.1, 0.2, 0.3 and 0.4 MeV, and each cropped-perturbed

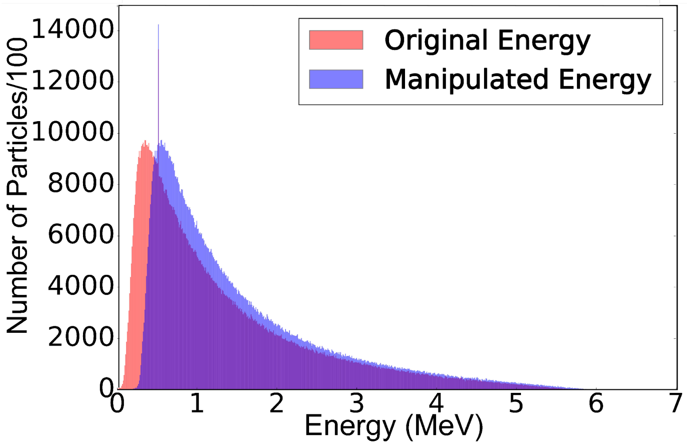

Figure 4.4: Energy manipulation: Energy spectra of the PhSp with original energies (red) and perturbed energies (blue). An increase of 0.2 MeV was applied to the original PhSp, meaning the energy of all particles, except the annihilation photons, was increased by 0.2 MeV. Reprinted with permission from [194].

PhSp was used for PDD simulations of the 10×10 cm^2 field. The results were compared to the measured curve and a local minimum of the cost value was found for simulations with the 0.3 and 0.4 perturbed PhSps. Subsequently, new cropped PhSps were perturbed with 0.32, 0.34, 0.36 and 0.38 MeV and used for PDD simulations, which were once again compared to measurement data. Once the optimal perturbation value was obtained, i.e. the value which yield a lower cost between simulation and measurement, it was applied to the entire LMU PhSp. The full perturbed PhSp was used to simulate PDD profiles for all field sizes in order to validate the manipulation. The influence of the energy manipulation on the lateral profiles was also assessed by simulating inline and crossline profiles at 15 mm depth in water, for all field sizes, with the full PhSp with optimized energy. The same process can be applied to decreases of energy, depending on the demands of the user.

Energy optimization: step-wise methodology

Summarizing, as part of the entire PhSp optimization methodology, the energy optimization can be described in 6 simple steps, listed below:

1. Apply large energy perturbations to all particles within different radially symmetric regions of the original PhSp, and simulate PDD profiles. Find the acceptable cropping radius (section 4.1.2).

2. Increase/decrease the energy of all particles (except annihilation photons) in the cropped PhSp in steps of 0.1 MeV, and simulate corresponding PDD profiles for a 10×10 cm^2 field. Calculate the cost values for each simulation.

3. Identify the two energy perturbations that yield the lowest cost values, and define small energy steps of 0.02 MeV between them. Perturb the energy of the cropped PhSp by the smaller energy steps and simulate PDD profiles for a 10×10 cm^2 field. Calculate the cost values.

4. Identify the perturbation yielding the lowest cost and use the respective cropped perturbed PhSp to simulate PDD for several squared fields. This value will be referred to as the optimal energy perturbation factor. This step is intended to avoid a perturbation that could potentially work for only one field size, i.e. result in a small cost value for the 10×10 cm^2 field, whilst producing high cost values for other field sizes.

5. The user can perform steps 3 and 4 multiple times, using even smaller perturbations, to achieve costs values of a desired accuracy.

6. Apply the optimal energy perturbation factor to the original PhSp and validate for all field sizes, using the full PhSp to simulate PDD and lateral profiles.

4.1.4 Penumbra cost and direction optimization

For the direction optimization step, the energy-optimized PhSp obtained from the method described in section 4.1.3 was used as the starting point.

The angular distribution of the particles, described in terms of their direction cosines P_x, P_y and P_z, is correlated to the focal spot and affects directly the lateral profiles [126, 156, 164, 196], particularly at the penumbra region. The penumbra is defined as the distance between the 20% and 80% dose values of the profile, normalized to the dose at central axis [85]. Hence, the cost values between measured and simulated penumbra of lateral profiles were used as a benchmark for particle's direction optimization. The cosine directions are the parameter that can be recorded in a PhSp file, and will be referred to as particle's direction for conciseness.

Similar to the PDD profiles, both measured and simulated lateral profiles were fitted to account for the different sampling spacing. As the penumbra is the most relevant parameter in this part of the optimization, the fitting function used for the lateral profiles needs to properly represent this region. The sigmoid function (Eq. 4.6) was used for fitting the curves, where k and x_0 are the slope and position of the center of the sigmoid, respectively, and $D(x)$ is the dose at position x. The slope and center of the sigmoid overlap the penumbra region of the lateral profile, as can be seen in Fig. 4.5 for simulated and measured crossline profiles at 15 mm depth in water, for the 10×10 cm^2 field.

$$D(x) = \frac{1}{1 + e^{-k(x-x_0)}} \tag{4.6}$$

Even though the sigmoid function is unable to properly fit the tail of the profiles, the remaining parts are well represented, especially the penumbra. Furthermore, the sigmoid fitting parameters k and x_0 are insensitive to simulation noise, as well as to the unfitted points at the low dose region (points with less than 20% of the maximum dose), where variations of up to ± 20% on the dose values resulted in less than 2% and 0.05% absolute fluctuations to the values of k and x_0, respectively. Hence, the sigmoid function provides

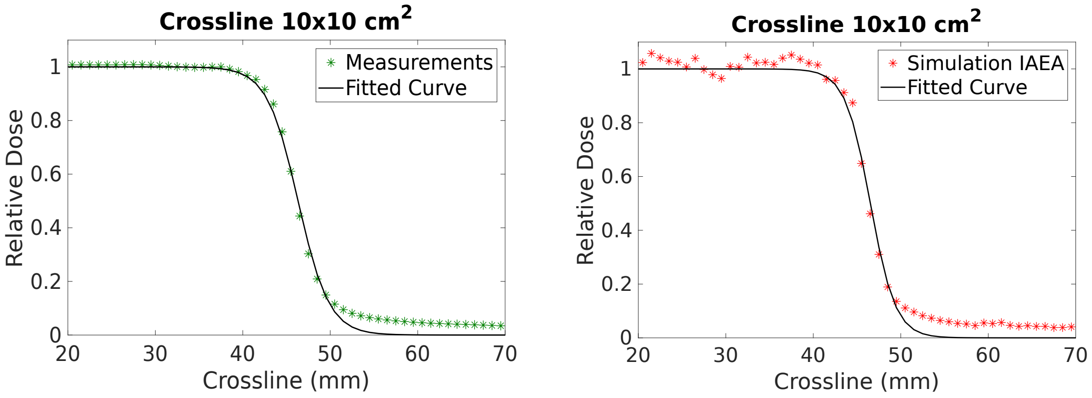

(a) Measurements **(b)** Simulations (IAEA PhSp)

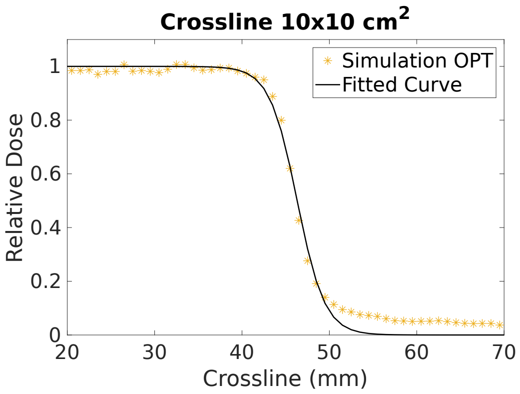

(c) Simulations (optimized PhSp)

Figure 4.5: Crossline profiles with fitted curves for the 10×10 cm^2 field. (a) Measurements, (b) simulations performed with the original IAEA PhSp and (c) simulations performed with optimized PhSp. For all three data sets and all fields analyzed, the adjusted R^2 values are higher than 0.988; RMSEs are smaller than 0.040. Reprinted with permission from [194].

reliable information on the penumbra region, and can be used for profiles at both crossline and inline directions, sampled at both 15 mm and 100 mm depth in water. For all data sets, the adjusted R^2 values are greater than 0.988 and RMSEs are smaller than 0.040.

The cost values k_{cost} and x_{0cost} were calculated as the percentage difference between the fitting parameters k and x_0 of measured profiles, and those of the simulated profiles, k' and x_0' (Eqs. 4.7 and 4.8). The k_{cost} and x_{0cost} exhibit the advantage over the χ^2 value that they provide information on the position and slope of the simulated penumbra, relative to the measured data, while the χ^2 only gives an absolute value of disagreement. Additionally, the sign of k_{cost} indicates whether the simulated profiles have steeper or shallower slopes in the penumbra area, serving as a guidance for manipulation.

$$k_{cost} = \frac{k - k'}{k}.100 \tag{4.7}$$

$$x_{0cost} = \frac{x_0 - x_0'}{x_0}.100 \tag{4.8}$$

Finally, to assess the overall improvement or worsening caused by the manipulations, each profile is fitted for both sides, and the average of the absolute values from k_{cost} and $x0_{cost}$ is considered. Additional gamma evaluation [179] was performed for both PDD and lateral profiles using the (3%, 3 mm) criteria, as adopted by other authors for quality evaluation of simulated PDD and lateral profiles [191].

The angular momentum of one particle can be thought of as the divergence of the particle's travel direction with respect to the beam traveling direction, as described by Eq. 4.9, with α representing the angle between the particle's and the beam's directions. Here, the beam traveling direction corresponds to the -Y axis, while the X axis is the crossline direction (patient's left-right in the HFS position) and Z is the inline direction (gun-target direction).

$$div = \tan\alpha = \sqrt{\frac{P_x^2 + P_z^2}{P_y^2}} \tag{4.9}$$

The divergence of a particle is linearly dependent on its radial position inside the PhSp, and increases with increasing distances from the PhSp center, as shown by the green line in Fig. 4.6 and described by Eq. 4.10. This linear relationship will be called beam divergence line.

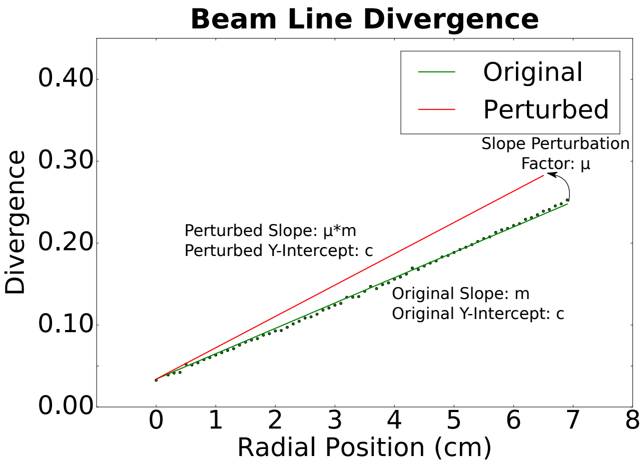

Figure 4.6: The slope of the original beam divergence line (green), is perturbed by a factor μ, transforming it into the perturbed beam divergence line (red). The intercept value c remains unchanged, to respect the constraint of low divergence of particles in the central region of the PhSp. Reprinted with permission from [194].

$$div(r) = mr + c, \tag{4.10}$$

The divergence of each individual particle describes the photon beam divergence macroscopically, which is a consequence of the virtual electron focal spot when traced back along the beam upstream. Hence, by changing the divergence of the individual particles in a controllable way, the macroscopic divergence of the entire beam can be steered, indirectly modifying the characteristics of the virtual focal spot.

Any changes applied to the particles' divergences must be consistent with the beam divergence line, i.e. the perturbed divergences must follow a linear relationship with their radial position. Thus, the only possible manipulation is to change the slope of the beam divergence line by a certain factor μ, described by Eqs. 4.11, 4.12 and 4.13, and represented by the red line in Fig. 4.6. Particles in the center of the beam have low divergence, and therefore the intersecting point c was kept unchanged. Hereon, the variables with the circumflex ($\hat{}$) represent the manipulated quantities.

To modify the slope of the beam divergence line, the divergence of all particles must be changed. Hence, the values of P_x, P_y and P_z stored in the original LMU PhSp need to be manipulated. The equations and constraints to be considered for this manipulation steps are listed from Eq. 4.11 to Eq. 4.15.

$$\hat{div}(r) = \hat{m}r + \hat{c}, \tag{4.11}$$

$$\hat{m} = \mu m, \tag{4.12}$$

$$\hat{c} = c, \tag{4.13}$$

$$\frac{\hat{P}_x}{\hat{P}_z} = \frac{P_x}{P_z}, \tag{4.14}$$

$$P_x^2 + P_y^2 + P_z^2 = \hat{P}_x^2 + \hat{P}_y^2 + \hat{P}_z^2 = 1. \tag{4.15}$$

The original beam divergence line is described by Eq. 4.10, while the modified beam divergence line is given by 4.11, with respect to the radial position r. The perturbations to be applied on the original beam divergence line are given by Eqs. 4.12 and 4.13, where Eq. 4.12 changes the slope of the line by a perturbation factor μ, with Eq. 4.13 keeping the divergence of the particles in the central of the PhSp region unaffected. To prevent the manipulated beam from curling around the central axis, Eq. 4.14 needs to be satisfied. Finally, Eq. 4.15 is the normalization condition of direction cosines and needs to be respected.

To identify the equations defining the changes in P_x, P_y and P_z, given a slope perturbation factor μ, Eqs. 4.10 and 4.11 are rearranged into Eqs. 4.16 and 4.17, and replaced into Eq. 4.12 to obtain Eq. 4.18.

$$m = \frac{div - c}{r}, \tag{4.16}$$

$$\hat{m} = \frac{\hat{div} - c}{r}, \tag{4.17}$$

$$\hat{div} = \mu(div - c) + c, \tag{4.18}$$

Inputting the original definition of divergence (Eq.4.9) into Eq. 4.18 leads to

$$\frac{\hat{P_x}^2 + \hat{P_z}^2}{\hat{P_y}^2} = \left((\mu(\sqrt{\frac{{P_x}^2 + {P_z}^2}{{P_y}^2}} - c) + c \right)^2, \tag{4.19}$$

With Eqs. 4.15, 4.19 and some mathematical manipulation one can obtain Eq. 4.20:

$$\frac{1}{\hat{P_y}^2} = \left((\mu(\sqrt{\frac{{P_x}^2 + {P_z}^2}{{P_y}^2}} - c) + c \right)^2 + 1, \tag{4.20}$$

By defining γ as

$$\gamma = \left((\mu(\sqrt{\frac{{P_x}^2 + {P_z}^2}{{P_y}^2}} - c) + c \right)^2 + 1, \tag{4.21}$$

The transformation needed for the direction cosine P_y is given by Eq. 4.22:

$$\hat{P_y} = \frac{1}{\gamma} \tag{4.22}$$

Defining ω as Eq. 4.23 and replacing γ and Eq. 4.22 into Eq. 4.19, the equation describing the transformation on P_z is obtained as Eq. 4.24:

$$\omega = \frac{P_x}{P_z}, \tag{4.23}$$

$$\hat{P_z}^2 = \frac{1}{1+\omega^2}\left(1 - \frac{1}{\gamma}\right), \tag{4.24}$$

Finally, the transformation for P_x is obtained by using Eqs. 4.24, 4.19 and ω:

$$\hat{P_x}^2 = \frac{\omega^2}{1+\omega^2}\left(1 - \frac{1}{\gamma}\right), \tag{4.25}$$

4.1.4.1 Sigmoid fitting parameters k and x_0 versus μ

With the transformations defined by Eqs. 4.22, 4.24 and 4.25, the influence of μ into the resultant lateral profiles needs to be analyzed. This can be assessed by the parameters k and x_0 of fitted profiles simulated using PhSps with modified particles' directions. For this task, it was necessary to create a great number of manipulated PhSps using different μ values, and use these PhSps to simulate lateral profiles.

Assuming that a change in the focal spot influences the simulated lateral profiles in a similar way, regardless of field size, measuring direction and depth in water, one can assume that the effect of μ on the lateral profiles is also independent from these factors. To optimize this process, which would be otherwise unfeasible in a reasonable amount of time, only the 2×2 cm^2 field was chosen. For this field, a cropping radius of 5 mm can be used to simulate lateral profiles (Eq. 4.3). Therefore, 800 5-mm cropped PhSp files were perturbed, with μ ranging from 0.001% ($\mu = 1.00001$) to 0.8% ($\mu = 1.008$), in steps of 0.001% ($\Delta\mu = 0.00001$). These PhSps were used to simulate inline profiles at 15 mm depth in a virtual water phantom, for the 2×2 cm^2 field.

The simulated profiles were fitted with the sigmoid function (Eq. 4.6). A negative linear correlation between the fitting parameters k and x_0 and the manipulation factor μ is observed, with a linear fit properly representing the scattered data (Fig. 4.7). This implies that an increase in the slope of the beam divergence line leads to a decrease in the sigmoid fitting parameters: i.e. the more divergent the beam is, the broader is the lateral profile. With this correlation identified, the fitting parameters of lateral profiles could be steered in the direction of decreasing costs, using appropriated μ values.

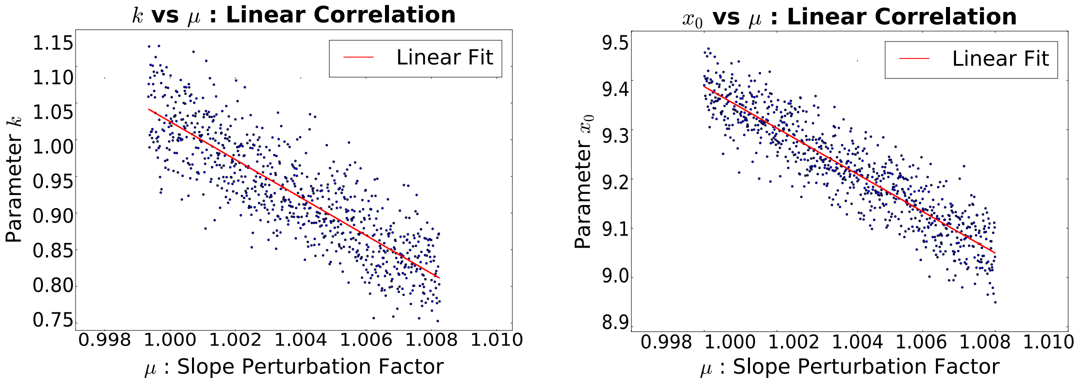

(a) k vs μ

(b) x_0 vs μ

Figure 4.7: Linear correlation between the manipulation factor μ and fitting parameters (a) k and (b) x_0. Reprinted with permission from [194].

Differences of up to 15% were observed on the sigmoid fitting parameters for multiple simulations with the same configuration. This statistical uncertainty is intrinsic from Monte Carlo methods [107], and arises due to the finite number of particles stored inside a PhSp, especially when cropped PhSps are used. To account for this fluctuation, 50 simulations of the same configuration were performed and the resultant profiles were fitted. Simulations with the same configuration will be hereon referred to as instances. The medians of k and x_0 were calculated using sets with different number of simulations, ranging from 1 to 50, in order to determine the minimum number of instances needed to obtain stable median values. The median is an adequate quantity, as its value is less sensitive to the presence of outliers. The median values become stable when 20 or more simulations are considered (Fig. 4.8). From hereon in this study, 24 instances of simulations for the same configuration were performed, to ensure a stable result and a reasonable trade-off with the computational time.

Direction optimization: step-wise methodology

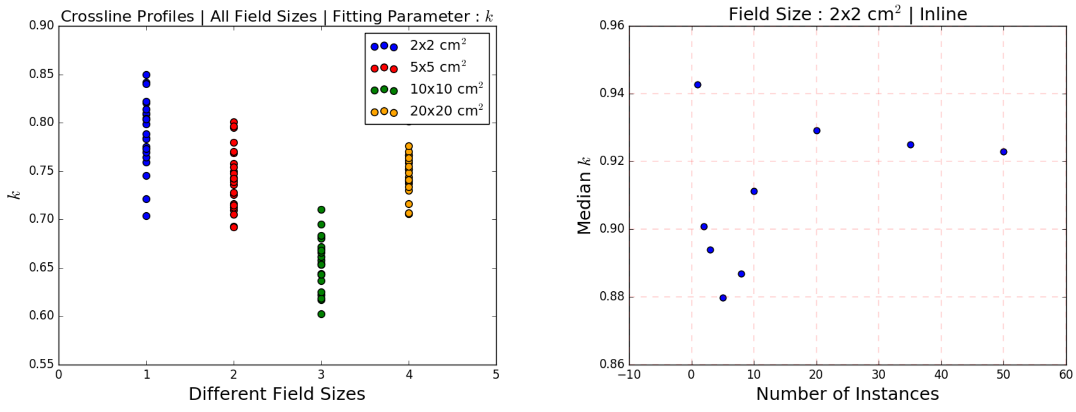

(a) **(b)**

Figure 4.8: (a) Different fitting parameters k, obtained by different simulations with same configuration and (b) the median value of k for the 2×2 cm^2 for different data set sizes.

Following the determination of the linear relationship between the sigmoid parameters and μ, the manipulation of particles' directions can be described by the steps listed below. While the 2×2 cm^2 field was used to determine the relationship between the fitting parameters and perturbations to μ, the 10×10 cm^2 field was used as the reference field for the optimization process, with a cropping radius of 50 mm.

1. The energy-optimized PhSp was cropped to an effective cropping radius (section 4.1.2).
2. A second cropped PhSp was generated and perturbed by a relatively large factor ($\mu = 0.8\%$) following Eqs. 4.22, 4.24 and 4.25.
3. 24 instances of lateral profiles were simulated at 15 mm depth on the virtual water phantom, for each PhSp.
4. All simulated profiles were fitted to sigmoid functions, and the medians of the resulting k and x_0 fitting parameters were calculated. A linear fit was applied to the two medians, as a function of μ (red line, Fig. 4.9).
5. The measured profile for the respective field size was fitted into a sigmoid curve and the parameters k and x_0 were extracted (green line in Fig. 4.9).
6. The intersections between the curves (green dots) represent the optimal perturbation factors obtained from k and x_0 (Figs. 4.9a and 4.9b respectively), i.e. the optimal values of μ that would result in $k' \approx k$ and $x_0' \approx x$, for the 10×10 cm^2 field.
7. Steps (1-6) were repeated for 2×2, 5×5 and 20×20 cm^2 fields, yielding different perturbation factors specific to every field.
8. The resulting perturbation factors were calculated as the average of the individual factors from each field, considering the results obtained using the k parameter only.

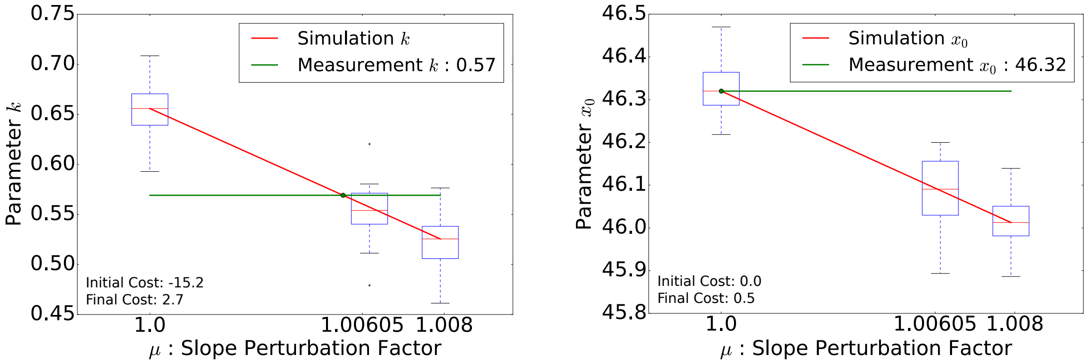

(a) **(b)**

Figure 4.9: Finding the optimal perturbation factor from (a) k and (b) x_0. The green lines represent the fitting parameters of the measured crossline profile. The intersection of both green and red lines (green dot) gives the optimal perturbation factor in the crossline direction for the 10×10 cm^2 field. The optimized fitting parameters are represented by the boxplots in the middle. The final cost values are given in the bottom of the plots. Reprinted with permission from [194].

This methodology results in two different optimal perturbation factors, one obtained from the k parameter and another one obtained from x_0, which were one order of magnitude different from each other. However, the k parameter is more significantly influenced by changes in the beam divergence than the x_0 parameter: the k value decreased by $\approx 20\%$ after a perturbation of $\mu = 0.8\%$, while x_0 decreased by only 0.75%. Using just the perturbation factor given by k greatly decreases the absolute k cost (from 15.2% to 2.7%), but slightly increases the x_0 cost (by $\approx 0.5\%$), as seen in Fig. 4.9. However, this increase results in a shift of less than 0.3 mm in the profile penumbra (Fig. 4.9b). This behavior was observed for all fields analyzed. Hence, only the perturbation factor obtained from k was considered in the next steps of this methodology. It was ensured that the final cost values for x_0 remained within reasonable limits for all fields.

In addition, the steps have to be applied to the inline and crossline profiles separately, resulting in one perturbation factor based on the inline direction μ_{in} and a second one based on the crossline direction, μ_{cross}. This essentially means that the divergence of the beam is not symmetric, as expected [126, 166].

To validate each individual perturbation factor, the full (not cropped) energy-optimized PhSp file was perturbed by μ_{cross} and μ_{in} separately. Using these PhSps, respective lateral profiles in the crossline and inline direction were simulated 24 times for all the squared fields considered. Fig. 4.9 shows the median values of k and x_0 for the crossline profiles of the 10×10 cm^2 field, together with the cost values obtained before (initial cost) and after (final cost) the manipulation with μ_{cross}. The same behavior was observed for inline profiles.

The result was two different PhSps, each one optimized for profiles in one direction. The goal was, however, to obtain one single optimized PhSp that can be used for any simulation. Therefore, both factors μ_{cross} and μ_{in} needed to be merged into one single factor, μ_{opt}. As the cost values for inline profiles were already low after the energy optimization step, special focus was put into the optimization of crossline profiles, while preserving the low cost values for inline profiles.

4.1.4.2 μ_{opt}: merging μ_{cross} and μ_{in}

The methodology for obtaining one single factor μ_{opt} derived from μ_{cross} and μ_{in} was a result of several steps investigated [193]. The initial assumption was that the crossline profiles were mostly dependent on the divergence of the particles positioned near the crossline axis (X), illustrated by the pink-shaded region in Fig. 4.10a. The same assumption was made for the inline profiles. This hypothesis can be described by Eq. 4.26, where the perturbation factor μ_ϕ applied to each particle's directions is dependent on its angular position ϕ with respect to the crossline axis. Eq. 4.26 ensures that particles positioned exactly on the crossline and inline axis are perturbed exclusively by μ_{cross} and μ_{in}, respectively, with proportional contribution of both to particles positioned in between both directions (Fig. 4.10b, blue curve).

$$\mu_\phi = cos(\phi)^2 \mu_{cr} + sin(\phi)^2 \mu_{in} \tag{4.26}$$

A full energy-optimized PhSp was perturbed using μ_ϕ and used to simulate lateral profiles for all fields. An increase in the resultant cost values was observed, in comparison to values obtained for the two separately optimized PhSps, suggesting that particles positioned in the immediate surroundings might be perturbed insufficiently (for crossline) or excessively (for inline).

To compensate for the fall-off of μ_ϕ, μ_{cross} and μ_{in} values were shifted to higher and lower values μ'_{cross} and μ'_{in}, respectively (Eqs. 4.27 and 4.28) (Fig. 4.10c, blue line).

$$\mu'_{cr} = \zeta \cdot \mu_{cr} \quad , \quad \zeta \geq 1 \tag{4.27}$$

$$\mu'_{in} = \eta \cdot \mu_{in} \quad , \quad \eta \leq 1 \tag{4.28}$$

Several combinations of μ'_{cr} and μ'_{in} were evaluated, with μ'_{cr} ranging from 1.00605 to 1.014 and μ'_{in} from 0.997 to 1.00051. A cross-correlation between μ_{cr} and inline profiles could be identified, with same behavior observed for μ_{in} and crossline profiles, implying that an attempt at optimizing the crossline profiles affected the resultant inline profiles considerably, and vice versa. This results in a trade-off between the quality of simulated profiles in the inline and crossline directions, as observed in Fig. 4.11 for some combinations of μ'_{cr} and μ'_{in}, thus confuting the initial assumption of lateral profiles depending strictly on particles positioned locally.

The cross-correlation, and subsequent trade-off, could be controlled by Eq. 4.29, a modification of Eq. 4.26, where the final optimization factor μ_{op} has a faster fall-off (Fig.

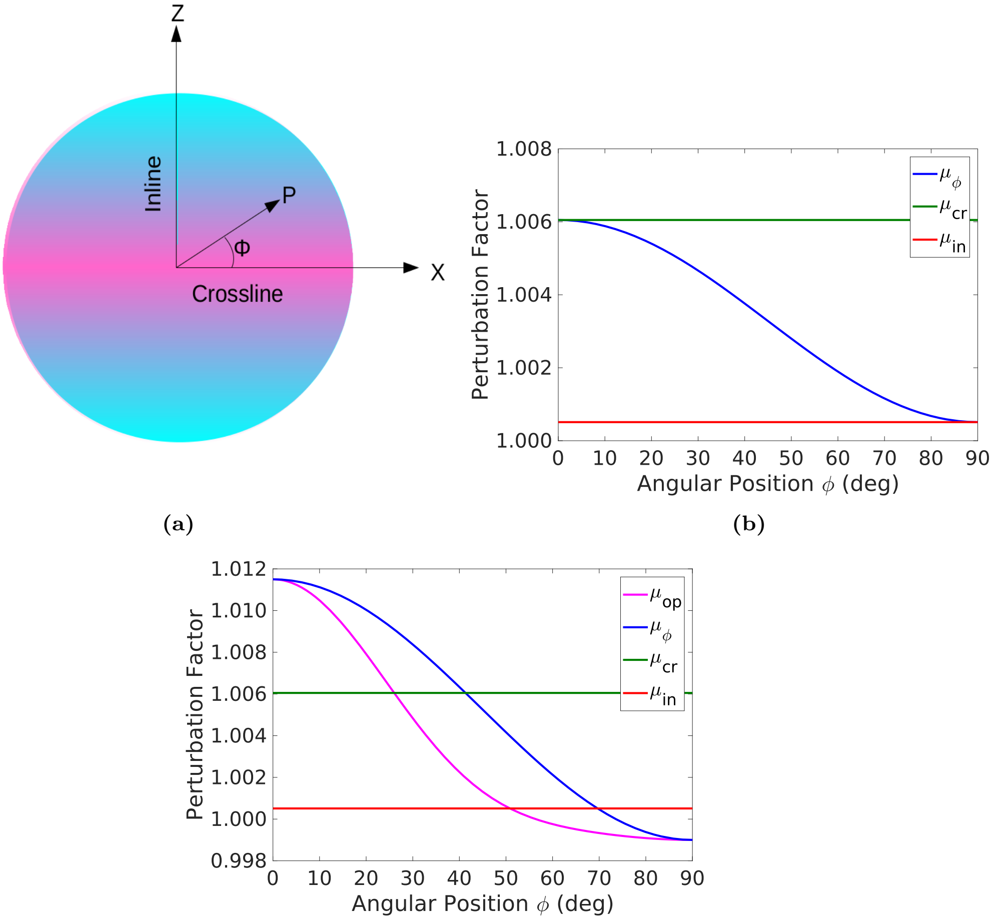

(a) **(b)**

(c)

Figure 4.10: Merging μ_{cr} and μ_{in}: (a) A schematic representation of the PhSp plane, where ϕ is the angular position of a particle with respect to the crossline direction. The pink shaded region represents the relevance of the particles for crossline dose profiles, as explained in the text (not shaded to scale); (b) The variation of μ_ϕ with ϕ; (c) μ_ϕ and μ_{op} for same μ'_{cr} and μ'_{in}, a faster fall off can be observed for μ_{op}. The green and red lines are the individual perturbation factors, μ_{cr} and μ_{in} respectively. Reprinted with permission from [194].

4.10c, pink curve).

$$\mu_{op} = cos(\phi)^6(\mu'_{cr} - 1) + sin(\phi)^6(\mu'_{in} - 1) + 1 \tag{4.29}$$

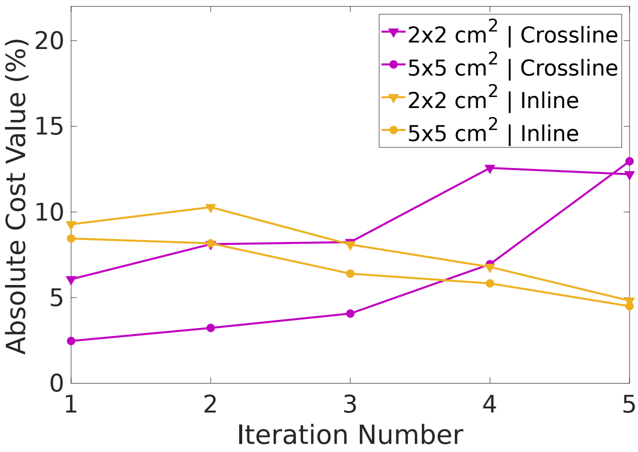

Figure 4.11: Crossline and inline trade-off: Absolute cost behavior with multiple combinations of μ'_{cr} and μ'_{in}, represented by the horizontal axes (iteration number), for μ_{ϕ}. A trade-off between the crossline and the inline cost values is evident. Reprinted with permission from [194].

4.2 Results

4.2.1 Cropping the Phase Space

To crop the PhSp for PDD simulations, the influence of particles located at different radial regions of the PhSp needed to be evaluated. Fig. 4.12 shows the comparison between PDDs curves simulated with the large-energy-perturbed PhSp (as described in section 4.1.2) and simulations performed with the original IAEA PhSp. When the perturbation was applied to particles located beyond 20 mm from the center (Fig. 4.12e and 4.12f), the cost values between both simulated curves drops to under 1.0%. This indicates a low influence of particles beyond this point to simulations of PDD profiles. Therefore, for all PDD simulations, a cropping radius of 20 mm was selected.

For lateral profiles, the cropping radius was defined based on field size and magnification, as described in section 4.1.2. The individual values for each field size are shown in Table 4.1. To ensure no relevant particle was disregarded, simulations performed with cropped and full PhSps were compared against each other, using the cost values between both simulations. The cropping radius was accepted upon the condition that the cost between both simulations was below 3.0%, keeping a good trade-off with the simulation time.

Table 4.1: Cropping radii versus field sizes.

Field Size (cm^2)	Cropping Radius (mm)
2×2	20
5×5	20
10×10	30
20×20	50

Furthermore, the influence of the cropping radius in the determination of μ_{in} and μ_{cross} was investigated. Two PhSps with 30-mm and 50-mm cropping radius were compared. A cropping radius of 50 mm is enough to simulate all four different squared fields, therefore all fields were used for the determination of μ_{in} and μ_{cross} in combination with the 50-mm cropped PhSp. For the 30-mm cropped PhSp, only the 2×2, 5×5 and 10×10 cm^2 fields were used, as this cropping is too small for the 20×20 cm^2 field. The absolute difference between the values obtained for μ_{in} and μ_{cross} for the two different cropping are under 0.01%, as shown in Table 4.2. Moreover, using the smaller cropping radius reduces the simulation time by half. Hence, a cropping radius of 30 mm combined to 3 different field sizes would be sufficient to determine the individual perturbation factors μ_{in} and μ_{cross}, with half computation time.

It was not possible to determine the exact gain on time efficiency obtained with the cropping technique, as the simulation time depends on many factors, such as the profile type (lateral or PDD), the field sizes and specially the specifications of the machine used

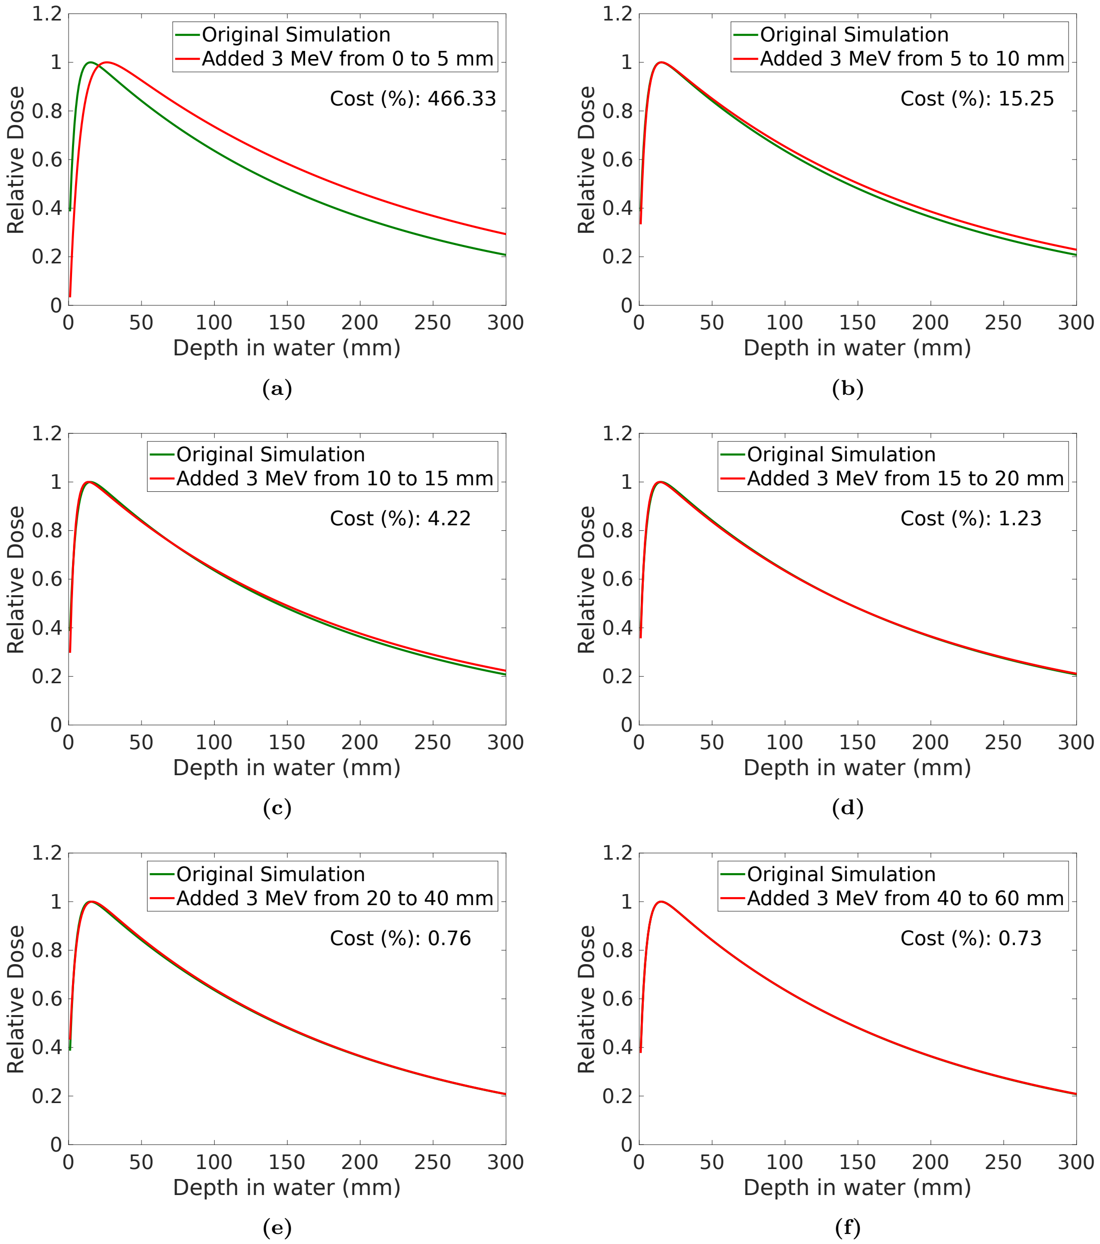

(a) (b) (c) (d) (e) (f)

Figure 4.12: PDD Cropping Radius: the effect of large energy perturbations to particles within different radial regions in the PhSp. The green curves represent the fitted PDD profiles simulated using the original IAEA PhSp, while the red curves represent the fitted PDDs simulated using PhSps with 3 MeV energy perturbation to particles within a radial distance of (a) 0 to 5 mm, (b) 5 to 10 mm, (c) 10 to 15 mm, (d) 15 to 20 mm, (e) 20 to 40 mm and (f) 40 to 60 mm. It can be inferred that, beyond 20 mm, the particles' influence on the PDD profiles drops significantly. Reprinted with permission from [194].

to run the simulation. In this study, a computer cluster was used, comprised of machines with different individual specifications. Additionally, the workload on these machines affects their performance. In general, using cropped PhSps reduced the simulation time by at least 60%, even when the largest cropping radius was considered (50 mm for lateral profiles of 20×20 cm^2 field). For PhSps with a 20-mm cropping radius, used for lateral profiles of 2×2 and 5×5 cm^2 and all PDDs, the simulations were up to 20 times faster.

Table 4.2: Optimal perturbation factors for crossline and inline directions, determined using PhSps with different cropping radii. The percentage differences are also presented.

Cropping Radius (mm)	μ_{cr}	μ_{in}
30	1.00601	1.00042
50	1.00605	1.00051
Difference (%)	0.004	0.009

4.2.2 Energy optimization

A perturbation of 0.32 MeV to all particles inside the PhSp produced the lowest cost values between PDD profiles simulated using the cropped-energy-perturbed PhSp and measurements, for all field sizes. Therefore, the 0.32 MeV was identified as the optimal energy perturbation factor and applied to the entire LMU PhSp. Fig. 4.13 shows the fitted PDD profiles (Eq. 4.4) from all field sizes, simulated using both the original IAEA PhSp (red curves) and the energy-optimized LMU PhSp (orange curves), in comparison to the fitted measurements (green curves). The cost values decrease significantly after the energy optimization, as shown in Table 4.3. The respective gamma passing rates, i.e. the percentage of points passing the (3 %, 3mm) criteria, are presented in Table 4.4.

Table 4.3: Initial and final cost values of PDD profiles, for the optimization of energy, when using the original and energy-manipulated full PhSp, respectively.

Field Size (cm^2)	PDD Initial Cost (%)	PDD Final Cost (%)
2 × 2	2.9	1.1
5 × 5	6.4	1.4
10 × 10	11.2	1.1
20 × 20	9.1	1.3

Table 4.4: Initial and final gamma passing rates for PDD profiles, for the optimization of energy, when using the original and energy-optimized full PhSp, respectively.

Field Size (cm^2)	PDD Initial passing rates (%)	PDD Final passing rates (%)
2 × 2	94	97
5 × 5	95	99
10 × 10	92	98
20 × 20	96	97

The influence of the energy optimization on the lateral profiles is show in Table 4.5, for inline and crossline profiles. For the inline profiles, the cost values decreased. For crossline profiles, no relevant changes were observed for the smaller fields, while important improvements were obtained for the 10×10 and 20×20 cm^2 fields.

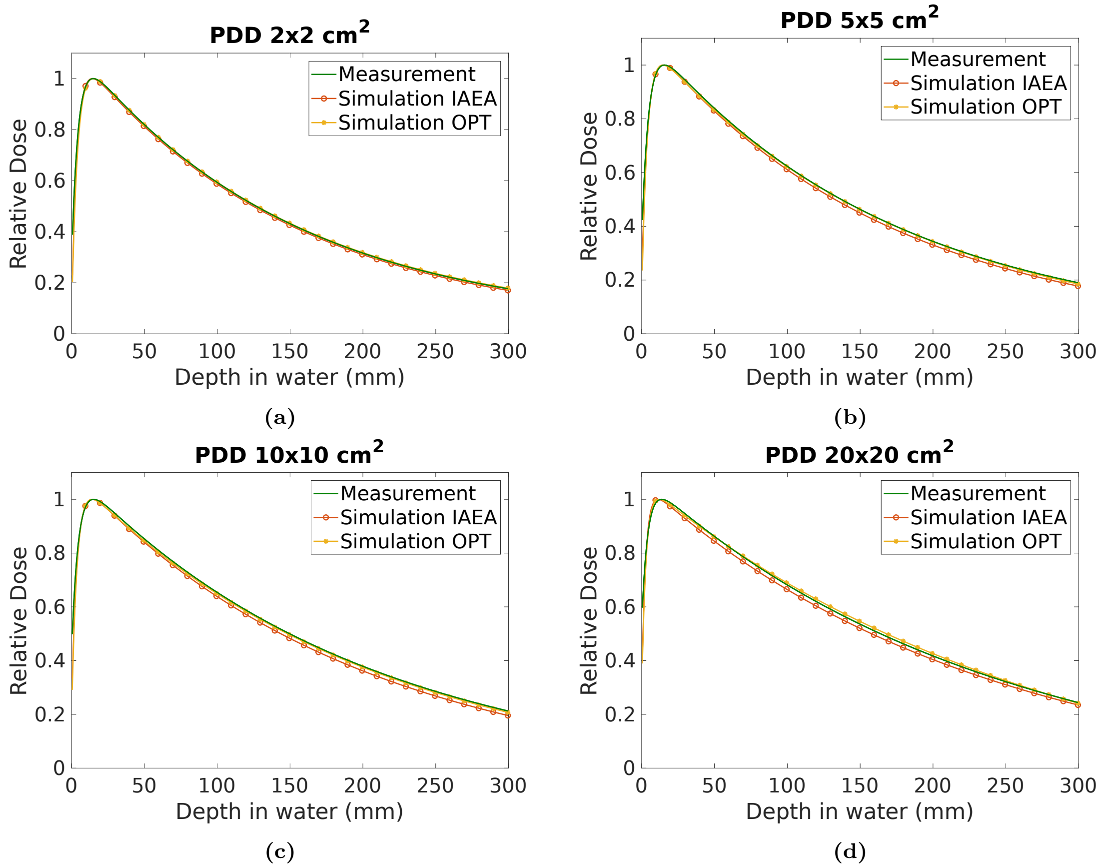

(a) **(b)** **(c)** **(d)**

Figure 4.13: PDD Comparison: Measured PDDs (green) compared to simulated PDDs using the original IAEA PhSp (red) and the energy-optimized PhSp (orange), for field sizes of (a) 2×2, (b) 5×5, (c) 10×10 and (d) 20×20 cm^2. All profiles have been generated using full PhSp for validation. The curves represent the fitted data (Eq. 4.4). Reprinted with permission from [194].

4.2.3 Direction optimization

The values obtained for the individual perturbation factors in the inline and crossline directions, as a result from the 8 steps described in Section 4.1.4, were $\mu_{cr} = 1.00605$ and $\mu_{in} = 1.00051$ (Table 4.2). The individual factors were merged into a single optimization factor μ_{opt}, as defined by Eq. 4.29. The best compromise between inline and crossline cost values was achieved for $\mu'_{cr} = 1.012$ and $\mu'_{in} = 1.000$, which were then taken as the optimal values for the definition of μ_{opt}.

A final full-optimized PhSp - the LMU PhSp, with energy and direction manipulation, was created and used to simulate lateral profiles in both directions, for all field sizes, at 15 mm and 100 mm depths in the virtual water phantom. The simulation results were compared to measurements performed under the same conditions with the Elekta Synergy, via the cost values (Tables 4.6 and 4.7). The same evaluation was performed to profiles simulated using the initial IAEA PhSp, validated for the Elekta Precise.

Table 4.8 shows the gamma passing rates with (3%, 3 mm) criteria for simulations performed with the LMU PhSp and the IAEA PhSp, in comparison to measurements, at

Table 4.5: Initial and final cost values of lateral profiles in the inline and crossline directions, for the optimization of energy, when using the original and energy-optimized full PhSp, respectively.

	Inline		Crossline	
Field Size (cm^2)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (Energy Opt.)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (Energy Opt.)
2 × 2	4.3	3.1	24.2	22.4
5 × 5	8.5	3.1	26.3	26.6
10 × 10	15.7	6.0	22.1	17.6
20 × 20	8.6	1.2	16.4	12.9

15 mm depth. For the 100-mm depth profiles, no relevant improvement was observed in the gamma evaluation: for all fields, the passing rates were above 94% for simulations performed with both the LMU Phsp and the original IAEA PhSp.

Table 4.6: Initial and final cost values of lateral profiles in the inline and crossline directions, at 15 mm depth, when using the original IAEA and final-optimized LMU PhSp, respectively.

	Inline - 15 mm		Crossline - 15 mm	
Field Size (cm^2)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (LMU PhSp)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (LMU PhSp)
2 × 2	4.3	1.6	24.2	13.2
5 × 5	8.5	2.9	26.3	10.6
10 × 10	15.7	4.6	22.1	8.1
20 × 20	8.6	0.8	16.4	8.0

Table 4.7: Initial and final cost values of lateral profiles in the inline and crossline directions, at 100 mm depth, when using the original IAEA and final-optimized LMU PhSp, respectively.

	Inline - 100 mm		Crossline - 100 mm	
Field Size (cm^2)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (LMU PhSp)	Initial Cost (%) (IAEA PhSp)	Final Cost (%) (LMU PhSp)
2 × 2	6.3	4.0	27.5	12.5
5 × 5	6.5	2.9	25.1	6.8
10 × 10	3.2	4.7	11.6	0.8
20 × 20	7.0	1.8	6.1	5.8

For a qualitative assessment of the improvement due to the entire optimization process, Figs. 4.14 and 4.15 show the penumbra region of profiles (data fitted according to Eq. 4.6) at 15 mm, for crossline and inline directions respectively. The same behavior was observed for profiles at 100 mm depth. Because the cost values for the lateral profiles were calculated as the median of individual values from 24 instances, the curves displayed in the figures are sigmoid curves produced using the median values of k and x_0. The curve representing the measured profile is also the fitted sigmoid, for consistency.

To ensure that the direction optimization did not negatively affect the PDD profiles, new PDD simulations were performed using the final optimized LMU PhSp. The cost values obtained remained low and comparable to the values obtained after energy optimization (Table 4.3), as illustrated by Fig. 4.16 for the 10×10 cm^2 field. The same behavior was observed for all field sizes.

Table 4.8: Initial and final gamma passing rates for lateral profiles in the inline and crossline directions, at 15 mm depth, when using the original IAEA and final-optimized full PhSp, respectively.

	Inline		Crossline	
Field Size (cm^2)	Initial gamma passing rates)%)	Final gamma passing rates (%)	Initial gamma passing rates (%)	Final gamma passing rates (%)
2×2	100	97	100	100
5×5	96	99	97	97
10×10	90	99	84	90
20×20	90	94	88	96

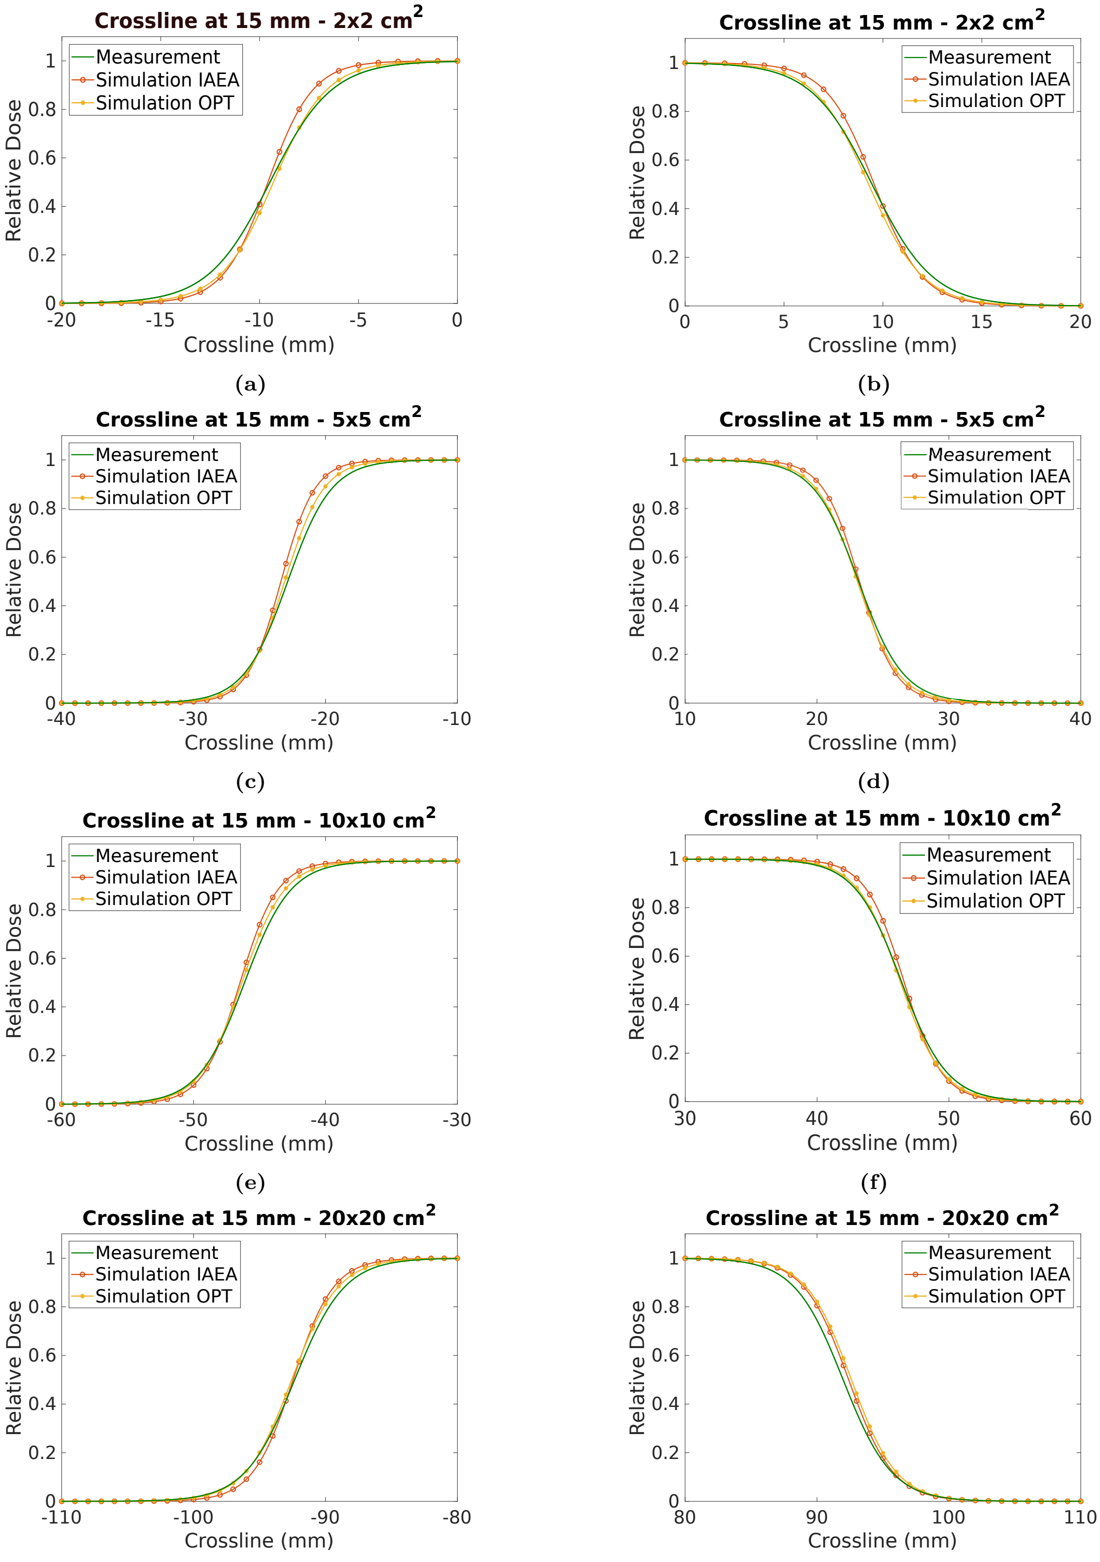

(a) **(b)** **(c)** **(d)** **(e)** **(f)** **(g)** **(h)**

Figure 4.14: Crossline Comparison at 15 mm depth: Measured crossline profiles (green) compared to the ones simulated using the original IAEA PhSp (red) and the final-optimized PhSp (orange), for field sizes of (a),(b) 2×2, (c),(d) 5×5, (e),(f) 10×10 and (g),(h) 20×20 cm^2. All profiles have been generated using full PhSp for validation. The curves represent the fitted data (Eq. 4.6). Reprinted with permission from [194].

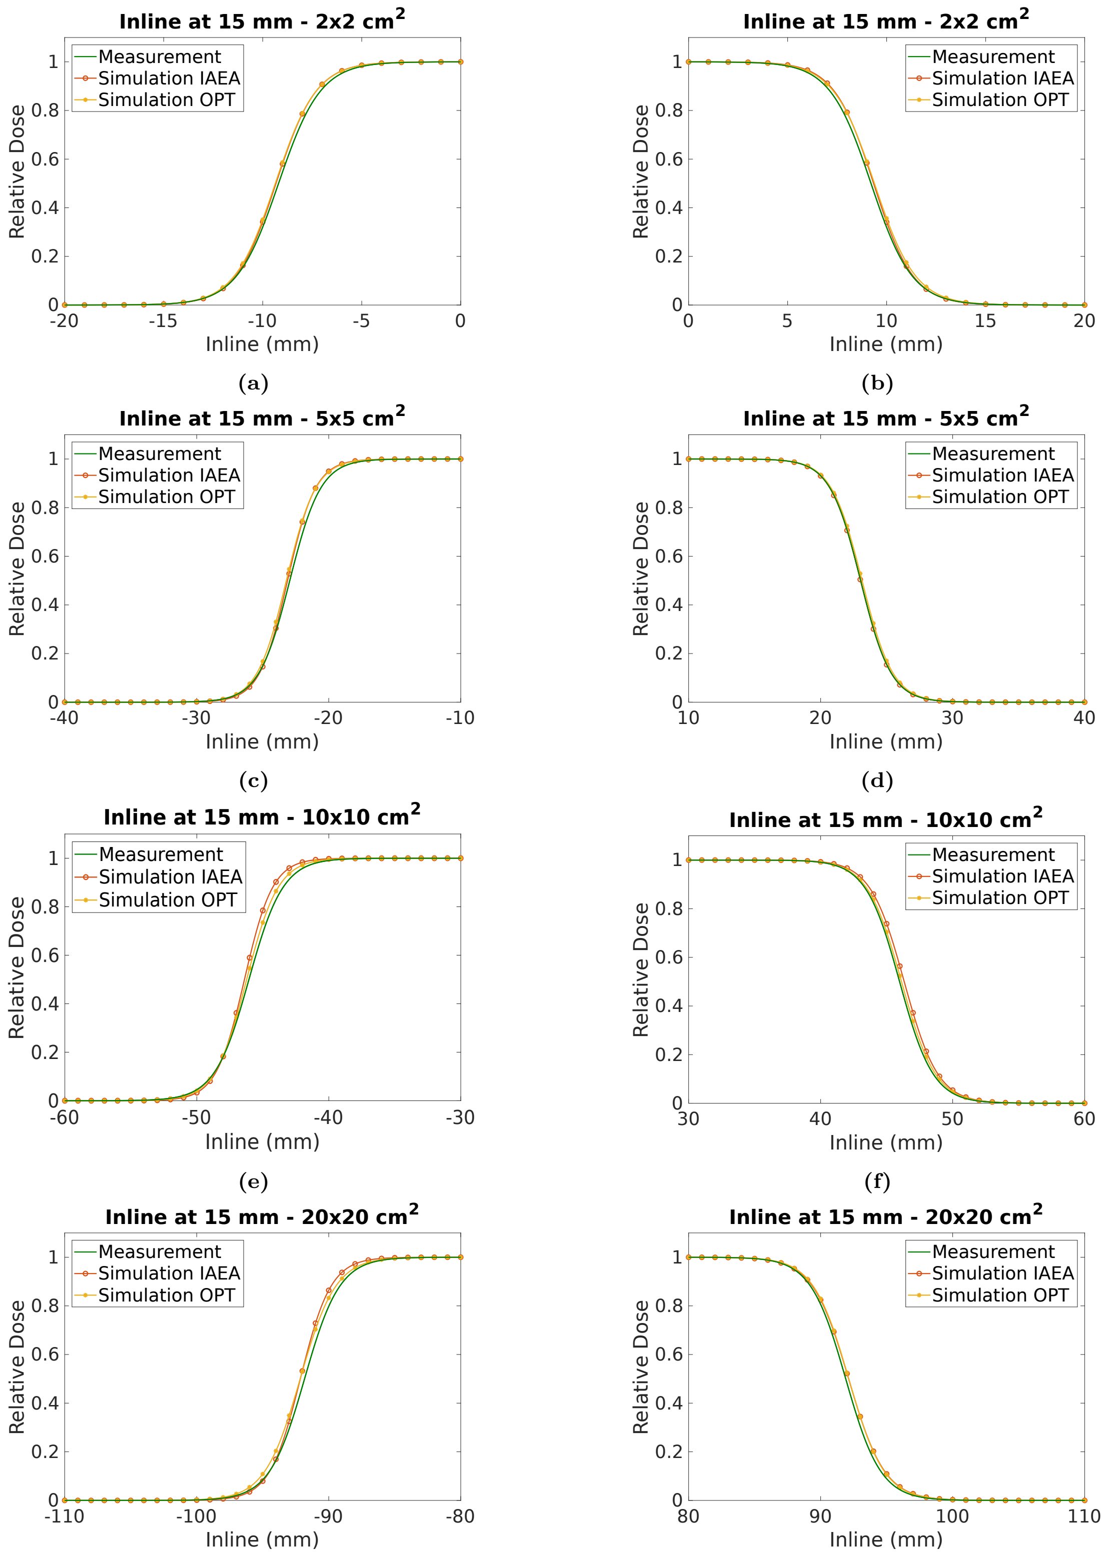

(a) **(b)**

(c) **(d)**

(e) **(f)**

(g) **(h)**

Figure 4.15: Inline Comparison at 15 mm depth: Measured crossline profiles (green) compared to the ones simulated using the original IAEA PhSp (red) and the final-optimized PhSp (orange), for field sizes of (a),(b) 2×2, (c),(d) 5×5, (e),(f) 10×10 and (g),(h) 20×20 cm^2. All profiles have been generated using full PhSp for validation. The curves represent the fitted data (Eq. 4.6). Reprinted with permission from [194].

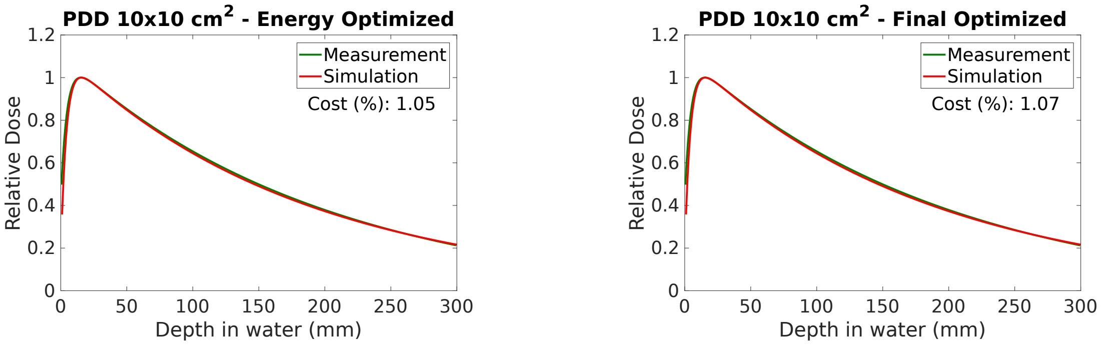

(a) **(b)**

Figure 4.16: PDD profiles for the 10×10 cm^2 field before (a) and after (b) the optimization of direction: The low cost values obtained after energy optimization are maintained. The curves represent the fitted data (Eq. 4.4). Reprinted with permission from .[194].

4.3 Discussion

A PhSp is essentially a matrix containing an enormous amount of information describing the particles it stores. This information is in the form of different interconnected parameters, such as the exact point where one particle crosses the PhSp (on the PhSp plane), the angular momentum with which this particle is traveling, its energy, among others. All these parameters depend on the geometric characteristics of the linac MC model used to generate the PhSp. The production of a PhSp for a specific linac model is therefore subject to the availability of detailed geometric information about the equipment's components, scarcely available in the literature and only disclosed by the vendors upon special agreements. This study investigated the possibility to modify some information inside a PhSp, generated and validated for one linac model, in order to optimize it for a different model. The methodology proposed focuses on manipulating the energy and angular distribution of the particles, through the manipulation of their direction cosines.

As the methodology is an iterative process, the cropping technique introduced in Sec. 4.1.2 was essential, as several hundreds of simulations were needed to estimate the optimal perturbation factors for both energy and direction. For the direction optimization, the manipulation factors μ_{cr} and μ_{in} can be obtained with a 30-mm cropping radius, for the linac model and field sizes considered in this study. With a cropping radius of 20 mm, it was possible to determine the optimal energy manipulation factor and the cost values for lateral profiles from the 2×2 cm^2 field. An even smaller cropping radius of 5 mm was unacceptable for calculating absolute cost values for any field, yet it was extremely useful for determining the linear relationship between μ and the sigmoid fitting parameters k and x_0 (Sec. 4.1.4, Fig. 4.7), which guided the method used for the manipulation of the particles directions.

For the manipulation of energy, a straightforward addition operation was sufficient. An additional 0.32 MeV to the energies of all particles inside the original LMU PhSp (corresponding to the Elekta Precise IAEA PhSp), except for annihilation photons whose energy was kept unchanged at 0.511 MeV, improved the PDD simulation results for an Elekta Synergy. This perturbation decreased the PDD cost values from up to 11.4% down to a maximum of 1.4% for all fields (Table 4.3). The energy optimization also increased the agreement between measured and simulated lateral profiles in the inline direction (Table 4.5).

Finding a proper methodology for the angular distribution optimization was, in contrast, very demanding. The linear relationship between the sigmoid fitting parameters of the lateral profiles and the slope of the beam divergence line provided a guidance on how to proceed with the direction optimization (Fig. 4.7). The technique to manipulate the directions could be summarized in 8 steps, resulting in two distinct optimization factors μ_{cr} and μ_{in} for crossline and inline directions separately, proving that the electron beam of the Elekta Synergy is indeed elliptical [126, 166]. As the goal was to find one manipulation that could optimize the PhSp for any kind of simulation, both factors needed to be merged into one single optimal manipulation factor μ_{opt}. It was observed that crossline and inline dose profiles do not depend strictly on the divergence of particles positioned locally, implying that an attempt to optimize solely the crossline profiles affects the inline profiles significantly, and vice versa, as displayed in Fig. 4.11 for some combinations of μ'_{cr} and μ'_{in}. Several attempts were investigated for the development of a merging technique, until the best trade-off between crossline and inline cost values was achieved through Eq. 4.29.

For validation, both energy and direction manipulations were applied to the original full LMU PhSp, producing a final optimized LMU PhSp. The final LMU PhSp was used to simulate PDD for all field sizes, as well as lateral dose profiles for both directions at 15 mm and 100 mm depth in a virtual water phantom. The simulations were compared to measurements performed under the same conditions, via cost values and gamma evaluation. Remarkable improvements were obtained for the inline profiles simulated with the final LMU PhSp, in comparison to the original IAEA PhSp: the absolute cost values dropped to under 4.6% and 4.7%, for all fields analyzed at 15 mm and 100 mm depths, respectively, representing reduction factors of at least 2.7 and 1.6 (Tables 4.6 and 4.7). A small increase on the cost value of the 10×10 cm^2 was observed for the profile at 100 mm depth, however its final cost (4.7%) was comparable to the final cost values of the other fields. At the crossline direction, all absolute cost values dropped to under 10.6% and 6.8% for dose profiles at 15 mm and 100 mm, respectively, except for the smallest 2×2 cm^2 field, whose final costs were 13.2% and 12.5% at 15 mm and 100 mm depth, respectively. Even though the final cost values for the crossline profiles were above the desirable range in this study (5%), they still represent a reduction of about 50% in comparison to simulations performed with the original IAEA PhSp. Moreover, the qualitative improvement on the simulations can be seen in Figs. 4.14 and 4.15.

The gamma evaluation was performed using a 3% dose-difference (DD) and a 3-mm distance-to-agreement (DTA) criteria (3%, 3 mm). The gamma passing rates were used to assess the quality of the simulations and the improvements achieved by the manipulations, however with the chosen criteria it failed in detecting small discrepancies between measured and simulated data, especially for the smallest 2×2 cm^2 field (Table 4.8). Stricter criteria could improve the sensitivity of the gamma evaluation, therefore passing rates with stricter DTA criteria were also calculated (3%, 2 mm and 3%, 1 mm), but the resultant values were similar. Stricter DD criteria could not be used due to the magnitude of the simulation uncertainties ($\approx 2.5\%$). Based on this information, the gamma evaluation method could be potentially used for PhSp optimization only if more stringent DD and DTA criteria are adopted. These can only be safely applied by reducing the statistical uncertainties of simulations, bound to an eventual increase on the computational time. On the other hand, the cost values introduced provide a valid alternative with less stringent statistical requirements. At the same time, the cost values as defined in this study offer relevant information to the user, in particular for the lateral profile, as they provide valuable insights on the nature of the discrepancies that can serve as a guidance during the optimization process.

Furthermore, the proposed methodology is an iterative process, meaning that the PhSp can be potentially further optimized as desired. As a good compromise between results and simulation time, we decided to stop the iterations once all the costs dropped to under 50% of their original values. Some potential - but not exclusive - measures that could improve the optimization process are listed below:

1. Steps 1 to 4 in Sec. 4.1.4 can be performed with more perturbed PhSps to obtain the linear fit (Fig. 4.9), which can potentially lead to better values of μ_{cr} and μ_{in};
2. An even higher number of combinations of μ'_{cr} and μ'_{in} can be investigated (Eqs. 4.27 and 4.28);
3. The evaluation of lateral profiles in different directions (e.g. diagonal) could provide further information regarding the divergence of the beam;
4. Additional information could be obtained by evaluating different collimator angles (e.g. 90°);
5. Manipulation of additional particle parameters could be considered, as a complement to the energy and angular distribution optimization.

It is important to notice that a specific PhSp will not necessarily fit all machines of the same model to which it is validated, as the geometric parameters are essentially machine-specific due to small differences in installations and setting [141], and so will be the PhSps generated by their MC models. The methodology developed here offers the flexibility to potentially optimize and tune a PhSp to a specific machine as much as desired. Moreover, the methodology was applied here to a unique set of linac models (Elekta Precise versus Elekta Synergy), but the general characteristics of PhSps are essentially the same regardless of the machine model. Therefore, the methodology is generic and can be applied to any equipment from any vendor.

Finally, the optimization of particles' energy and angular distribution relates mostly to the electron source and focal spot of the linac MC model used to generate the PhSp. Further differences in the geometry of the Elekta Precise and Elekta Synergy, which could greatly influence the output of the simulations [133], could not be assessed by this method and might be a potential limiting factor on the improvements that can be achieved.

4.4 Conclusions

The present study proposes a general methodology for manipulation of an existing IAEA PhSp, generated and validated for a certain linac model, for the production of a new PhSp, which could be used for simulating a different linac. Starting from the Elekta Precise IAEA PhSp, the methodology produced an optimized PhSp that can better reproduce the measurements from an Elekta Synergy linac. The optimization occurred by manipulating particles energies and direction cosines, using optimal perturbation factors determined throughout this work. Remarkable improvements were obtained for PDD and lateral profiles in the inline direction, for several squared fields investigated at two different depths in water. For crossline dose profiles the improvement was more subtle, yet a reduction of approximately 50% on the cost values were obtained for all the fields and depths analyzed. The PhSp can be potentially further optimized, as the proposed method is an iterative process. Even better agreement between measurements and simulations could be achieved by the full-linac-head simulation, if geometric details are available (Chapter 3). The optimized PhSp should not replace the full modelling of the linac head, which should remain as the primary method of choice, however it offers an alternative for MC linac simulations when neither the geometric details, nor the validated IAEA PhSp files, are available for the user.

CHAPTER 5

Deep Dose Estimation

A methodology for in vivo dosimetry can only be clinically feasible if the necessary steps are reasonably simple, if it can be performed in a reasonable time, and if the workload on the clinical staff is not increased unnecessarily. Several methods for in vivo dosimetry proposed in the past years can predict the dose delivered to the patient within acceptable time, however most of them involve laborious corrections. In addition, several methods require an in-air pre-treatment EPID irradiation, which means an increase in the overall time slot for each patient treatment, on the workload on the clinical staff, and could add an extra source of uncertainty, in case the machine does not behave exactly in the same way in the in-air and in vivo irradiation. Finally, many methods cannot properly account for tissue inhomogeneity, considering the patient as a homogeneous water-equivalent volume.

Monte Carlo (MC) methods, on the other hand, provide a very accurate description of energy deposition to a medium by ionizing radiation, properly accounting for tissue inhomogeneities. However, despite the great improvements in computational capabilities, MC methods are still time consuming and can hardly fit into the clinical routine. Motivated by the great capabilities demonstrated by Neural Networks (NN) in the field of medical physics, and especially inspired by the recently proposed Deep Dose Estimation (DDE) [71] network, which predicts almost in real time the dose received by patients undergoing Computed Tomography (CT) imaging exams, this work proposes a method able to combine the strength of both: the accuracy of MC and the computational efficiency of NN. Parts of this chapter were previously published in Physica Medica [197].

5.1 Methodology

Inspired by the success of the Deep Dose Estimation (DDE) network in estimating patient-specific dose distributions for CT acquisitions in the context of radiological applications [71, 80], the DDE was extended to radiotherapy treatments and was proposed as a potential method for EPID-based in vivo dosimetry. The network was trained to predict dose distributions inside a patient CT for step-and-shoot clinical IMRT treatment fields, giving as input the patient's CT image and a first-order dose (FOD) approximation. An accurate dose distribution (ADD) inside the patient CT, produced by MC simulations, was given as the training target.

For ADD simulations, the patient CT was imported in an in-house validated Geant4 MC model of an Elekta Synergy® linac, equipped with an Elekta Agility^{TM} Multi-Leaf Collimator (MLC) (Elekta Oncology Systems, Crawley, UK), as described in chapter 3. The transmitted signals reaching the virtual EPID were simultaneously simulated and recorded as 2D absorbed dose distributions in water. These 2D EPID signals contain intrinsic information regarding the dose deposited in the patient geometry and were used to produce 3D FODs. Therefore, for each IMRT field simulated, one closely-related ADD-FOD pair was produced. The DDE network, as well as the necessary steps involved in the production of the used dataset, are explained in details in this chapter.

5.1.1 Patient CT

The patient CT database was comprised of 83 different thorax-abdominal diagnostics CT scans, 58 of male patients and 25 of female patients. The patient couch was present in 45 scans, while metallic objects either inside (metallic orthopedic prosthesis or fiducial markers) or outside the patient (localization markers) were visible in 11 scans. The CT scans were acquired on a Siemens SOMATOM Force CT scanner (Siemens Healthcare AG, Forchheim, Germany), with 150 kV X-ray tube voltage, 0.6 mm slice thickness and 512 × 512 pixels (axial plane). For diagnostic CT exams, the scan parameters can vary depending on the patient's anatomy and the purpose of the diagnostic image. Therefore, the pixel sizes on the axial plane were not the same for all 83 scans, with values ranging from 0.69 mm to 0.98 mm. To account for such differences, the images were pre-processed before being used, as described in the following.

First, each one of the 83 CTs was visually examined and 320 slices, centered on the pelvic bones, were manually selected, corresponding to a total length of 192 mm in the axial direction. This reduced the number of images used in the simulations and as input to the DDE, while the size is still large enough to be irradiated with the clinical IMRT prostate plans considered in this work, since the largest field extends over 90 mm in the axial direction. Next, the selected slices were resampled with cubic interpolation to an isotropic 2 × 2 × 2 mm^3 pixel size. The resampled pixel size was chosen as a compromise between resolution and computational cost for the MC simulations.

As the original pixel spacing was different among the available CT images, after resampling, the pixel number in the axial plane varied among the dataset. Moreover, the CT images, as extracted from the scanner, were given on a scale ranging from zero to approximately 2200, where the air voxels were assigned to zero. Therefore, the resampled images were symmetrically padded with zeros at the edges (i.e. in air, outside the patient), such that all scans had 256 × 256 × 96 voxels. The CT images were rescaled to Hounsfield Units (HU) using scaling factors given inside the CT DICOM head.

The final resampled and padded CT images have 256 × 256 × 96 voxels with 2 × 2 × 2 mm^3 isotropic voxel size. The voxels were distributed symmetrically around the isocenter, extending from -256 mm to 256 mm in the x and y directions (transaxial plane), and from -96 mm to 96 mm in the z direction (axial direction). The X axis corresponds to patient's right-left direction, the Y axis is the posterior-anterior direction and the Z axis is the cranio-caudal direction, for a patient in the head-first supine position (HFS), as defined in Chapter 3. The final CT isocenter corresponds to the linac isocenter in the MC model, to which all the simulated IMRT fields are defined.

To discriminate voxels inside and outside the patient contours, a binary mask was created for each patient, where 0 and 1 correspond to voxels outside and inside the patient, respectively. The threshold value of -700 HU was chosen so that voxels with CT numbers ≤ -700 HU were set to 0 and voxels with CT numbers > -700 HU were set to 1. This threshold does not cut the CT couch on the scans where it is present, therefore the couch was classified by the masks as part of the patient.

5.1.2 IMRT treatment plans

As the patient CT database used consists of pelvic scans, two clinical step-and-shoot IMRT prostate treatment plans were selected for consistency. The plans were calculated with the research treatment planning system (TPS) Hyperion V.2.4.5 (Equivalent to Elekta Monaco 5.1) [178] by the Department of Radiation Oncology from the University Hospital of the Ludwig-Maximilians-Universität München (Klinikum Grosshadern). Both plans were calculated with a 6 MV nominal photon beam energy, and together have 125 different sub-fields, spread over different beams, each beam delivered at a fixed gantry angular position. Each sub-field, also called a control point (CP), has a unique combination of gantry angle, jaws positions, MLC shape (defined by the position of each leaf) and number of monitor units (MU). All these parameters are specified in the RTPlan.dcm file (DICOM file outputted by the TPS) and were extracted using an in-house MATLAB script (Appendix A). Tables 5.1 and 5.2 show some parameters of the plans.

Table 5.1: Characteristics of the first prostate step-and-shoot IMRT plan. The plan was calculated with a 6 MV photon beam nominal energy.

Prostate IMRT Beam	1	2	3	4	5	6	7	8
Gantry Angle (°)	140	100	60	20	340	300	260	220
Control Points (CP)	8	5	6	11	11	9	7	7
Total MU/Beam	85.9	77.2	66.5	99.9	97.2	80.9	81.9	77.7

Table 5.2: Characteristics of the second prostate step-and-shoot IMRT plan. The plan was calculated with a 6 MV photon beam nominal energy.

Prostate IMRT Beam	1	2	3	4	5	6	7	8	9
Gantry Angle (°)	180	140	100	60	20	340	300	260	220
Control Points (CP)	7	7	7	6	8	7	7	5	7
Total MU/Beam	74.9	85.1	101.4	76.4	77.9	82.8	103.3	70.1	92.2

5.1.3 Monte Carlo simulations

All simulations were performed using a validated MC model of the Elekta Synergy® linac equipped with an Elekta Agility^{TM} Multi-Leaf Collimator (MLC) (Elekta Oncology Systems, Crawley, UK), with the Geant4 toolkit (version 10.05.p01) [198] (Chapter 3). The linac head components were replaced by a validated phase space file (PhSp), while the MLC was modeled in detail. The virtual EPID was modeled as a 41 × 18 × 41 cm^3 water box, placed at 157 cm source-to-surface distance (SSD), with the scoring layer placed at 3 cm depth in water, at 160 cm source-to-detector distance (SDD) (Chapter 3). The patient geometry was reconstructed using the pre-processed CT DICOM images (subsection 5.1.1).

Three randomly selected CPs were assigned to each one of the 83 pelvic CTs, assuring that the same CP was not attributed more than once to the same patient CT. Using the linac MC model, the accurate dose distribution inside each patient CT (ADD) and the corresponding transmitted signal at the virtual EPID were simulated for the assigned CPs, at a fixed gantry angle of 0° (patient antero-posterior direction in HFS) for simplicity. The EPID simulations were used as input to analytically produce the first-order dose estimates (FODs, subsection 5.1.4), resulting in 249 unique ADD-FOD combinations. The resulting ADD and FOD dose distributions were linearly combined for data augmentation and divided into training and test sets (subsubsection 5.1.3.1).

An additional dataset of fields irradiated with the gantry positioned at 90° (patient left-right direction) was produced, and will be referred to as the lateral dataset. The lateral dataset was comprised of eight different fields simulated for eight different CTs, resulting in one ADD-FOD pair per patient CT. The DDE was trained using exclusively dose distributions produced with the gantry fixed at 0° (subsection 5.1.6), hence the lateral dataset was used to evaluate the performance of the trained network in predicting dose distributions from fields coming from a different direction.

Simulated doses inside the virtual patient and EPID were recorded using scoring meshes (Chapter 3). The patient scoring mesh was created as a three-dimensional grid of 256 × 256 × 96 voxels, with 2 × 2 × 2 mm^3 voxel size, covering and perfectly matching the virtual patient grid. The physical EPID panel has 1024 × 1024 voxels with 0.4 × 0.4 mm^2 voxel size. As a compromise between resolution and computational cost, the EPID scoring mesh was created as a 41 × 41 cm^2 grid, with 1 × 1 mm^2 pixel size (X-Z plane). In the

y direction, the mesh was 10 mm thick to reduce the statistical fluctuation. Both dose distributions were scored as dose-to-medium, but since the EPID was modeled as a water box, the dose scored at the virtual EPID was equivalent to dose to water.

For both the patient and EPID grids, the deposited dose (in Gy) and the deposited dose squared (in Gy^2) were scored for each voxel. The Gy^2 values were used for uncertainty estimation, calculated using the history-by-history method [175]. The resultant values were converted into absolute dose values (Chapter 3), to account for the MU values of the respective CPs, yielding the final ADD and the corresponding transmitted EPID signal. Finally, the patient-specific mask, which discriminates the voxels inside and outside the patient contour, is applied to the ADD grid to set the dose values to zero in the voxels outside the patient.

As a compromise between computational cost and simulation uncertainty, each CP was simulated 30 times with 30 different PhSp files, each one generated with 10^8 initial electrons (Chapter 3), so that each CP was simulated with 3×10^{9} independent electron histories. Particle recycling was used to further reduce the statistical noise of the results, with each particle being recycled 24 times. For photons, this is a reasonable recycling factor [175]. The estimated uncertainty was calculated considering only voxels which scored at least 50% of the maximum dose at both ADD and EPID dose distributions, as in lower dose areas fewer interactions occurred and a non-representative increase in the uncertainty was observed. The estimated relative uncertainty was 2.8% for the ADDs and 7.7% for the EPID signals, averaged over all simulated fields. The uncertainty of the EPID simulation is higher due to the smaller voxel size of the EPID scoring grid (1 mm grid spacing instead of 2 mm as used for ADDs) and due to the attenuation of the beam by the virtual patient volume (fewer scoring events happening at the EPID plane). Finally, the 30 independent simulations were performed in parallel on a computer cluster comprised of machines with distinct specifications. Furthermore, as the machines are not used exclusively for simulation purposes, the workload affects their performance. Hence, the simulation time from each field varied considerably, but on average 14 hours were spent for each CP.

5.1.3.1 Data Augmentation

To increase the size of the database, three different linear combinations of the simulated ADDs from each patient CT were produced, as a post-processing data augmentation tool, following Eq. 5.1, with different combinations of a and b ranging from 0.5 to 2.5.

$$\mathrm{ADD}_{LC} = a \cdot \mathrm{ADD}_m + b \cdot \mathrm{ADD}_n \tag{5.1}$$

ADD_{LC} is the linearly combined ADD, where ADD_m and ADD_n are the originally simulated ADDs, with m and $n \in \{1, 2, 3\}$, and $m \neq n$. A fourth linear combination was created using the three original distributions (Eq. 5.2, c, d and e ranging from 0.5 to 2.5).

$$\mathrm{ADD}_{LC} = c \cdot \mathrm{ADD}_1 + d \cdot \mathrm{ADD}_2 + e \cdot \mathrm{ADD}_3 \tag{5.2}$$

The MU values of the combined CPs were scaled accordingly. The same was performed for the FODs. This data augmentation was implemented only for the simulations performed with gantry at 0°, increasing the database to 581 different ADD-FOD pairs, seven pairs for each patient CT.

The linear combinations of the dose distributions can be interpreted as different irradiation fields, and can therefore be used as different datasets for training the neural network. Additionally, in a real IMRT treatment, several CPs with different MUs are irradiated to the patient from the same direction. Therefore, the linear combinations used for data augmentation reflect a clinical scenario.

5.1.4 First Order Dose approximation (FOD)

The simulated 2D EPID transmitted signal contains intrinsic information on the dose deposited to the patient geometry, and is thus related to the ADD. From this 2D distribution, recorded as dose to water, a simplified 3D dose distribution is produced as a ray-driven backprojection into the patient CT, called the first-order dose approximation (FOD). This approximation accounts for geometric magnification, inverse square law (ISL) and attenuation of the beam when crossing the patient. Other effects are disregarded, such as the build-up effect, scattering inside the patient and beam hardening. The steps and tools used in the FOD production are explained in the following.

5.1.4.1 Radiological path length

To account for the attenuation of the beam when crossing the patient geometry, the radiological path length is needed. The radiological path length is given by the physical path length, or accumulated intersection length, multiplied by some factor related to the material of the voxel being crossed. Each voxel was reconstructed with material and density based on its HU number, as imported from the CT DICOM image into the MC linac model. The scaling factor between physical and radiological path length is the Relative Electron Density (RED) of each voxel's material with respect to water [199].

The radiological path length L (Eq. 5.3) can be described as the sum of the physical intersection lengths l_i of a ray with each voxel crossed on its way towards the EPID, scaled by the relative electron density of that voxel's material RED_i [199]. The l_i values were obtained using a 3D ray tracing algorithm (subsubsection 5.1.4.2). To obtain the RED of each voxel, reference values from a clinical TPS, correlating the CT number (in HU) to RED, were used. From these reference values, three distinct regions could be identified, and a linear fit was applied to each region (Figure 5.1 and Eq. (5.4)), where p_1 and p_2 are the fitted parameters. Table 5.3 shows the values of p_1 and p_2 for each of the three fitting regions. The linear fits were used as conversion curves.

$$L = \sum_{i=1} \mathrm{RED}_i \cdot l_i \tag{5.3}$$

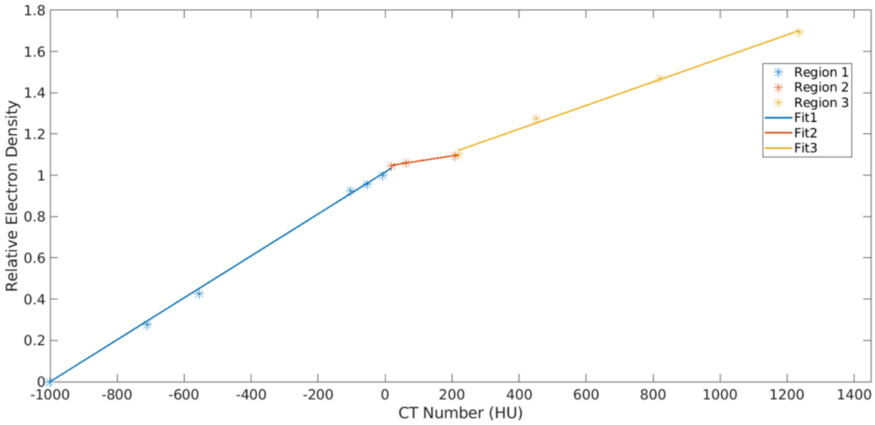

Figure 5.1: Conversion values from Hounsfield Units (HU) to Relative Electron Density (RED), obtained from a clinical TPS. The data points could be divided into three distinct regions, and a linear fit was produced for each region. The linear fits were then used as conversion curves between the HU values from the patient CTs and the RED of their respective materials.

$$RED = p_1 \cdot CT + p_2 \tag{5.4}$$

Table 5.3: The linear fit parameters p_1 and p_2 for each one of the three different regions, correlating HU and RED (Eq. 5.4).

	Region 1	Region 2	Region 3
p_1	0.0010	0.0003	0.0006
p_2	1.028	1.042	0.975

5.1.4.2 The ray tracing algorithm

In the ray-driven backprojection approach, the radiation source and the detector elements are connected through straight lines. The projections are calculated as a weighted sum of all the intersections lengths of each voxel crossed by these lines [200]. A ray tracing algorithm calculates the straight line trajectories of each ray between the radiation source and the detector elements, passing through an object placed in the way, and outputting the physical intersection length of the rays with each voxel of the object (l_i).

In this work, a modified 3D ray tracing algorithm [201] based on the Siddon method [202] was used to calculate the trajectories from the EPID scoring volume towards the radiation source (linac target, considered here as a point source). 3D matrices containing the RED of the objects being crossed were given as input to the algorithm, so that the physical intersection lengths (l_i) were converted into radiological intersection lengths ($RED_i \cdot l_i$). Moreover, the algorithm was set to output both the radiological intersection lengths and cumulative radiological intersection lengths from the exit of the patient CT towards its entrance. The results for each object were saved separately in two 256 × 256 × 96 matrices. The outputted cumulative radiological intersection length corresponds to the radiological path length given by Eq. 5.3. Ray tracing was performed for every patient CT and for a voxelized water box of the same dimensions.

To illustrate the outputs of the ray tracing algorithm, Fig. 5.2b and Fig. 5.2c show the radiological intersection and cumulative radiological intersection lengths, respectively, of one ray crossing the water box of 256 × 256 × 96 voxels in the central axis (CAX), and reaching the central pixel of the EPID scoring volume with no divergence, represented by the blue ray in Fig. 5.2a. As the RED for water is equal to one ($RED_w = 1$), the physical lengths correspond to the radiological lengths. The voxel size is $2 \times 2 \times 2$ mm^3, therefore the radiological intersection lengths are 2 mm for all voxels and the cumulative radiological intersection length is 512 mm at the entrance of the volume.

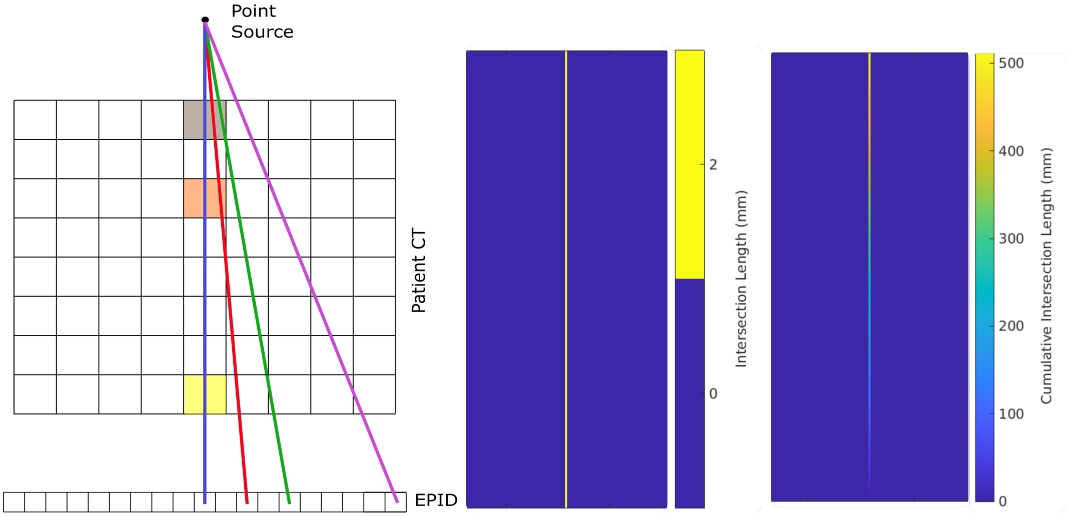

(a) **(b)** **(c)**

Figure 5.2: Representation of the ray tracing algorithm. (a) An exemplary grid with four straight lines (rays) connecting the radiation source and distinct detector elements. (b) The radiological intersection length of central voxels of a water grid crossed by a ray traveling in the central axis. (c) The cumulative radiological intersection length of the same voxels crossed by the ray at central axis. The figures are not to scale.

For the one dimensional case, i.e. when only one ray is considered, the resultant values are clear as illustrated in Fig. 5.2. When several rays are considered and the ray tracing is extended to two and three dimensions, not all voxels in the trajectories will be crossed by the same number of rays. Due to the divergence of the radiation beam, voxels closer to the source will be crossed by several rays, while voxels closer to the detector array will be crossed by fewer rays, or might even be completely missed by them, as illustrated by the highlighted voxels in Fig. 5.2a. This causes an artifact known as Moiré effect, commonly observed in ray-driven back projections. Fig. 5.3a and Fig. 5.3b show the central slice of the (physical) intersection length and cumulative intersection length matrices inside the water box, considering 410 × 410 rays, i.e. each voxel of the EPID hit by one ray, where the Moiré artifact can be clearly identified. This artifact can be corrected by normalizing the matrices by the sum of the (physical) intersection lengths observed by each voxel. Figs. 5.3c and 5.3d illustrate the (physical) intersection and cumulative intersection lengths on the central slice of the water box after correction. All matrices created with the ray tracing algorithm have been corrected before being used for the FOD production.

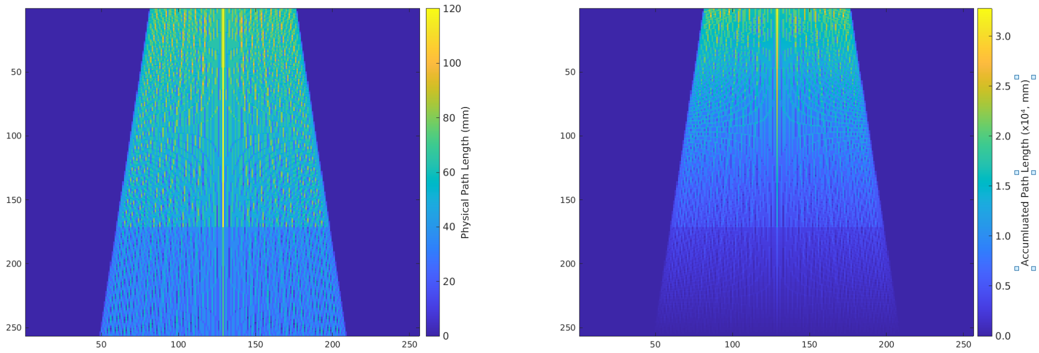

(a) Intersection length, before correction **(b)** Cumulative intersection length, before correction

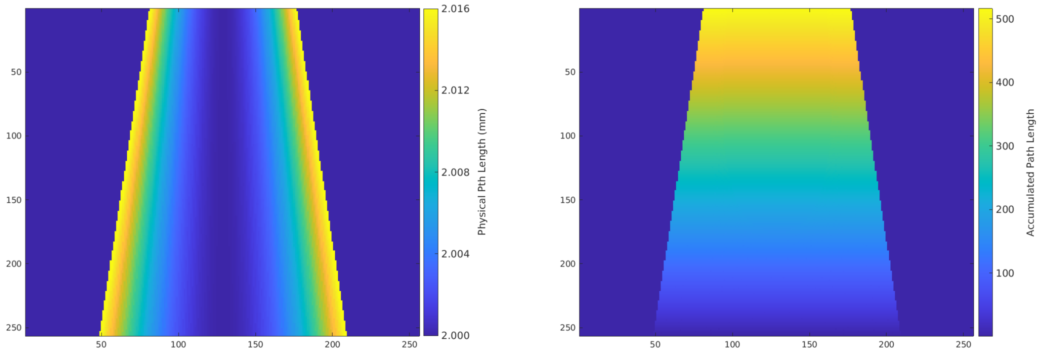

(c) Intersection length, after correction **(d)** Cumulative intersection length, after correction

Figure 5.3: (a) Central slice of the intersection length and (b) cumulative intersection length matrices inside the water matrix, before correction, when 410 × 410 rays are considered. The Moiré artifacts are clearly visible in both cases. (c) and (d) are the intersection length and cumulative intersection length after correction, respectively.

5.1.4.3 Field-size energy-dependent linear attenuation coefficient in water

Once the radiological path length has been calculated with the ray tracing algorithm, the appropriate value for the linear attenuation coefficient μ needs to be determined, to properly account for the attenuation of the beam when passing through the patient CT. μ is defined for a monoenergetic beam crossing a certain material. As the radiological path length is given in terms of water, the linear attenuation coefficient of water μ_w was considered. However, the beam produced by a clinical linac is polyenergetic and the energy spectrum depends on the field size, due to scattering radiation produced by the MLC and other parts of the linac [203]. The mean energy of the beam from each one of the MC-simulated 125 IMRT sub-fields was then used to produce field-size energy-dependent $\mu_{w,\overline{E_f}}$.

In the research conducted by J. Fleckenstein [132], a Geant4 model of an Elekta Synergy linac, coupled to a MLCi^{TM} MLC (Elekta Oncology Systems, Crawley, UK), was developed and used to simulate several squared fields. A PhSp file was created below the linac head to score, among others, the energy of the particles exiting the linac, and the results were used to determine the mean energy of the beams for different field areas. As the majority of the particles scored in the PhSp correspond to photons (over 95%), the mean energy of the beam was calculated as the mean energy of the photons. Despite the different MLC model used, the results obtained by Fleckenstein and presented in Table 5.4 were used in this work to find a relationship between field area (FA in mm^2) and mean energy $\overline{E_f}$ (in MeV). The data was fitted with a double exponential function (Eq. (5.5), Fig. 5.4), and can be used as a mapping to estimate the mean energy of the beam $\overline{E_f}$ for different fields f.

Table 5.4: Field area FA (in mm^2) and mean energy $\overline{E_f}$ (in MeV) of particles exiting the Elekta Synergy linac, coupled with an Elekta MLCi MLC [132].

Field area FA (mm^2)	Mean energy $\overline{E_f}$ (MeV)
100	1.92
400	1.92
900	1.91
2500	1.88
10000	1.80
40000	1.66

$$\overline{E_f} = 0.114e^{-0.00014 \cdot FA} + 1.811e^{-2.184.10^{-6} \cdot FA} \tag{5.5}$$

For each simulated EPID signal, the irradiated area was determined by the number of irradiated voxels, considering only the voxels receiving more than 50% of the maximum signal value. As each voxel has a surface of 1 mm^2, the sum of these voxels results in the irradiated area in mm^2. Finally, the MLC field area corresponds to the field at the isocenter, so the EPID irradiated area was scaled to the isocentric value.

After estimating the mean energy of the beam, the corresponding $\mu_{w,\overline{E_f}}$ needed to be calculated. Several mass attenuation coefficient values for water $(\mu/\rho)_{w,E}$ (in cm^2/g)

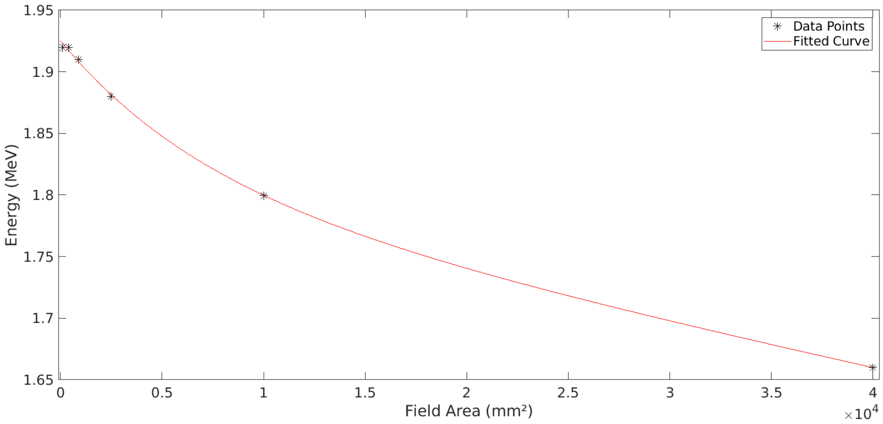

Figure 5.4: Double exponential fit between field area FA (mm^2) and mean energy $\overline{E_f}$ (MeV) of photons exiting the linac head, used as a conversion curve to obtain the field-size dependent mean energy of the beam.

at different beam energies were gathered from the National Institute of Standards and Technology (NIST) database [204], and are listed in Table 5.5. Considering the density of water $\rho_w = 1g/cm^3$, $\mu_{w,E}$ can be obtained in cm^{-1}. As the ray tracing algorithm outputs the radiological path in mm, the $\mu_{w,E}$ values were scaled to mm^{-1}. This gathered data were fitted using a double-exponential function, and the fitted curve was used as a mapping function between $\overline{E_f}$ and $\mu_{w,\overline{E_f}}$ (Eq. 5.6, Fig. 5.5). Finally, using Eqs. (5.5) and (5.6), $\mu_{w,\overline{E_f}}$ can be estimated for different field areas with corresponding field-size dependent beam energies.

Table 5.5: Beam energy and respective mass attenuation coefficients values for water $(\mu/\rho)_{w,E}$ obtained from NIST [204]. The energy-dependent linear attenuation coefficients $\mu_{w,E}$ (mm^{-1}) are calculated based on $(\mu/\rho)_{w,E}$ and the density of water ($\rho_w = 1g/cm^3$).

Beam energy (MeV)	$(\mu/\rho)_{w,E}$ (cm^2/g)	$\mu_{w,E}$ (mm^{-1})
0.60	0.0896	0.00896
0.80	0.0787	0.00787
1.00	0.070	0.0070
1.25	0.0632	0.00632
1.50	0.0575	0.00575
2.00	0.0494	0.00494
3.00	0.0397	0.00397
4.00	0.0340	0.00340

$$\mu_{w,\overline{E_f}} = 0.0082e^{-1.39\cdot\overline{E_f}} + 0.0058e^{-0.14\cdot\overline{E_f}} \tag{5.6}$$

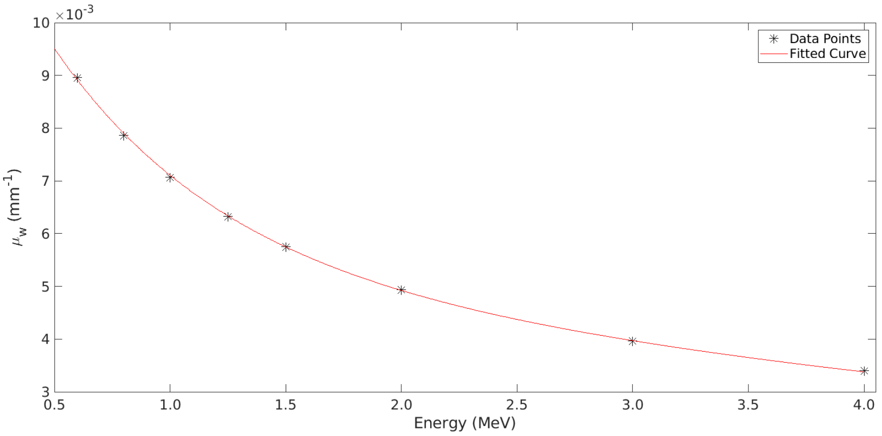

Figure 5.5: Double exponential fit between $\mu_{w,E}$ (mm^{-1}) and beam energy (MeV), used as a conversion curve to obtain the field-size energy-dependent linear attenuation coefficients of water $\mu_{w,\overline{E_f}}$ for the simulated IMRT sub-fields.

5.1.4.4 FOD Summary

All the previously described steps led to the calculation of the FODs, which can be defined by Eq. 5.7. First, the 2D EPID signal was backprojected into several parallel planes d towards the radiation source, creating a 3D matrix of dose-to-water distribution ($D_{w,d}$ on Eq. 5.7). This step was performed using the described 3D ray-tracing algorithm, which accounted for the geometric magnification at depth d. Next, the intensity of the signal at every depth was corrected to the inverse square law effect (second term on Eq. 5.7), where r_d and r_{EPID} are the distances from the radiation source to the reconstruction depth d and the EPID scoring plane, respectively. The isocenter is positioned at 1000 mm from the radiation source, while the EPID has a SDD of 1600 mm (r_{EPID} = 1600 mm). As all the patient CTs used in this work (subsection 5.1.1) extend over 512 mm in the antero-posterior direction and are positioned at the isocenter, $r_d \in \mathbb{R} : [745, 1255]\ mm$.

$$\mathrm{FOD}_d = D_{w,d} \cdot \left(\frac{r_d}{r_{\mathrm{EPID}}}\right)^{-2} \cdot e^{(\mu_{w,\overline{E_f}} \cdot L)} \tag{5.7}$$

Finally, the 3D dose-to-water distribution, already corrected for magnification and intensity, was also corrected for the attenuation inside the patient CT (last term on Eq. 5.7), with $\mu_{w,\overline{E_f}}$ determined for each field size and energy with Eqs. 5.5 and 5.6. L is the cumulative intersection length matrix of the patient CT, calculated with the ray tracing algorithm, describing the radiological path traveled by each ray, from any voxel inside the virtual patient into the EPID. FOD$_d$ is the first order dose approximation at reconstruction depth d inside the virtual CT. The resultant FOD for all depths d was a $256 \times 256 \times 96$ matrix, with $2 \times 2 \times 2$ mm^2 voxel size, to be consistent with the ADD and CT datasets. The patient mask was applied to the resultant FOD matrix, setting the dose values to zero at voxels outside the patient.

5.1.5 Datasets

A total of 581 ADD-FOD pairs were produced using the 83 different patient CTs, corresponding to 7 fields per patient, all fields simulated with the gantry at 0°. This dataset was divided into a training set, composed of 67 patients (469 fields), and a test set, with the remaining 16 patients (112 fields). The training set was used for training the network, while the test set was an independent set used to evaluate the performance of the final trained network parameters with unknown data. This helps verifying that the network was not over-fitted to the training set and whether it could be generalized to other examples not used in the training process. Additionally, the lateral dataset, composed of eight ADD-FOD pairs of fields simulated with the gantry at 90°, using eight different patient CTs, was used to evaluate the performance of the trained network in predicting dose distributions from fields coming from a different direction compared to the training scenario.

5.1.6 Deep Dose Estimation and Network Training

The Deep Dose Estimation (DDE) algorithm is a modified 3D U-Net architecture [205], represented in Figure 5.6. Its capability was already proven for dose estimation of CT scans at different patient anatomies and different setups [71, 80], and is here extended to IMRT treatment plans.

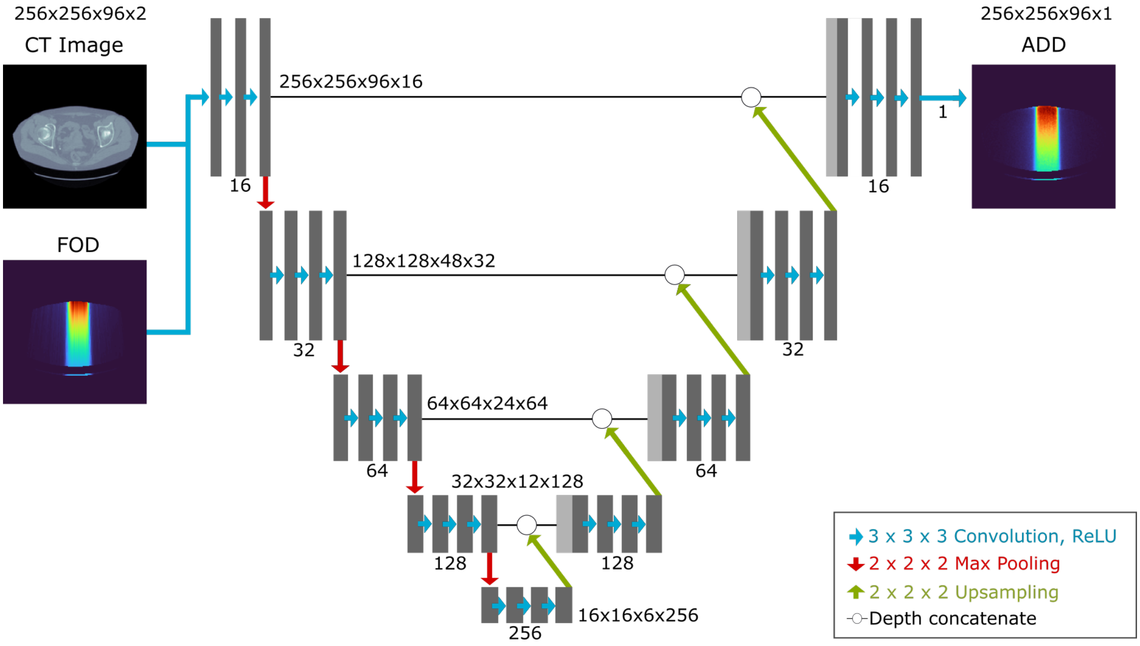

Figure 5.6: Schematic representation of the Deep Dose Estimation (DDE) architecture. The input to the 3D U-Net is given as a two-channel volume, consisting of the first-order dose approximation (FOD) and a scaled CT image. The respective MC accurate dose distribution (ADD) is given as the target data. The size of the output matrix at each convolution stage and the respective number of filters used are also displayed. Reprinted with permission from [197].

The DDE network was implemented using Tensorflow/Keras framework and trained using the training set. The training was performed on an Nvidia Quadro P5000 GPU for 500 epochs, using the Adam optimizer, a batch size of 1 and an initial learning rate of 0.0001. The learning rate was progressively reduced by 25% every 50 epochs. The initial learning rate and the reduction strategy was chosen based on previous experiences with the DDE network [71, 80].

The FOD and patient CT matrices were given as a two-channel input to the network (256 × 256 × 96 × 2), while the respective ADD (256 × 256 × 96 × 1) was the training target. The input CT data were scaled according to Eq. 5.8, so that all the input data, namely CT, ADD and FOD, were in the same order of magnitude. This step is known as feature scaling, and makes the training process more stable [119]. The CT provides the network with patient anatomic information, while the FOD contains intrinsic information regarding the irradiated field and respective deposited dose inside the patient CT, and the ADD is the desired dose distribution, which the DDE should learn to reproduce.

$$\mathrm{CT}_{scaled} = \left(\frac{\mathrm{CT} + 1000\ \mathrm{HU}}{2000\ \mathrm{HU}}\right) \tag{5.8}$$

The mean absolute percentage error (MAPE) between the DDE predicted dose distribution (DDEP) and the ADD was used as the loss function. Therefore, during the training process, the network parameters were optimized towards the minimization of the MAPE. The patients' masks were fed into the DDE, to assure that the network was optimizing the dose distribution inside the patient, i.e. that only voxels within the patient were used to calculate the MAPE. The masks set the cost values outside the patients to zero, and these voxels therefore did not interfere on the optimization of the loss function.

5.1.7 Evaluation

The test dataset, composed of dose distribution from fields incident at patient antero-posterior direction, similar to the data used for training, was used to evaluate the performance of the trained DDE. The CTs and FODs of the test set were fed into the network, together with the final trained parameters. The output of the network, i.e. the DDEP, was compared to the AADs of the test set, and the quality of the predictions was assessed using a modified implementation of the gamma evaluation method [206], able to account for under- and overestimation of the dose, as described in Chapter 3. A maximum dose difference of 3% and a 3 mm distance-to-agreement criteria (3%, 3 mm) are recommended for 3D dose distributions in pre-treatment IMRT plan verification in a clinical routine [13, 207]. To have a more sensitive evaluation, the distance-to-agreement criterion was reduced to 2 mm, while the dose difference criterion was kept at 3%. These values were limited by the level of statistical uncertainty in the ADD distributions (2.8%) and by the voxel sizes of the CT, FOD, ADD and DDEP matrices (2 mm). The gamma passing rates, i.e. the percentage of voxels passing the given criteria, were calculated using only voxels which scored at least 20% of the field's maximum dose. The passing rates reflected the agreement between predictions and respective target ADDs, and consequently the performance of the DDE in predicting dose distributions inside a patient geometry with MC-like accuracy, starting from a FOD and a patient CT. The gamma evaluation was also performed for the FOD and respective ADD of the test set, to assess the improvement obtained with the DDE. The same evaluation was performed for the lateral set of fields simulated with the gantry at 90°.

5.2 Results

The network training process in the presented configuration took approximately 47 hours. To evaluate if the training was proceeding as expected, i.e. in the direction of minimizing the cost, the cost values for both training and test sets were recorded every 10 epochs (Fig. 5.7). The initial cost of 5.64% and 2.77% for the training and test set, respectively, dropped to under 1% for both sets after 50 epochs. The final cost values for the training and test sets after the 500 epochs were 0.62% and 0.68%, respectively. As expected, the final training cost value is smaller, since the network learns the parameters based on the data of the training set. However, the final test cost is very close to the final training cost, demonstrating that the network is able to work on generalized data that were not present in the training process. This is a good indication that the network parameters were not over-fitted to the training data and learned how to work on generic data, provided they are similar to the training data set. Moreover, the trained network predicted the DDEPs within 0.6 s per field/patient combination, using the architecture described in 5.1.6. In contrast, the MC simulations took approximately 14 hours per CP.

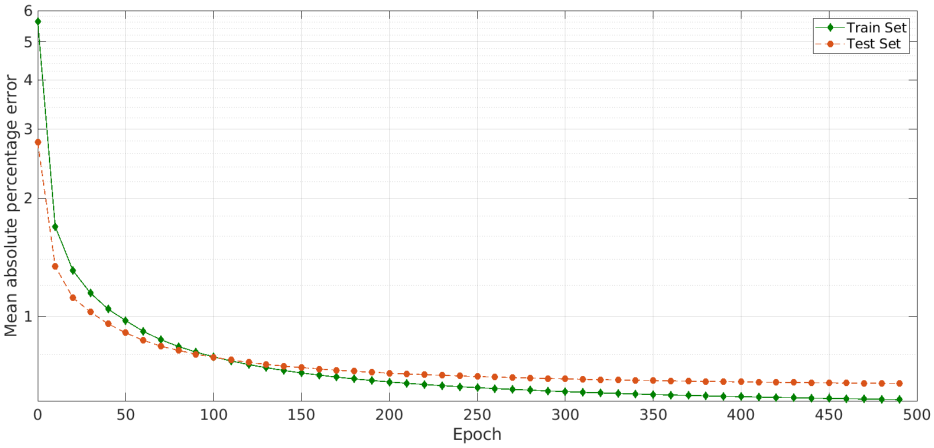

Figure 5.7: Cost values (MAPE) from the training and test sets, along the different epochs of the training process. The vertical axis is in logarithmic scale.

The gamma passing rates for the test dataset, for both FOD and DDEP with respect to the ADD, are presented in Fig. 5.8. In general, a clear improvement was observed on the passing rates when comparing the initial FOD and the DDEP distributions with the reference ADD. For the FOD distributions, the passing rates were as low as 46% for some fields, while all the DDEPs had passing rates above 97%. The maximum dose values of each ADD distribution ("Max. Dose (Gy)") are presented in the central row. The mean CT number ("Mean CT N. (HU)"), averaged over all voxels inside the patient CT which scored $\geqq$ 20% of the maximum dose, are shown in the lower graphic. A clear correlation between the passing rates and the maximum dose values was observed, while no correlation with the mean CT number could be identified.

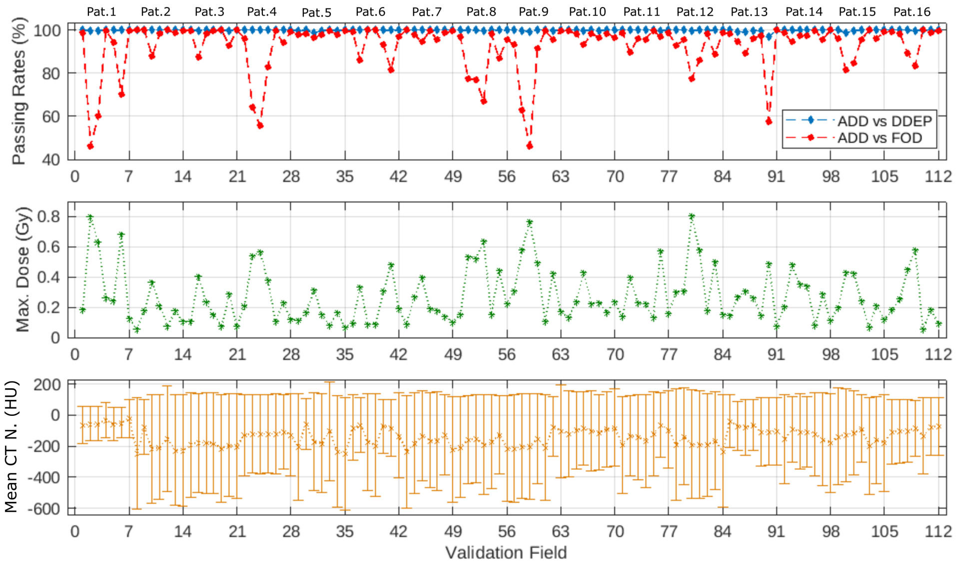

Figure 5.8: Gamma passing rates (3%, 2mm) for the fields in the test dataset, for ADD vs. FOD and ADD vs. DDEP (top). The maximum dose values of the ADDs (center) and the mean CT number value (with standard deviation) of the voxels scoring at least 20% of the maximum dose (bottom) are also presented. A clear correlation is observed for the FOD passing rates and the maximum dose values. No obvious correlation can be identified between the mean CT number values and the passing rates for both FODs and DDEPs. Reprinted with permission from [197].

Fig. 5.9 shows an exemplary field of one patient from the test set. The ADD, FOD, DDEP (superimposed over the patient CT) and respective gamma distributions are shown in different planes for a qualitative evaluation. Voxels represented in green correspond to voxels passing the (3%, 2 mm) criteria, while voxels represented in blue and red fail the criteria by underdose or overdosage with respect to the reference values, respectively. Additionally, Fig. 5.10 shows the dose profiles through patient's antero-posterior direction, sampled through the CAX and crossing the isocenter, for the ADD, FOD and DDEP dose distributions, for the same exemplary case.

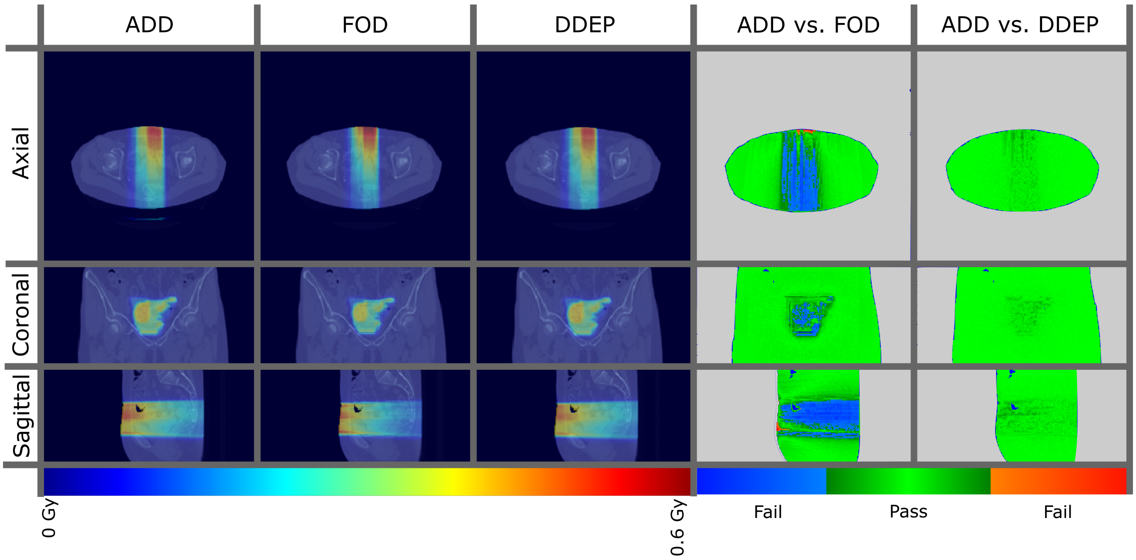

Figure 5.9: The ADD, FOD and DDEP for an exemplary patient irradiated with the gantry at 0°. The respective gamma distributions for ADD vs. FOD and ADD vs. DDEP are also displayed, where only voxels inside the patient anatomy are considered. At the gamma distributions, voxels represented in green correspond to voxels passing the (3%, 2 mm) criteria. Red and blue voxels are voxels which did not satisfy the criteria by an over or underdosage with respect to the reference ADD, respectively. The DDEP is more closely related to the ADD than the FOD, as can be seen by the increased number of green voxels in the gamma distributions. Reprinted with permission from [197].

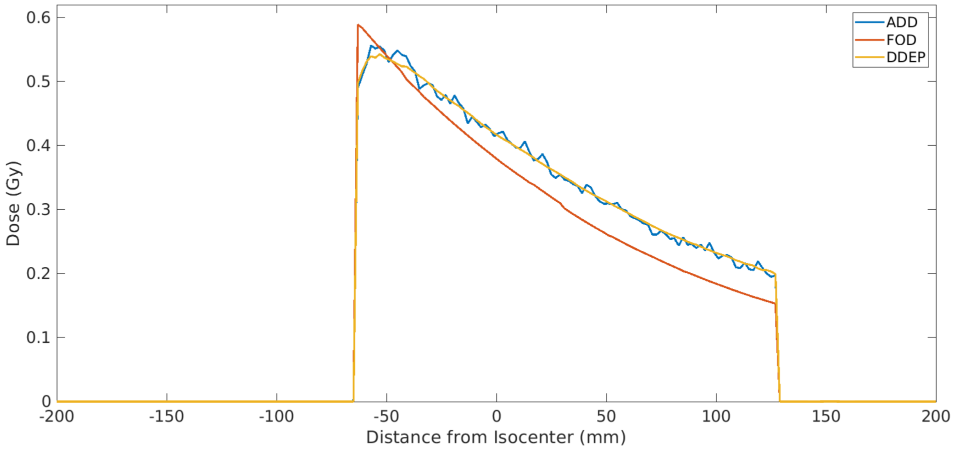

Figure 5.10: Profiles along the CAX and crossing the isocenter, through the patient antero-posterior direction, for ADD, FOD and DDEP dose distributions, for a simulation performed with the gantry at 0°. Reprinted with permission from [197].

Regarding the fields from the lateral dataset, a clear improvement on the passing rates can also be observed when comparing FODs and DDEPs with the reference ADDs (Fig. 5.11), although less pronounced than the improvements obtained for the test dataset. The lowest FOD passing rate was 88% (field 2), while the DDEP passing rate for the same field was 95%. Similar to the test dataset, a correlation between FOD passing rates and maximum ADD dose values is observed, while no correlation with the mean CT number was identified. Fig. 5.12 shows the ADD, FOD, DDEP and respective gamma distributions for one patient from the lateral dataset, and Fig. 5.13 shows the dose profiles through the patient left-right direction, sampled through the CAX and crossing the isocenter, for the ADD, FOD and DDEP dose distributions.

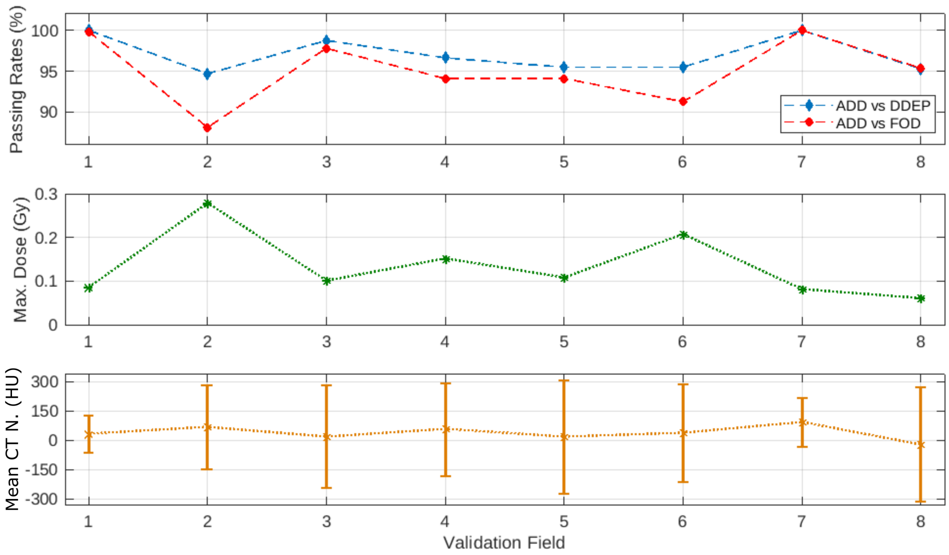

Figure 5.11: Gamma passing rates (3%, 2mm) for the fields in the lateral dataset, for ADD vs. FOD and ADD vs. DDEP (top). The maximum dose values of the ADDs (center) and the mean CT number value (with standard deviation) of the voxels scoring at least 20% of the maximum dose (bottom) are also presented. A clear correlation is observed for the FOD passing rates and the maximum dose values. No obvious correlation can be identified between the mean CT number values and the passing rates for both FODs and DDEPs. Reprinted with permission from [197].

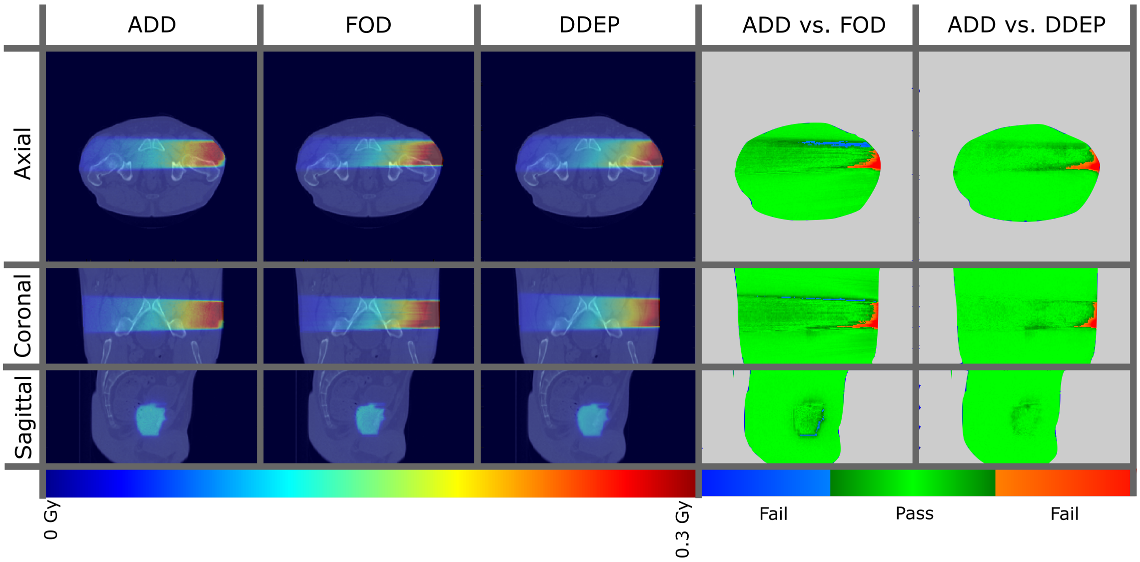

Figure 5.12: The ADD, FOD and DDEP for an exemplary patient irradiated with gantry at 90°. The respective gamma distributions for ADD vs. FOD and ADD vs. DDEP are also displayed, where only voxels inside the patient anatomy are considered. At the gamma distributions, voxels represented in green correspond to voxels passing the (3%, 2 mm) criteria. Red and blue voxels are voxels which did not satisfy the criteria by an over or underdosage with respect to the reference ADD, respectively. The DDEP is more closely related to the ADD than the FOD, as can be seen by the increased number of green voxels in the gamma distributions. Reprinted with permission from [197].

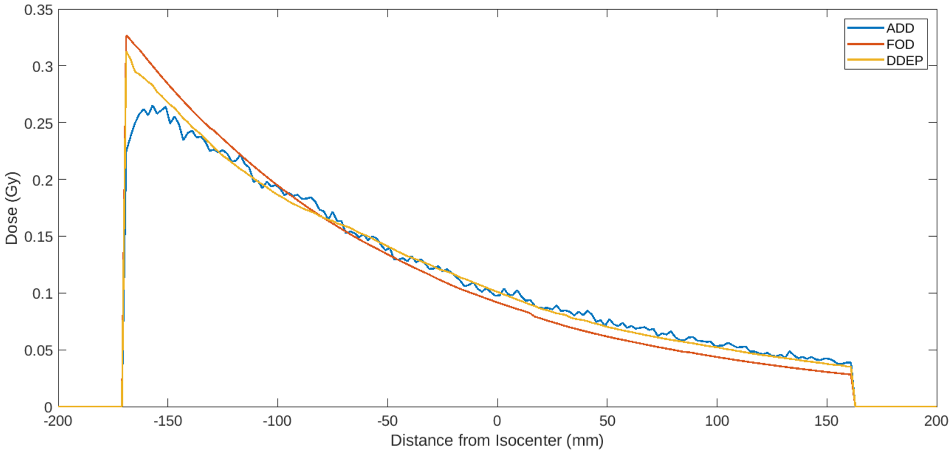

Figure 5.13: Profiles along the central axis through the patient left-right direction, for ADD, FOD and DDEP dose distributions, for a simulation performed with the gantry at 90°. Reprinted with permission from [197].

5.3 Discussion

This study presents a proof-of-principle of a methodology for EPID-based in vivo dosimetry which combines the accuracy of Monte Carlo methods to the efficiency of convolutional neural networks. The DDE algorithm, originally developed for patient-specific dose estimation in radiological CT examinations [71, 80], has been extended to dose estimation for clinical radiotherapy treatments. A patient CT image and a first order dose estimation were given as input to the network, which was trained to map this two-channel input to accurate dose distributions simulated with Monte Carlo.

The FOD, as described by Eq. 5.7, does not include corrections for the build-up effect, observed at the entrance of the patient, nor scattering effects arising from the entire patient volume. Furthermore, the $\mu_{w,E}$ values were obtained considering the mean energy of the beams in air for different field sizes, not accounting for the beam hardening effects that occur when a patient is present. This is especially relevant for the fields irradiated at 90°, since the beam goes through a thicker and more dense anatomy than the beam irradiated at 0°, due to the increased presence of higher density voxels from the pelvic bones. In MC methods, on the other hand, all these effects are very well described. Therefore, as observed in Figs. 5.10 and 5.13, the FOD dose profiles considerably differ from the ADD profiles for both fields from the test and lateral datasets, respectively, where the FOD dose values are higher at the entrance of the patient and lower at deeper depths. This was also reflected by the red and blue voxels in the gamma distribution of "ADD vs. FOD" (Figs. 5.9 and 5.12, fourth column), and by the lower gamma passing rates for the FODs (Figs. 5.8 and 5.11). The DDEP distribution, on the other hand, agrees with the ADD for the exemplary patient in the test set throughout the patient volume (Fig. 5.10), which is also reflected by the increased number of green voxels in the gamma distribution "ADD vs. DDEP" (Fig. 5.9, fifth column) and by the increased passing rates for the DDEPs (Fig. 5.9). This demonstrates the ability of the DDE network to reproduce MC-like dose distributions for the simulations at 0°, starting from a simple FOD, accounting for build-up, scattering and beam hardening effects.

Additionally, the mean energies of different fields used to determine $\mu_{w,E}$ were obtained based on the results from Fleckenstein (2013) [132], where the Elekta Synergy head was coupled to an Elekta MLCi MLC, which might have introduced uncertainties in the FOD estimation. The same linac head model was considered in this work, however the Elekta Agility MLC was used instead. Due to many similarities among both MLC models in terms of material composition and density, the values obtained by Fleckenstein can be used as a good approximation for the field-size energy dependence. However, better results could potentially be obtained if the mean energies for different field sizes were calculated using the in-house validated Geant4 model of the Elekta Synergy with Elekta Agility MLC. Despite all this, the DDE was able to properly reproduce the ADD using as input the FOD as it is, and the resultant DDEPs of the test set properly account for all the missing effects from the FODs.

In general, a clear improvement is observed in the gamma passing rates for the test set, when comparing the initial FOD and the DDEP distributions with the reference ADD (Fig. 5.9). For the FOD distributions, the passing rates were as low as 46% for two fields, and lower than 80% for 12 fields; the passing rates for the DDEP increased considerably, and were above 97% for all fields. The FOD passing rates were above 95% in some cases, especially for smaller fields, where the majority of voxels are outside the primary beam (open field) region, scoring very low dose values coming from scattering and/or transmission through the MLC. This can potentially lead to an overestimation of the gamma passing rate, since most of the low-dose voxels from FOD will still pass the (3%, 2 mm) criteria with respect to the ADD. This bias can be reduced by defining a dose threshold to select only voxels receiving at least a certain fraction of the maximum dose. In this work, a 20% dose threshold was chosen. Higher values can be used if the user is interested in high dose regions, however some voxels that are on the edge of the primary beam region could be excluded.

For some patients in the test set, the FOD passing rates were notable lower (patients 1, 4, 8, 9, 12 and 13 in Fig. 5.8), but no special trend was observed on these patients' CT images regarding gender, presence of metal and/or couch: no metal was present in any of them, the couch was present on the CTs of patients 1, 4, 9 and 13, two of them were male patients (1 and 4), and the rest female (8, 9, 12, 13). However, all cases with FOD passing rates below 80% correspond to linearly combined fields created for data augmentation with Eqs. 5.1 and 5.2. This suggests that the limitations of the FOD, in the way defined in this work, were more pronounced at bigger and more complex fields, as the discrepancies from individual fields accumulated when combined to each other. This justifies the observed correlation between lower FOD passing rates and higher ADD maximum values, since the linearly combined fields were simply the summation of two or three individual fields and had therefore higher dose values. Moreover, IMRT and VMAT both exhibit treatment plans comprised of many sub-fields combined and irradiated at different gantry angles. The good performance of the DDE for linearly-combined fields, despite the initial poor FODs, suggests that the DDE will be able to properly predict the dose distributions due to such complex treatment plans.

Regarding the fields in the lateral dataset, it can be observed that the DDEP was more closely related to the reference ADD than the FOD, as reflected by the higher passing rates of DDEPs in comparison to those of FODs (Fig. 5.11) and by the dose profiles (Fig. 5.13). However, both the FOD and DDEP dose profiles were higher than the ADD profile, which could be an indication that the DDE was unable to properly account for the build-up effect for the exemplary patient from the lateral dataset. This can be justified by the absence of simulations at 90° in the dataset used for training the network. Even though the improvements on the lateral dataset were less pronounced than for the test dataset, they were not negligible and indicate the ability of the DDE to generalize to irradiation at other gantry angles. The DDE network was trained using exclusively simulations with the gantry at 0° and still performed well with the lateral dataset, which present higher heterogeneities, as reflected by the higher mean CT values in Fig. 5.11. Nevertheless, the performance of the DDE would certainly benefit from the inclusion of irradiation from different directions and different anatomical regions in the training process.

For the simulations performed with the gantry at 0°, the beam goes through a fairly homogeneous region of the pelvis, mostly between the pelvic bones, as illustrated in Fig. 5.9. This is reflected by the mean CT values of the voxels in Fig. 5.8. For simulations with gantry at 90°, the beam goes through the pelvic bones (Fig. 5.12), and the amount of higher density voxels, i.e. voxels with higher CT values, increases (Fig. 5.11). Even though the pelvis region is fairly homogeneous, the presence of higher density voxels is not negligible. Moreover, the DDE performed well in the regions of density inhomogeneities, i.e. in the boundaries between soft tissue and bones, as can be observed from the "ADD vs. DDEP" gamma distributions for both the test and lateral dataset patients (Figs. 5.9 and 5.12, respectively). Furthermore, no correlation was observed between the gamma passing rates and the mean and standard deviation of the CT values of the voxels scoring at least 20% of the maximum dose (Figs. 5.8 and 5.11), suggesting that the DDE performance, in the scope of this work, does not depend on the heterogeneity of the patient geometry. Finally, it has been previously proved that the DDE is able to estimate dose distributions for clinical CT scans acquired with different tube voltages and at different anatomical sites [71]. With a proper training set, the DDE can also be extended to accommodate a broader range of clinical scenarios for dose estimation in radiotherapy.

The FOD distribution is not noisy, in contrast to the ADD reference (Figs. 5.10 and 5.13), even though it is generated with the EPID simulated signal, which is also subject to the statistical noise inherent from MC. When the 2D simulated EPID signal is backprojected onto a 3D matrix using the ray tracing algorithm, the resultant distribution needs to be corrected for Moiré artifacts, commonly observed in ray-driven backprojections. These artifacts were corrected by normalizing the matrix by the sum of the intersections of each ray with each voxel. As a result, the backprojected distribution is smoothed and so is the subsequent FOD distribution. Since the FOD is one of the inputs for the DDE, the predicted dose DDEP agrees with the ADD distribution, however without the statistical noise present in the ADD.

Regarding run time, the DDE clearly outperformed MC simulations. For the proposed DDE network, the training process took approximately two days on a Nvidia Quadro P5000 GPU, for 500 epochs and 469 training examples. But once the network was trained, it predicted the dose distribution inside the CT within 0.6 s for each field, in a 256 × 256 × 96 voxel volume. Thus, the dose delivered to a patient due to a clinical IMRT prostate plan, with approximately 60 CPs as the ones used in this work, can be predicted in approximately 36 s, which is clinically viable. Furthermore, a meaningful reduction on training and prediction times is expected if the network is implemented in the newest GPU models available. In contrast, the MC simulations took around 14 hours per CP.

Finally, as the training targets used are 3D dose-to-medium distributions inside patient CTs, the DDEPs are also given in dose-to-medium, properly accounting for the inhomogeneities inside the patient. If preferred, DDEPs in terms of dose-to-water can be obtained, provided the network is trained with dose-to-water target distributions.

5.4 Conclusion

This study shows the feasibility of training a deep neural network (DDE), originally developed for radiological applications, to predict the dose delivered to patients for IMRT treatments, using as input an approximated dose distribution (FOD), reconstructed from the signal recorded by the EPID, and patients' CT images. The FOD does not account for several effects, such as beam hardening, build-up and scatering inside the patient CT, which are all properly described in the target Monte Carlo dose distribution (ADD). After training, the DDE was able to reproduce the ADDs for the patients in the test set (gantry at 0°), and the resultant predictions (DDEPs) properly account for all the missing effects from FODs. For patients in the lateral set (gantry at 90°), the DDE also performed well, however the improvement was less pronounced, justified by the absence in the training dataset of fields simulated at different gantry angles. Even though this is a simulation study, it shows the viability of the proposed method as an EPID-based in vivo dosimetry approach. Furthermore, it can potentially predict the dose received by the patient at each fraction or even for each CP within seconds, with MC accuracy. The proposed method is thus time-efficient and does not increase the workload on the personnel, making it clinically applicable.

CHAPTER 6

General Discussion and Outlook

This study presents a methodology for EPID-based in vivo dosimetry combining the accuracy of Monte Carlo methods to the efficiency of convolutional neural networks. The DDE network, originally developed for patient-specific dose estimation in radiological CT imaging [71, 80], was extended to dose estimation for clinical radiotherapy treatments. A patient CT image and a first-order dose (FOD) estimation are given as input to the network, which is trained to map this two-channel input to an accurate dose distribution (ADD) inside the same patient CT, simulated with Monte Carlo (MC).

The quality of the DDE predictions (DDEP) strongly depends on the quality of the target dose used for training the network: if the MC model of the linac used to simulate the ADDs does not represent the linac in the clinic, neither will the predictions obtained by the trained DDE. Therefore, a reliable, validated MC model of the Elekta Synegy and the coupled Agility MLC system was crucial for the development of this project.

To produce a reliable MC model, the geometry and the material composition of all components of the linac need to be faithfully described. To simplify the development of the MC model, the linac was divided in two parts: the static patient-independent part, and the dynamic patient-dependent part, each part being modeled separately. This division enabled the production of PhSp files using the static model, as described in details in Chapter 3. Such PhSp files contain relevant information of the particles crossing the PhSp plane, and can be used as primary particle source for the patient-dependent model, replacing the static components of the linac head. PhSp files are very useful to make the simulation more efficient. First, because the production of X-rays in the beam target is a very inefficient process where only a small portion of the electron beam energy leads to X-ray production, while the rest is converted into heat at the target [85]; second, because the particles transport through all the static components of the linac head do not need to be repeatedly simulated every time, a time-consuming process which can be avoided by replacing the static part of the linac with PhSp files.

Using the information provided by the vendor after the establishment of a non-disclosure agreement, a detailed MC model of the static components of the Elekta Synergy linac was produced using Geant4. The dynamic part (or patient-dependent part) was modeled based solely on information found after an extensive literature review. Details are provided in Chapter 3. Even with the geometric information provided by the vendor, modeling the static part of the linac head was a time consuming process. The parameters of the electron beam incident on the target needed to be optimized, namely the energy and the spatial distribution of the beam (focal spot). By the time of this optimization, a consensus on the proper methodology was not found in literature. Therefore, PDD profiles of different squared fields were initially used for the energy optimization, while lateral profiles of squared fields at different depths in a water phantom were used as guidance for the focal spot optimization. It was observed that the PDD profiles are not sufficient for the energy optimization, as the lateral profiles of bigger fields have shown a clear dependence on the beam energy, especially at the off-axis region. For future electron beam energy optimization, PDD profiles should be used in combination with at least one set of inline and crossline profiles for a bigger field, for example 30 × 30 cm^2. On the other hand, inline and crossline lateral profiles at two different depths in water for several squared fields have served as a good guidance for the focal spot optimization, and the penumbra region was especially useful for the angular spread optimization of the electron beam. Even with the available information provided by the vendor, the optimized beam energy and focal spot parameters obtained were significantly different from the nominal values provided by the vendor, showing that these values are equipment-specific and should be obtained for each equipment individually. Finally, in this work, only the 6 MV photon beam of the Elekta Synergy was modeled. Other beam energies available for the Synergy linac could be modeled following the method described in Chapter 3. The methodology is general and could also be extended for other linac models. The final optimized set of electron beam parameters was used to produce PhSp files, recorded between the flattening filter and the monitor chamber arrays, and used as primary particle source for the patient-dependent model.

The dynamic patient-dependent part of the linac head, the Elekta Agility MLC system, was modeled using solely information found during an extensive literature review. For some parameters, namely the diaphragms length and width along with the T&G width, no value was found in the literature and they had to be calculated or simply tested until a reasonable value was obtained. Moreover, for some parameters, several different values were encountered in the literature, confirming that due to differences in installation and settings, some parameters are machine-specific, and similar to the electron beam parameters, they had to be optimized to reproduce the equipment used in this work. The optimized MLC system parameters were the leaf's height, the density of the tungsten alloy (used for both leaves and diaphragms modeling) and the shift of the MLC bank to the virtual focus. The optimization was once more an iterative process, which used several measurements as reference. The values of each parameter were changed individually and the resultant simulations were compared to the reference measurements until a good visual agreement was found between measurements and simulations.

Once both static and dynamic models were finalized, they were extensively validated using several strategies which verified all the different parts of the MC model: electron beam parameters, MLC geometry and positioning, construction of the patient geometry (in this case, phantom) inside the Geant4 model, absolute dose calibration, uncertainty estimation and reproducibility of IMRT fields in the MC model, extracted from RTPlan.dcm files and implemented in Geant4 through user-defined commands. The electron beam and MLC system parameters were optimized using a qualitative visual evaluation with respect to reference measurements. Some steps of the validation process also relied on qualitative visual inspection to reference data. However, the 3D simulated dose distributions for two complex clinical IMRT plans inside the cylindrical water equivalent phantom were evaluated, with respect to reference data, using the gamma methodology, which is a well-established method for quantitative evaluation of 3D dose distributions [179]. This was the most crucial validation step and confirmed that the developed MC model can produce dose distributions that represent the linac of interest, which can therefore be used as the target doses to DDE training.

At the beginning of this PhD project, the geometric information of the linac static components was not available. An alternative method of PhSp manipulation was developed as a master thesis project, under daily supervision of the PhD applicant, as described in details in Chapter 4. The goal was to manipulate the information stored inside IAEA PhSp files, produced and validated for a different linac model, until the information was optimized and the manipulated PhSp could describe the photon beam of the Elekta Synergy used. The energy and angular direction of the particles stored inside the IAEA PhSp files were perturbed, and the resultant PhSp files were used to simulate PDD and lateral profiles of different squared fields in a water phantom. The simulated profiles were compared to reference measurements, and the process continued iteratively, following the decrease of cost values between simulations and measurements. Once a good agreement between them was achieved, as quantified by the minimum cost value, the PhSp was considered optimized. While the energy optimization was simply consisting of a change on its magnitude, the angular direction manipulation was more complicated, and involved several steps. Finally, an optimized PhSp was produced that could satisfactory describe the measured data. Since this optimization was an iterative method, it involved several steps and countless simulations, only made possible by the cropping technique introduced, in which a smaller PhSp file containing only the relevant particles in a “cropped” central region of the original PhSp was used for certain simulations, significantly reducing the simulation time needed in each iteration. This PhSp optimization methodology is a good alternative for the scenario where the user does not have the geometric information provided by the vendor. However, it is limited to changes applied at the particles’ energies and angular distributions, and might not consider other factors which influence the spectrum of the photon beam produced by a clinical linac, such as the flattening filter material and composition. Moreover, as it is an iterative method, further iterations could result in potentially better optimized PhSps, and it is up to the user to find a compromise between results and simulation time. The manipulated PhSps were used in combination with the Agility MLC model constructed based on the methodology described in Chapter 3. The combination of the PhSp optimization method and the Agility model, constructed using only information found in the literature, enables the production of a MC model of an entire linac head without the need of confidential vendor information.

In order to implement in vivo dosimetry in the clinical routine, the methodology needs to be reasonably simple, should fit into the patient's treatment routine, should be time-efficient and should not increase the workload on the staff. In the proposed methodology, the most laborious and time-consuming part is the network training, involving the collection and pre-processing of a substantial amount of patient data, the development of a MC model of the linac of interest including patient and EPID models (which itself can be difficult to produce), and a backprojection algorithm to produce an approximated dose distribution based on simulated EPID signals. The time spent on the production of FODs also needs to be considered. The ray tracing algorithm used in this work is the bottleneck for the FOD calculation, taking about 1 hour when applied to one CT image with 256 × 256 × 96 voxels. Once the ray tracing matrices are available, the FOD is calculated within 20 s per CP, given the simulated EPID image as input. It is important to mention that the entire FOD production, including the ray tracing algorithm, was performed sequentially in a Intel® Core$^{\mathrm{TM}}$ i7-6700 CPU. Possible optimization of the ray tracing algorithm and FOD calculation process, such as parallelisation and/or GPU implementations, can certainly reduce the FOD production time to less than 1 s for all CPs of the IMRT plans used. Furthermore, the training process itself can be time consuming, even when GPUs are used, however a meaningful reduction on training and prediction times is expected if the network is implemented in the newest GPU models available.

Yet, all these steps can be performed before patient irradiation, and as demonstrated, once the network is trained, it can predict dose distributions inside a patient CT image within 0.6 s for one sub-field of an IMRT treatment. Hence, for a prostate IMRT plan with approximately 60 CPs, as the ones used in this study, the total dose delivered to a patient geometry in one treatment fraction could be predicted in less than one minute, which is clinically viable. Moreover, this methodology does not require any extra step to be introduced during the clinical routine, as e.g. in-air EPID acquisition, thus not increasing the necessary time slot reserved for each patient treatment. The workload on the staff does not increase, and its implementation requires little or no extra personnel training. These are very relevant aspects to be considered, since many facilities work on a packed patient schedule and sometimes struggle with personnel shortage.

The accuracy of the resultant DDEP is limited by the accuracy of the target data used for training the DDE. The accuracy of the MC simulations can be increased if more particles are used in the simulation, potentially increasing the quality of the predicted dose.

It is up to the user to decide on a good trade-off between computational cost and accuracy of the target dose, and consequently on the accuracy of the predictions. Moreover, dose distributions from clinical TPSs could also be used as target data, in case MC models of the linac are not available or deemed too time consuming. However, clinical TPSs are usually not calibrated to calculate doses at the EPID level (SDD of 160 cm), required for the production of FODs, and corrections are needed [208]. Hence, the use of TPS for simulating the transmitted signal at the EPID is not straightforward, and was not considered sufficiently reliable for the IMRT fields used in this work. Additionally, TPS-independent verification approaches are desirable, as they can identify several problems associated with the TPSs (heterogeneity calculations, data corruption, among others) [56, 209]. Finally, dose distributions produced with MC simulations are known to provide the most accurate estimation available, and therefore MC simulations were given priority over analytical TPS dose distributions in this work, despite the clear disadvantages regarding time efficiency.

By the time of the development of this project, a big dataset of diagnostic patient CT images was available, and was therefore used as the patient geometric information throughout this work. These images are similar to planning CTs, i.e. the CT image used for treatment planning with clinical TPSs, which can easily be used as input for both the MC model of the linac, to produce the ADDs and EPID signals, as well as input to the DDE, together with the respective FODs. The planning CT however does not necessarily correspond to the patient geometry at the moment of treatment, as it is usually acquired several days in advance. The method can be extended to use cone-beam CTs (CBCT), which provide updated information of the patient geometry during the course of treatment. CBCT scanners are widely available on modern clinical linacs and are often already part of the clinical routine, with daily or weekly acquisitions used for patient positioning verification. However, the image quality of CBCTs is lower compared to clinical CTs, including some artifacts and unreliable CT number values [210]. This could introduce additional uncertainties on the FOD production, since it relies on the conversion of CT numbers into relative electron density. Many correction methods have been investigated [211–213], with deep learning approaches being able to perform such corrections within seconds [214]. If CBCT scans can be corrected within seconds and plugged into the trained network, together with FODs reconstructed from EPID images acquired during patient irradiation, a prediction of the delivered dose to the patient could be obtained quasi real-time. Finally, by combining both approaches presented here and by Maier et al.[71], the total dose received by patients throughout the entire treatment process could be estimated, including diagnostic imaging, planning CT, CBCT and treatment irradiation. This can be useful not only for treatment adaptation, but also for assessment of risk and late effects that might result from the radiotherapy treatment chain.

As already explored by other groups, EPID signals and convolutional neural networks can be used together to develop EPID-based in vivo dosimetry for the newest

MR-linacs [70, 79], which use magnetic resonance images (MRI) to monitor the patient's geometry in real-time during the irradiation. Therefore, the proposed methodology could also be potentially extended to use images from MR-linacs as geometric patient information, paving the way towards real-time in vivo dosimetry.

So far, the methodology has been applied to one anatomical region, trained with fields simulated at one gantry angle, using treatment fields for one nominal beam energy (6 MV). Further investigation is needed to determine whether the network needs to be trained for each individual anatomical region and beam energy, or if the network can be trained to account for several different scenarios. Thinking about clinical implementation, having only one DDE trained parameter set capable of dealing with all anatomical sites and clinical scenarios prevents using the wrong model in a busy clinical routine. Moreover, in many cases the tumor can extend to several different anatomical regions, and this might be a problem if the network is trained for dedicated anatomical parts. Additionally, the time-consuming task of training the network needs to be performed only once. Maier et al. [71] has demonstrated that the DDE can be trained to estimate dose distributions for clinical CT scans, with different acquisition parameters and at different anatomical regions, if a comprehensive data set is provided. Moreover, Maspero et al. [214] has proposed a single deep learning network able to correct CBCT images of different anatomical regions. Both results suggests that convolutional neural networks can be trained to accommodate several anatomical sites simultaneously. Therefore, once a more comprehensive training set is available, the DDE could also be potentially extended to accommodate a broader range of clinical scenarios for dose estimation in radiotherapy.

This study is a proof-of-principle for EPID-based 3D in vivo dosimetry, based on simulated EPID signals. To make the step towards clinical implementation, the methodology needs to be extended and tested with measured EPID data, potentially involving verification using anthropomorphic phantoms. As the proposed method uses EPID signals simulated as dose-to-water, the acquired EPID signals should also be corrected and calibrated to represent dose-to-water, following one of the several methods described in literature [32, 35, 215, 216]. In contrast to other EPID-based in vivo methodologies, no further corrections to the EPID signal should be necessary, and in-air acquisitions prior to patient irradiation are not required. Moreover, the possibility of using the EPID signal without conversion to dose-to-water can also be explored, although some corrections at the EPID level (e.g. dead pixel, pixel sensitivity, energy response) will still be periodically necessary, as these parameters can vary over time with the usage of the EPID panel [32, 148].

There is no indication that the method proposed here cannot be extended to other treatment modalities, such as Volumetric Modulated Arc Therapy (VMAT), especially since the DDE performs well for CT acquisitions, where the radiation source rotates constantly around the patient. For VMAT applications, besides potential adaptations needed in the DDE and training data, the linac MC model needs to be modified to accommodate gantry rotation simultaneously to MLC movement and beam irradiation.

Finally, the DDE has been trained to output dose distributions in terms of dose-to-medium, properly accounting for the inhomogeneities inside the patient. DDEPs in terms of dose-to-water can be obtained, provided the network is trained with dose-to-water target distributions. Nevertheless, the method is flexible and the decision to train with dose-to-medium or dose-to-water distributions is left to the user, based on the individual needs.

6.1 Outlook

As previously discussed, this project has produced promising results, with many potential extensions in sight and many exciting topics for further development.

The MC model of the Elekta Synergy coupled to the Elekta Agility MLC system, developed and described in details in Chapter 3, has been extensively validated for measurements and TPS calculations performed in water or in a cylindrical water-equivalent phantom. To validate for heterogeneous media, more complex geometries need to be used, as for example anthropomorphic phantoms. The validation can be performed either using TPS calculations as reference, of with measurements performed with detectors introduced inside or in the surface of the phantoms. For example, a head phantom can be used (RTSafe PseudoPatient$^{\mathrm{TM}}$ Prime, RTSafe, Athens, Greece), in combination with an ionization chamber, to measure point doses inside the phantom due to different squared and complex irradiation fields, and compared to simulations performed under the same conditions. Another possible approach is to use termoluminescent dosimeters (TLDs), which can be introduced inside the cavities present in the slices of an antrhopomorphic phantom (Alderson Phantom, Radiology Support Devices, Long Beach, USA) to measure point doses due to different irradiation scenarios, and compare to simulations performed at the same location and under the same conditions. Dose measurements in these scenarios are limited to point dose values. Film dosimetry could also be used to obtain 2D dose distribution. Alternatively, TPS dose calculations on planning CT images of both phantoms could be used as the reference dose distributions for 3D validations.

Throughout this project, only the 6 MV photon beam of the Elekta Synergy has been used. Despite being the most frequently used photon beam energy for patient treatment in Klinikum Grosshadern, this ultimately limits the application of the EPID-based in vivo methodology to treatment plans delivered with this single beam. The first step to expand the applicability of the proposed method to other possible clinical scenarios is to produce MC models of the other beam types and energies available at the Synergy linac, such that the required dataset to extend the DDE can be produced.

The methodology described in Chapter 3 can be used to model the other beam modalities available for this linac model. For example, photon beams with different nominal energies could be potentially modeled, provided proper changes are applied in the patient-independent MC model, such as changes in the target and flattening filter geometries, among others. To perform such changes, knowledge of geometric information of the elements of interest might be needed, involving the use of confidential data from the vendor. Once the modifications have been made, respective PhSp files can be created, optimized and validated using the same MC model of the Agility MLC system, since the components of the patient-dependent part of the linac do not change for different photon beams. This would enable the expansion and the applicability of the proposed in vivo dosimetry to irradiation with all possible beams available in the linac of interest. Furthermore, the MC model can be modified to accommodate time-dependent simulations [217]. This is a requirement for extension of the model to dynamic beam delivery modalities, such as VMAT treatment plans. Finally, the Geant4 code can be adapted to run using GPUs, with great potential to significantly decrease the simulation time [218–221]. This would considerably facilitate the production of datasets needed for training and validation of the DDE network.

During the development of this project, a database consisting of patient planning CT images with segmented structures was unavailable, and instead, diagnostic CT images were used. By using planning CT images with segmented structures, as targets volumes and organs at risk (OAR), dose volume histograms (DVH) for every structure of interest can be calculated for the ADDs and DDEPs. DVH is a representation of the dose distribution within a particular volume of interest (VOI), and has become an extensively used tool in clinical practice for treatment plan evaluation [85, 222, 223]. DVH values are also used to define prescription doses to the target and/or tolerance dose values for OARs. Combined to the gamma evaluation, the DVHs can potentially provide a more complete information of the dose received by the patient during in vivo dosimetry, helping to identify possible relevant deviations between planned and delivered dose. Depending on the identified discrepancies, the radiation oncologist can, for instance, decide whether or not the treatment plan can continue, or if a plan adaptation is necessary. This is a crucial step for implementation of adaptive radiation therapy.

The DDE should be extended to replace simulated by measured EPID data. This is essential for implementing the proposed EPID-based in vivo dosimetry methodology into clinical practice. As already discussed, the EPID measured signals should be corrected and calibrated to dose-to-water, such that the measured and corrected EPID signal can be directly used for FOD calculation, using the method described in Chapter 5. To perform the EPID measurements, phantoms can be irradiated on the linac with different fields, while the EPID signals are acquired simultaneously. Different phantoms could be used, but ideally heterogeneous and patient-like phantoms represent a more realistic scenario. CT images of the phantoms, either planning and/or CBCT, should be performed and used for ADD and FOD production, as well as input to the DDE. Ultimately, point dose or dose distribution measurements performed during phantom irradiation can offer an additional validation for the DDE network.

Another interesting idea is to determine how sensitive the DDEPs are to changes in the irradiation scenario. For instance, the CT image used for FOD production can be intentionally perturbed in a controlled way, e.g. artificially rotated by some degrees, shifted by some millimeters and/or locally deformed. Such perturbations could be interpreted as miss-positioning or anatomical changes of patients in a clinical scenario. The perturbed CT image and the respective FOD would then be given as input to the DDE. The resultant DDEP should be compared to the target ADD, produced with the original unperturbed CT image. Using gamma evaluation and DVH values, the sensitivity of the proposed in vivo methodology to changes in the patient geometry can be investigated. This can be helpful for medical physicists and radiation oncologists on defining critical parameters that would require some kind of action, like adaptation of the treatment plan.

Finally, the implementation of CBCTs in the DDE algorithm and/or the extension to MR-linacs are the ultimate step towards a real-time in vivo dosimetry. The combination of the EPID-based in vivo dosimetry proposed here for radiation therapy, to the approach presented by Maier et al. [71], can result in a comprehensive quality assurance methodology that encompasses the entire patient treatment chain, from the diagnostic and planning CTs, to the CBCTs performed for position verification, until the treatment irradiation. This could be particularly relevant for assessment of risk and late effects.

CHAPTER 7

Summary and Conclusions

Modern treatment techniques, such as Intensity Modulated Radiation Therapy (IMRT) and Volumetric Modulated Arc Therapy (VMAT), deliver highly conformal dose distributions, aiming at improving coverage of target volumes while sparing the surrounding normal tissue and organs at risk (OARs). These requirements lead to steep dose gradients between the target and the OARs, and the treatment quality relies on highly accurate machine operation, as well as proper patient positioning during irradiation and reproducible patient anatomy on each treatment fraction. In this context, the demands for quality assurance of the clinical linear accelerators (linac) and dosimetric verification have increased, and much effort has been put into developing in vivo dosimetry methodologies that could be introduced in clinical practice, with EPID-based in vivo dosimetry appearing as a viable alternative in the past decade.

This study shows the feasibility of training a deep neural network (DDE), originally developed for radiological applications, to predict the dose delivered to patients for IMRT treatments. The network uses as input an approximated dose distribution (FOD) reconstructed from simulated EPID signals and a patient CT image, and tries to map this two-channel input to an accurate dose distribution (ADD) inside the same patient CT, produced with Monte Carlo (MC) simulations.

All the steps necessary for the implementation of the proposed methodology have been developed in this thesis. First, a reliable and detailed MC model of the linac used has been constructed based both on information available in the literature and on confidential information, provided by the vendor upon the establishment of a non-disclosure agreement. The MC model was divided in a patient-independent and a patient-dependent part. The first comprised the static parts of the linac head, and was used to produce validated PhSp files describing the photon beam generated by the linac; the latter comprised the patient-specific part, i.e. the multi-leaf collimator (MLC) system, and used the validated PhSp files as primary generators to simulate dose distributions inside virtual phantoms and patients. A virtual model of the EPID panel was also included in the patient-dependent part, to simultaneously record the transmitted signal through the virtual phantom or patient volume. Both MC models have been extensively validated and can properly describe the physical equipment considered in this work. Additionally, a methodology for optimization of existing IAEA PhSp files was presented, as a viable alternative for linac modeling when confidential vendor information is not available.

Next, a method for producing the FODs from simulated EPID signals was required. The FOD was developed using a modified version of a ray-tracing algorithm. Adaptations have been made to the algorithm to output not only the physical intersection lengths of the radiation with each voxel of the virtual patient volume, but also the radiological path lengths, calculated using the relative electron densities of each voxel's material. The final FOD calculation accounts for magnification and inverse square law corrections, and also considers the attenuation through the virtual patient model, while other effects were disregarded. Nevertheless, the goal of the FOD is to be fast and deliver an approximated dose distribution that is sufficient for the DDE network to map to a more detailed, accurate distribution. Despite the simplifications in the FOD construction, the trained DDE network was able to reproduce MC-like dose distributions using the FODs as inputs. In this scenario, the time needed to produce the FOD is more relevant than the discrepancies observed between FOD and ADD, and efforts should be focused on optimizing and reducing the FOD production time rather than on improving its quality.

Finally, the required datasets for training and validation of the DDE network, the training and validation sets respectively, were produced using the validated MC model and the FOD calculation approach. An extra dataset, the lateral set, containing fields simulated with a different gantry angle was also produced, to evaluate the performance of the trained network in predicting dose distributions from a field coming from a different direction, not used on the training process. The DDE was implemented in a GPU and trained based on previous experiences from a group of collaborators, who supported the extension of the DDE to radiation therapy treatment fields. The final trained network was validated using the validation set, and a remarkable improvement was observed in the agreement between the prediction (DDEPs) and the target (ADDs) dose distributions, when compared to the input FODs, quantified in terms of the gamma passing rates. This is evidence that the trained DDE can correctly map the two-channel FOD-CT input into the MC-accurate target dose distribution ADD. The lateral set was also given as input to the trained network. Although the resultant DDEPs from the lateral set were more comparable to the target ADDs than the input FODs, the improvements obtained were less pronounced than for the validation set, as reflected by the gamma distribution and by the dose profiles presented. This can be justified by the absence of fields from different directions in the training set. The extension of the dataset is currently in progress, and the performance of the DDE will certainly benefit from the inclusion of simulations performed at different directions and different anatomical sites.

In conclusion, the proposed method combines the accuracy of MC simulations, yielded by the DDE predictions, with the efficiency of convolutional neural networks, reflected by the prediction time of less than one second per control point. The results obtained are promising and show the viability of the method as a clinical EPID-based in vivo dosimetry approach. Several steps need to be developed and investigated in order to move towards clinical implementation. Nevertheless, the potential of the methodology in estimating the dose delivered to a patient, at each treatment fraction or even at each control point within seconds and with MC-like accuracy, has been demonstrated.

APPENDIX A

MATLAB script for RTPlan.dcm

MATLAB script to extract information from the RTPlan.dcm files and save it as .txt files, to be used as input for simulations performed with the patient-dependent MC model of the linac head.

```
RTPLAN_FILE=dicominfo('path_to_file/RTPlan_file_name.dcm'); % Load the RTPLAN
    file

format long

No_Beams=length(struct2cell(RTPLAN_FILE.BeamSequence)); % Number of beams

for i=1:No_Beams
  z ='Item_%d';
  v=i;
  Item_No=sprintf(z,v);

  TOTAL_BEAMS(i).Beam=RTPLAN_FILE.BeamSequence.(Item_No);
  TOTAL_BEAMS(i).angle=RTPLAN_FILE.BeamSequence.(Item_No).ControlPointSequence.
      Item_1.GantryAngle; % Gantry angle of each beam
  TOTAL_BEAMS(i).ISO = RTPLAN_FILE.BeamSequence.(Item_No).ControlPointSequence.
      Item_1.IsocenterPosition; % Position of isocenter of the plan
  TOTAL_BEAMS(i).No_segments=RTPLAN_FILE.BeamSequence.(Item_No).
      NumberOfControlPoints; % Number of control points of each beam
  TOTAL_BEAMS(i).Total_Beam_MU = RTPLAN_FILE.FractionGroupSequence.Item_1.
      ReferencedBeamSequence.(Item_No).BeamMeterset; % Total MU of each beam (
      sum of MUs from all CPs)

  for g=1:((TOTAL_BEAMS(i).No_segments).*0.5)
    v=g*2;
    Seg_No=sprintf(z,v);
```

```
    Seg_No_odd = sprintf(z,v-1);

    TOTAL_BEAMS(i).Seq(g).MU =double(TOTAL_BEAMS(i).Beam.ControlPointSequence.(
        Seg_No).CumulativeMetersetWeight-...
    TOTAL_BEAMS(i).Beam.ControlPointSequence.(Seg_No_odd).
        CumulativeMetersetWeight); % MU for each CP individually

  end
end

for i=1:No_Beams

  l ='MLC_B%d_%d_%d.txt'; % Format for saving the MLC position files
  j='Jaws_B%d_%d_%d.txt'; % Format for saving the Jaw position files
  k='iso_%d.txt'; % Format for saving the isocenter position files
  ang=TOTAL_BEAMS(i).angle;
  a=ang;

  MU_per_Beam{i} =[TOTAL_BEAMS(i).Total_Beam_MU];

  FileName3=sprintf(k,a);
  dlmwrite( FileName3,TOTAL_BEAMS(i).ISO) % Writing the isocenter position files

  for s=1:size(TOTAL_BEAMS(i).Seq,2)

    b=s;
    c = TOTAL_BEAMS(i).Total_Beam_MU;

    FileName=sprintf(l,i,a,b);
    FileName2=sprintf(j,i,a,b);

    dlmwrite( FileName,TOTAL_BEAMS(i).Seq(s).Leaf) % Writing the MLC position
        files
    dlmwrite( FileName2,TOTAL_BEAMS(i).Seq(s).Jaws) % Writing the Jaw position
        files

  end
end
```

APPENDIX B

List of publications and conference contributions

List of Publications

- **Juliana Cristina Martins**, Joscha Maier, Chiara Gianoli, Sebastian Neppl, George Dedes, Abdulaziz Alhazmi, Stella Veloza, Michael Reiner, Claus Belka, Marc Kachelrieß, Katia Parodi. Towards Real-Time EPID-based 3D in vivo Dosimetry for IMRT with Deep Neural Networks: a Feasibility Study. *Physica Medica 114 (2023):103148.*
- Lydia J Wilson, Wayne D Newhauser, Christopher W Schneider, Florian Kamp, Michael Reiner, **Juliana C Martins**, Guillaume Laundry, Augusto Giussani, Ralf-Peter Kapsch, Katia Parodi. Method to quickly and accurately calculate absorbed dose from therapeutic and stray photon exposures throughout the entire body in individual patients. *Medical Physics 47 (2020):2254-2266.*
- **Juliana Cristina Martins**, Rangoli Saxena, Sebastian Neppl, Abdulaziz Alhazmi, Michael Reiner, Stella Veloza, Claus Belka, Katia Parodi. Optimization of Phase Space files from clinical linear accelerators. *Physica Medica 64 (2019):54-68.*
- Abdulaziz Alhazmi, Chiara Gianoli, Sebastian Neppl, **Juliana Martins**, Stella Veloza, Mark Podesta, Frank Verhaegen, Michael Reiner, Claus Belka, Katia Parodi. A novel approach to EPID-based 3D volumetric dosimetry for IMRT and VMAT QA. *Physics in Medicine & Biology 63 (2018):115002.*

Conference Contributions

- **Oral presentation**: **Juliana Martins**, Joscha Maier, Chiara Gianoli, Abdulaziz Alhazmi, Sebastian Neppl, Michael Reiner, Claus Belka, Stella Veloza, Marc Kachelrieß, Katia Parodi. Towards real-time EPID-based 3D in-vivo dosimetry using machine learning. *3rd European Crongress of Medical Physics, 2021, online conference.*

- **"Blue Ribbon" poster**: **Juliana Martins**, Sebastian Neppl, Albulaziz Alhazmi, Jens Fleckenstein, George Dedes, Michael Reiner, Claus Belka, Katia Parodi. Analytical Modeling of the Parameters of a Multileaf Collimator for Monte Carlo Dose Simulations in Radiotherapy. *2020 Joint AAPM | COMP Meeting, 2020, online conference.*

- **Poster**: **Juliana Martins**, Sebastian Neppl, Albulaziz Alhazmi, Jens Fleckenstein, George Dedes, Michael Reiner, Claus Belka, Katia Parodi. Analytical determination of Multileaf Collimator parameters for Monte Carlo dose simulation. *Annual conference of the German Society for Medical Physics, DGMP 2020, online conference.*

- **Poster**: **Juliana Martins**, Rangoli Saxena, Sebastian Neppl, Albulaziz Alhazmi, Michael Reiner, Claus Belka, Katia Parodi. Perturbation techniques for optimizing IAEA phase spaces for different medical linacs. *ESTRO 38, 2019, Milan, Italy.*

- **Oral presentation**: **Juliana Martins**, Rangoli Saxena, Abdulaziz Alhazmi, Sebastian Neppl, Michael Reiner, Claus Belka, Stella Veloza, Katia Parodi. Optimization of Phase Space Files from clinical Linear Accelerators. *2nd European Crongress of Medical Physics, 2018, Copenhagen, Denmark.*

- **e-Poster**: Abdulaziz Alhazmi, Chiara Gianoli, **Juliana Martins**, Sebastian Neppl, Stella Veloza, Michael Reiner, Claus Belka, Katia Parodi. Proposal of a comprehensive pre-treatment QA procedure in IMRT/VMAT techniques. *2nd European Crongress of Medical Physics, 2018, Copenhagen, Denmark.*

- **Poster**: Albulaziz Alhazmi, Chiara Gianoli, Sebastian Neppl, **Juliana Martins**, Stella Veloza, Mark Podesta, Frank Verhaegen, Michael Reiner, Claus Belka, Katia Parodi. EPID-based 3D dose reconstruction for IMRT/VMAT pre-treatment QA. *Annual conference of the German Society for Medical Physics, DGMP 2018, Nuremberg, Germany.*

- **Oral presentation**: Rangoli Saxena, **Juliana Martins**, Albulaziz Alhazmi, Sebastian Neppl, Michael Reiner, Claus Belka, Stella Veloza, Katia Parodi. Optimization of IAEA Phase Space Files for Monte Carlo Dose Calculation. *Annual conference of the German Society for Medical Physics, DGMP 2018, Nuremberg, Germany.*

- **Poster**: **Juliana Martins**, Albulaziz Alhazmi, Sebastian Neppl, Claus Belka, Michael Reiner, Katia Parodi, Stella Veloza. An equipment-specific Geant4 model for the Elekta Agility collimator. *International Conference on Monte Carlo Techniques for Medical Applications, MCMA 2017, Naples, Italy.*

- **Poster**: Albulaziz Alhazmi, Chiara Gianoli, Sebastian Neppl, Stella Veloza, **Juliana Martins**, Mark Podesta, Frank Verhaegen, Michael Reiner, Claus Belka, Katia Parodi. Sensitivity analysis of EPID-based 3D dose reconstruction for VMAT QA. *ESTRO 37, 2017, Barcelona, Spain.*

Co-supervised Master Thesis

- **Rangoli Saxena**: Optimization of LINAC beam modeling and clinical translation of a novel EPID-based QA Tool. *Ludwig-Maximilians-Universität München, 2018*

Bibliography

[1] Wilhelm Conrad Röntgen. On a new kind of rays. *Science*, 3(59):227–231, 1896.

[2] Indra J Das and Harald Paganetti. Introduction and history of proton therapy. In *Principles and Practice of Proton Beam Therapy*, chapter 1, pages 1–20. Medical Physics Publishing Inc, 2015. ISBN 9781936366439.

[3] Iowa State University. History of radiography. URL `https://www.nde-ed.org/NDETechniques/Radiography/Introduction/history.xhtml`. Accessed: 04.02.2023.

[4] Richard F Mould. Emil herman grubbé (1875–1960) with special reference to priority for x-ray cancer therapy. *Nowotwory. Journal of Oncology*, 68(5-6):286–289, 2018.

[5] American Cancer Society. History of cancer treatments: Radiation therapy. URL `https://www.cancer.org/treatment/understanding-your-diagnosis/history-of-cancer/cancer-treatment-radiation.html`. Accessed: 04.02.2023.

[6] Silvia Liprandi. *Development and performance evaluation of detectors in a Compton camera arrangement for ion beam range monitoring in particle therapy.* dissertation, Ludwig-Maximilians-Universität München, 2018.

[7] Particle Therapy Co-Operative Group (PTCOG). Particle therapy centers. URL `https://www.ptcog.ch/index.php`. Accessed: 04.02.2023.

[8] International Atomic Energy Agency (IAEA). Directory of radiotherapy centres (dirac). URL `https://dirac.iaea.org/`. Accessed: 04.02.2023.

[9] Michael M Goitein and Martin Jermann. The relative costs of proton and x-ray radiation therapy. *Clinical Oncology*, 15(1):S37–S50, 2003.

[10] Damien Charles Weber, Johannes A Langendijk, Cai Grau, and Juliette Thariat. Proton therapy and the european particle therapy network: The past, present and future. *Cancer/Radiothérapie*, 24(6-7):687–690, 2020.

[11] Wilbert Cruz, Ganesh Narayanasamy, Morgan Regan, Panayiotis Mavroidis, Niko Papanikolaou, Chul S Ha, and Sotirios Stathakis. Patient specific imrt quality assurance with film, ionization chamber and detector arrays: Our institutional experience. *Radiation Physics and Chemistry*, 115:12–16, 2015.

[12] Caroline Lafond, Sophie Chiavassa, Cindy Bertaut, Nicolas Boussion, Nathalie Chapel, Lucie Chapron, Frédéric Coste, Sylvain Crespin, Gilles Dy, Papa Abdoulaye Faye, et al. Demat: a multi-institutional dosimetry audit of rotational and static intensity-modulated radiotherapy. *Physica Medica*, 32(5):664–670, 2016.

[13] Gary A Ezzell, Jay W Burmeister, Nesrin Dogan, Thomas J LoSasso, James G Mechalakos, Dimitris Mihailidis, Andrea Molineu, Jatinder R Palta, Chester R Ramsey, Bill J Salter, et al. Imrt commissioning: multiple institution planning and dosimetry comparisons, a report from aapm task group 119. *Medical physics*, 36(11):5359–5373, 2009.

[14] Alan C Hartford, James M Galvin, David C Beyer, Thomas J Eichler, Geoffrey S Ibbott, Brian Kavanagh, Christopher J Schultz, and Seth A Rosenthal. American college of radiology (acr) and american society for radiation oncology (astro) practice guideline for intensity-modulated radiation therapy (imrt). *American journal of clinical oncology*, 35(6):612–617, 2012.

[15] Casey Bojechko, Mark Phillps, Alan Kalet, and Eric C Ford. A quantification of the effectiveness of epid dosimetry and software-based plan verification systems in detecting incidents in radiotherapy. *Medical physics*, 42(9):5363–5369, 2015.

[16] Chieh-Ju Tung, Hui-Chun Wang, Shih-Han Lo, Jin-Ming Wu, and Chang Jiang Wang. In vivo dosimetry for external photon treatments of head and neck cancers by diodes and tlds. *Radiation protection dosimetry*, 111(1):45–50, 2004.

[17] Esther J Bloemen-van Gurp, Ben J Mijnheer, Tom AM Verschueren, and Philippe Lambin. Total body irradiation, toward optimal individual delivery: dose evaluation with metal oxide field effect transistors, thermoluminescence detectors, and a treatment planning system. *International Journal of Radiation Oncology* Biology* Physics*, 69(4):1297–1304, 2007.

[18] Mahin Tariq, Cindy Gomez, and Adam C Riegel. Dosimetric impact of placement errors in optically stimulated luminescent in vivo dosimetry in radiotherapy. *Physics and Imaging in Radiation Oncology*, 11:63–68, 2019.

[19] Dominique Huyskens, Ria Bogaerts, Jan Verstraete, Marika Lööf, Håkan Nyström, Claudio Fiorino, Sara Broggi, Núria Jornet, Montserrat Ribas, and DI Twaithes. Practical guidelines for the implementation of in vivo dosimetry with diodes in external radiotherapy with photon beams (entrance dose). 2001.

[20] Cynthia F Chuang, Lynn J Verhey, and Ping Xia. Investigation of the use of mosfet for clinical imrt dosimetric verification. *Medical physics*, 29(6):1109–1115, 2002.

[21] Nils Kadesjö, Tufve Nyholm, and Jörgen Olofsson. A practical approach to diode based in vivo dosimetry for intensity modulated radiotherapy. *Radiotherapy and Oncology*, 98(3):378–381, 2011.

[22] Tina Marie Briere, Michael T Gillin, and A Sam Beddar. Implantable mosfet detectors: evaluation of a new design. *Medical physics*, 34(12):4585–4590, 2007.

[23] Landon Wootton, Rajat Kudchadker, Andrew Lee, and Sam Beddar. Real-time in vivo rectal wall dosimetry using plastic scintillation detectors for patients with prostate cancer. *Physics in Medicine & Biology*, 59(3):647, 2014.

[24] Cielle Collins, Jacob Kodra, Suk Whan Yoon, Robert Coakley, John Adamovics, and Mark Oldham. Preliminary investigation of a reusable radiochromic sheet for radiation dosimetry. In *Journal of Physics: Conference Series*, volume 1305, page 012032. IOP Publishing, 2019.

[25] Yi-Fang Wang, Kevin Liu, John Adamovics, and Cheng-Shie Wuu. An investigation of dosimetric accuracy of a novel presage radiochromic sheet and its clinical applications. In *Journal of Physics: Conference Series*, volume 1305, page 012041. IOP Publishing, 2019.

[26] Daniel Létourneau, Misbah Gulam, Di Yan, Mark Oldham, and John W Wong. Evaluation of a 2d diode array for imrt quality assurance. *Radiotherapy and oncology*, 70(2):199–206, 2004.

[27] Bjoern Poppe, Arne Blechschmidt, Armand Djouguela, Ralf Kollhoff, Antje Rubach, Kay C Willborn, and Dietrich Harder. Two-dimensional ionization chamber arrays for imrt plan verification. *Medical Physics*, 33(4):1005–1015, 2006.

[28] Wouter Van Elmpt, Leah McDermott, Sebastiaan Nijsten, Markus Wendling, Philippe Lambin, and Ben Mijnheer. A literature review of electronic portal imaging for radiotherapy dosimetry. *Radiotherapy and oncology*, 88(3):289–309, 2008.

[29] Arthur L Boyer, Larry Antonuk, Aaron Fenster, Marcel Van Herk, Harm Meertens, Peter Munro, Larry E Reinstein, and John Wong. A review of electronic portal imaging devices (epids). *Medical physics*, 19(1):1–16, 1992.

[30] Larry E Antonuk. Electronic portal imaging devices: a review and historical perspective of contemporary technologies and research. *Physics in Medicine & Biology*, 47(6):R31, 2002.

[31] Mike C Kirby and Andrew G Glendinning. Developments in electronic portal imaging systems. *The British journal of radiology*, 79(special_issue_1):S50–S65, 2006.

[32] Abdulaziz Alhazmi, Chiara Gianoli, Sebastian Neppl, Juliana Martins, Stella Veloza, Mark Podesta, Frank Verhaegen, Michael Reiner, Claus Belka, and Katia Parodi. A novel approach to epid-based 3d volumetric dosimetry for imrt and vmat qa. *Physics in Medicine and Biology*, 63(11), 2018.

[33] Ben Mijnheer. Epid-based dosimetry and its relation to other 2d and 3d dose measurement techniques in radiation therapy. In *Journal of Physics: Conference Series*, volume 847, page 012024. IOP Publishing, 2017.

[34] Bo Liu, Justus Adamson, Anna Rodrigues, Fugen Zhou, Fang-fang Yin, and Qiuwen Wu. A novel technique for vmat qa with epid in cine mode on a varian truebeam linac. *Physics in Medicine & Biology*, 58(19):6683, 2013.

[35] Mark Podesta, Sebastiaan M J J G Nijsten, Lucas C G G Persoon, Stefan G Scheib, Christof Baltes, and Frank Verhaegen. Time dependent pre-treatment epid dosimetry for standard and fff vmat. *Physics in Medicine & Biology*, 59(16):4749, 2014.

[36] Peter McCowan, G Asuni, Timothy Van Beek, Eric Van Uytven, K Kujanpaa, and Boyd MC McCurdy. A model-based 3d patient-specific pre-treatment qa method for vmat using the epid. *Physics in Medicine & Biology*, 62(4):1600, 2017.

[37] Peter B Greer. 3d epid based dosimetry for pre-treatment verification of vmat–methods and challenges. In *Journal of Physics: Conference Series*, volume 444, page 012010. IOP Publishing, 2013.

[38] Igor Olaciregui-Ruiz, Sam Beddar, Peter Greer, Nuria Jornet, Boyd McCurdy, Gabriel Paiva-Fonseca, Ben Mijnheer, and Frank Verhaegen. In vivo dosimetry in external beam photon radiotherapy: Requirements and future directions for research, development, and clinical practice. *Physics and Imaging in Radiation Oncology*, 15: 108–116, 2020.

[39] Belen Juste, Rafael Miró, Sergio Díez, Juan M Campayo, and Gumersindo Verdú. Dosimetric capabilities of the iview gt portal imager using mcnp5 monte carlo simulations. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3743–3746. IEEE, 2009.

[40] Krista Chytyk-Praznik, Eric Van Uytven, Timothy A Van Beek, Peter B Greer, and Boyd MC McCurdy. Model-based prediction of portal dose images during patient treatment. *Medical physics*, 40(3):031713, 2013.

[41] Sean L Berry, Cynthia Polvorosa, Simon Cheng, Israel Deutsch, KS Clifford Chao, and Cheng-Shie Wuu. Initial clinical experience performing patient treatment verification with an electronic portal imaging device transit dosimeter. *International Journal of Radiation Oncology Biology Physics*, 88(1):204–209, 2014.

[42] Jihyung Yoon, Jae Won Jung, Jong Oh Kim, and Inhwan Yeo. A monte carlo calculation model of electronic portal imaging device for transit dosimetry through heterogeneous media. *Medical physics*, 43(5):2242–2250, 2016.

[43] Sebastiaan MJJG Nijsten, Wouter JC Van Elmpt, M Jacobs, Ben J Mijnheer, André LAJ Dekker, Philippe Lambin, and André WH Minken. A global calibration model for epids used for transit dosimetry. *Medical physics*, 34(10):3872–3884, 2007.

[44] Tae Seong Baek, Eun Ji Chung, Jaeman Son, and Myonggeun Yoon. Feasibility study on the verification of actual beam delivery in a treatment room using epid transit dosimetry. *Radiation Oncology*, 9(1):1–8, 2014.

[45] Markus Wendling, Leah N McDermott, Anton Mans, Jan-Jakob Sonke, Marcel van Herk, and Ben J Mijnheer. A simple backprojection algorithm for 3d in vivo epid dosimetry of imrt treatments. *Medical physics*, 36(7):3310–3321, 2009.

[46] Raul Pecharromán-Gallego, Anton Mans, Jan-Jakob Sonke, Joep C Stroom, Ígor Olaciregui-Ruiz, Marcel van Herk, and Ben J Mijnheer. Simplifying epid dosimetry for imrt treatment verification. *Medical physics*, 38(2):983–992, 2011.

[47] Markus Wendling, Robert JW Louwe, Leah N McDermott, Jan-Jakob Sonke, Marcel van Herk, and Ben J Mijnheer. Accurate two-dimensional imrt verification using a back-projection epid dosimetry method. *Medical physics*, 33(2):259–273, 2006.

[48] Markus Wendling, Leah N. McDermott, Anton Mans, Ígor Olaciregui-Ruiz, Raul Pecharromán-Gallego, Jan-Jakob Sonke, Joep Stroom, Marcel van Herk, and Ben J. Mijnheer. In aqua vivo epid dosimetry. *Medical physics*, 39(1):367–377, 2012.

[49] Bo Sterckx, Isabell Steinseifer, and Markus Wendling. In vivo dosimetry with an electronic portal imaging device for prostate cancer radiotherapy with an endorectal balloon. *Physics and Imaging in Radiation Oncology*, 12:7–9, 2019.

[50] Igor Olaciregui-Ruiz, Julia-Maria Osinga-Blaettermann, Karen Ortega-Marin, Ben Mijnheer, and Anton Mans. Extending in aqua portal dosimetry with dose inhomogeneity conversion maps for accurate patient dose reconstruction in external beam radiotherapy. *Physics and Imaging in Radiation Oncology*, 22:20–27, 2022.

[51] Wendel Dean Renner, Kevin Norton, and Timothy Holmes. A method for deconvolution of integrated electronic portal images to obtain incident fluence for dose reconstruction. *Journal of applied clinical medical physics*, 6(4):22–39, 2005.

[52] Geneviève Jarry and Frank Verhaegen. Patient-specific dosimetry of conventional and intensity modulated radiation therapy using a novel full monte carlo phase space reconstruction method from electronic portal images. *Physics in Medicine & Biology*, 52(8):2277, 2007.

[53] Wouter van Elmpt, Sebastiaan Nijsten, Ben Mijnheer, André Dekker, and Philippe Lambin. The next step in patient-specific qa: 3d dose verification of conformal and intensity-modulated rt based on epid dosimetry and monte carlo dose calculations. *Radiotherapy and oncology*, 86(1):86–92, 2008.

[54] Eric Van Uytven, Timothy Van Beek, Peter M McCowan, Krista Chytyk-Praznik, Peter B Greer, and Boyd MC McCurdy. Validation of a method for in vivo 3d dose reconstruction for imrt and vmat treatments using on-treatment epid images and a model-based forward-calculation algorithm. *Medical physics*, 42(12):6945–6954, 2015.

[55] Wouter van Elmpt, Sebastiaan Nijsten, Steven Petit, Ben Mijnheer, Philippe Lambin, and André Dekker. 3d in vivo dosimetry using megavoltage cone-beam ct and epid dosimetry. *International Journal of Radiation Oncology* Biology* Physics*, 73 (5):1580–1587, 2009.

[56] Stephen Steciw, Brad Warkentin, Satyapal Rathee, and Biagio G Fallone. Three-dimensional imrt verification with a flat-panel epid. *Medical physics*, 32(2):600–612, 2005.

[57] Ben Mijnheer, Igor Olaciregui-Ruiz, Roel Rozendaal, JJ Sonke, H Spreeuw, Ron Tielenburg, M Van Herk, Ron Vijlbrief, and Anton Mans. 3d epid-based in vivo dosimetry for imrt and vmat. In *Journal of Physics: Conference Series*, volume 444, page 012011. IOP publishing, 2013.

[58] Boyd McCurdy, Peter Greer, and James Bedford. Electronic portal imaging device dosimetry. In *Clinical 3D dosimetry in modern radiation therapy*, pages 169–198. CRC Press, Boca Raton, FL, 2017.

[59] Anton Mans, Markus Wendling, Leah N McDermott, Jan Jakob Sonke, Rene Tielenburg, Ron Vijlbrief, Ben Mijnheer, Marcel Van Herk, and Joep C Stroom. Catching errors with in vivo epid dosimetry. *Medical physics*, 37(6Part2):2638–2644, 2010.

[60] Philippe Meyer, Vincent Noblet, Christophe Mazzara, and Alex Lallement. Survey on deep learning for radiotherapy. *Computers in biology and medicine*, 98:126–146, 2018.

[61] Berkman Sahiner, Aria Pezeshk, Lubomir M Hadjiiski, Xiaosong Wang, Karen Drukker, Kenny H Cha, Ronald M Summers, and Maryellen L Giger. Deep learning in medical imaging and radiation therapy. *Medical physics*, 46(1):e1–e36, 2019.

[62] Huan-Hsin Tseng, Yi Luo, Randall K Ten Haken, and Issam El Naqa. The role of machine learning in knowledge-based response-adapted radiotherapy. *Frontiers in oncology*, 8:266, 2018.

[63] Xiaoke Zhang, Zongsheng Hu, Guoliang Zhang, Yongdong Zhuang, Yuenan Wang, and Hao Peng. Dose calculation in proton therapy using a discovery cross-domain generative adversarial network (discogan). *Medical Physics*, 48(5):2646–2660, 2021.

[64] Timo M Deist, Frank JWM Dankers, Gilmer Valdes, Robin Wijsman, I-Chow Hsu, Cary Oberije, Tim Lustberg, Johan van Soest, Frank Hoebers, Arthur Jochems, et al. Machine learning algorithms for outcome prediction in (chemo) radiotherapy: An empirical comparison of classifiers. *Medical physics*, 45(7):3449–3459, 2018.

[65] Sunan Cui, Yi Luo, Huan-Hsin Tseng, Randall K Ten Haken, and Issam El Naqa. Combining handcrafted features with latent variables in machine learning for prediction of radiation-induced lung damage. *Medical physics*, 46(5):2497–2511, 2019.

[66] Joscha Maier, Stefan Sawall, Michael Knaup, and Marc Kachelrieß. Deep scatter estimation (dse): accurate real-time scatter estimation for x-ray ct using a deep convolutional neural network. *Journal of Nondestructive Evaluation*, 37(3):1–9, 2018.

[67] Joscha Maier, Elias Eulig, Tim Vöth, Michael Knaup, Jan Kuntz, Stefan Sawall, and Marc Kachelrieß. Real-time scatter estimation for medical ct using the deep scatter estimation: Method and robustness analysis with respect to different anatomies, dose levels, tube voltages, and data truncation. *Medical Physics*, 46(1):238–249, 2019.

[68] Adrian B Levine, Colin Schlosser, Jasleen Grewal, Robin Coope, Steve JM Jones, and Stephen Yip. Rise of the machines: advances in deep learning for cancer diagnosis. *Trends in cancer*, 5(3):157–169, 2019.

[69] Julien Erath, Tim Vöth, Joscha Maier, Eric Fournié, Martin Petersilka, Karl Stierstorfer, and Marc Kachelrieß. Deep learning-based forward and cross-scatter correction in dual-source ct. *Medical Physics*, 2021.

[70] Yongbao Li, Fan Xiao, Biaoshui Liu, Mengke Qi, Xingyu Lu, Jiajun Cai, Linghong Zhou, and Ting Song. Deep learning-based 3d in vivo dose reconstruction with an electronic portal imaging device for magnetic resonance-linear accelerators: a proof of concept study. *Physics in Medicine & Biology*, 66(23):235011, 2021.

[71] Joscha Maier, Laura Klein, Elias Eulig, Stefan Sawall, and Marc Kachelrieß. Real-time estimation of patient-specific dose distributions for medical ct using the deep dose estimation. *Medical Physics*, 49(4):2259–2269, 2022.

[72] Issam El Naqa, Jim Irrer, Tim A Ritter, John DeMarco, Hania Al-Hallaq, Jeremy Booth, Grace Kim, Ahmad Alkhatib, Richard Popple, Mario Perez, et al. Machine learning for automated quality assurance in radiotherapy: A proof of principle using epid data description. *Medical physics*, 46(4):1914–1921, 2019.

[73] Madoka Sakai, Hisashi Nakano, Daisuke Kawahara, Satoshi Tanabe, Takeshi Takizawa, Akihiro Narita, Takumi Yamada, Hironori Sakai, Masataka Ueda, Ryuta Sasamoto, et al. Detecting mlc modeling errors using radiomics-based machine learning in patient-specific qa with an epid for intensity-modulated radiation therapy. *Medical Physics*, 48(3):991–1002, 2021.

[74] Seied Rabie Mahdavi, Asieh Tavakol, Mastaneh Sanei, Seyed Hadi Molana, Farshid Arbabi, Aram Rostami, and Sohrab Barimani. Use of artificial neural network for pretreatment verification of intensity modulation radiation therapy fields. *The British journal of radiology*, 92(1102):20190355, 2019.

[75] Dao Lam, Xizhe Zhang, Harold Li, Yang Deshan, Brayden Schott, Tianyu Zhao, Weixiong Zhang, Sasa Mutic, and Baozhou Sun. Predicting gamma passing rates for portal dosimetry-based imrt qa using machine learning. *Medical physics*, 46(10):4666–4675, 2019.

[76] Troy P Teo, Syed Bilal Ahmed, Philip Kawalec, Nadia Alayoubi, Neil Bruce, Ethan Lyn, and Stephen Pistorius. Feasibility of predicting tumor motion using online data acquired during treatment and a generalized neural network optimized with offline patient tumor trajectories. *Medical physics*, 45(2):830–845, 2018.

[77] Cecile JA Wolfs, Richard AM Canters, and Frank Verhaegen. Identification of treatment error types for lung cancer patients using convolutional neural networks and epid dosimetry. *Radiotherapy and Oncology*, 153:243–249, 2020.

[78] Peng Dong and Lei Xing. Deep dosenet: a deep neural network for accurate dosimetric transformation between different spatial resolutions and/or different dose calculation algorithms for precision radiation therapy. *Physics in Medicine & Biology*, 65(3):035010, 2020.

[79] Igor Olaciregui-Ruiz, Iban Torres-Xirau, Jonas Teuwen, Uulke A van der Heide, and Anton Mans. A deep learning-based correction to epid dosimetry for attenuation and scatter in the unity mr-linac system. *Physica Medica*, 71:124–131, 2020.

[80] Joscha Maier, Elias Eulig, Sabrina Dorn, Stefan Sawall, and Marc Kachelrieß. Real-time patient-specific ct dose estimation using a deep convolutional neural network. In *2018 IEEE Nuclear Science Symposium and Medical Imaging Conference Proceedings (NSS/MIC)*, pages 1–3, 2018.

[81] Shrikant Deshpande, Samuel J Blake, Aitang Xing, Peter E Metcalfe, Lois C Holloway, and P Vial. A simple model for transit dosimetry based on a water equivalent epid. *Medical physics*, 45(3):1266–1275, 2018.

[82] Penelope Allisy-Roberts and Jerry Williams. Chapter 1 - radiation physics. In Penelope Allisy-Roberts and Jerry Williams, editors, *Farr's Physics for Medical*

Imaging (Second Edition), pages 1–21. W.B. Saunders, second edition edition, 2008. ISBN 978-0-7020-2844-1.

[83] Faiz M Khan and John P Gibbons. *Khan's the physics of radiation therapy.* Lippincott Williams & Wilkins, 2014.

[84] Glenn F Knoll. *Radiation detection and measurement.* John Wiley & Sons, 2010.

[85] Ervin B Podgorsak et al. *IAEA Radiation Oncology Physics: A Handbook for Teachers And Students.* IAEA Vienna, 2005.

[86] Frank Herbert Attix. *Introduction to radiological physics and radiation dosimetry.* John Wiley & Sons, 2008.

[87] Ervin B Podgoršak et al. *Radiation physics for medical physicists*, volume 1. Springer, 2006.

[88] Hanno Krieger. *Strahlungsquellen für Technik und Medizin.* B.G. Teubner Verlag, 2005.

[89] David I Thwaites and John B Tuohy. Back to the future: the history and development of the clinical linear accelerator. *Physics in Medicine & Biology*, 51(13):R343, 2006.

[90] Wolfgang Schlegel, Christian P. Karger, and Oliver Jäkel. *Medizinische Physik: Grundlagen - Bildgebung - Therapie - Technik.* Springer Spektrum, 2018.

[91] Peter J. Rosemark. *A Primer on Theory and Operation of Linear Accelerators in Radiation Therapy by C. J. Karzmark and R. J. Morton*, volume 12. Medical Physics, 1985.

[92] American Association of Physicists in Medicine (AAPM and others). Aapm report 72-basic applications of multileaf collimators. *Med Phys*, 2001.

[93] Christopher M Thompson, Steven J Weston, Vivian C Cosgrove, and David I Thwaites. A dosimetric characterization of a novel linear accelerator collimator. *Medical Physics*, 41(3), 2014.

[94] Simeon Nill, Ralf Hinderer, and Uwe Oelfke. X-imrt. In Wolfgang Schlegel, Thomas Bortfeld, and Anca-Ligia Grosu, editors, *New Technologies in Radiation Oncology*, pages 289–298. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. ISBN 978-3-540-29999-8.

[95] Ross I Berbeco. *Beam's Eye View Imaging in Radiation Oncology.* CRC Press, 2017.

[96] Mark Podesta, Lucas CGG Persoon, and Frank Verhaegen. A novel time dependent gamma evaluation function for dynamic 2d and 3d dose distributions. *Physics in Medicine & Biology*, 59(20):5973, 2014.

[97] Peter B Greer. Correction of pixel sensitivity variation and off-axis response for amorphous silicon epid dosimetry. *Medical physics*, 32(12):3558–3568, 2005.

[98] Peter Winkler, Alfred Hefner, and Dietmar Georg. Dose-response characteristics of an amorphous silicon epid. *Medical physics*, 32(10):3095–3105, 2005.

[99] Charles Kirkby and Ron Sloboda. Consequences of the spectral response of an a-si epid and implications for dosimetric calibration. *Medical physics*, 32(8):2649–2658, 2005.

[100] Boyd MC McCurdy, Kurt Luchka, and Stephen Pistorius. Dosimetric investigation and portal dose image prediction using an amorphous silicon electronic portal imaging device. *Medical physics*, 28(6):911–924, 2001.

[101] Samuel J. Blake, Philip Vial, Lois Holloway, Peter B. Greer, Aimee L. McNamara, and Zdenka Kuncic. Characterization of optical transport effects on epid dosimetry using geant4. *Medical Physics*, 40(4):1–14, 2013.

[102] C Kirkby and R Sloboda. Comprehensive monte carlo calculation of the point spread function for a commercial-si epid. *Medical physics*, 32(4):1115–1127, 2005.

[103] R. Mohan and T. Bortfeld. The potential and limitations of imrt: A physicist's point of view. In Thomas Bortfeld, Rupert Schmidt-Ullrich, Wilfried De Neve, and David E. Wazer, editors, *Image-Guided IMRT*, pages 11–18. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.

[104] Thomas Bortfeld, Arthur L Boyer, Wolfgang Schlegel, Darren L Kahler, and Timothy J Waldron. Realization and verification of three-dimensional conformal radiotherapy with modulated fields. *International Journal of Radiation Oncology* Biology* Physics*, 30(4):899–908, 1994.

[105] Karl Otto. Volumetric modulated arc therapy: Imrt in a single gantry arc. *Medical physics*, 35(1):310–317, 2008.

[106] Tao Zhang, Zhi-Wen Liang, Jun Han, Jian-Ping Bi, Zhi-Yong Yang, and Hong Ma. Double-arc volumetric modulated therapy improves dose distribution compared to static gantry imrt and 3d conformal radiotherapy for adjuvant therapy of gastric cancer. *Radiation oncology*, 10(1):1–8, 2015.

[107] Joao Seco and Frank Verhaegen. *Monte Carlo techniques in radiation therapy.* CRC press Boca Raton, FL:, 2013.

[108] John S Hendricks. A monte carlo code for particle transport. *Los Alamos Science*, 22:42–33, 1994.

[109] Chiara Gianoli, Sebastian Meyer, Lorena Magallanes, Chiara Paganelli, Guido Baroni, and Katia Parodi. Analytical simulator of proton radiography and tomography for different detector configurations. *Physica Medica*, 59:92–99, 2019.

[110] Katia Parodi, A Ferrari, F Sommerer, and Harald Paganetti. Clinical ct-based calculations of dose and positron emitter distributions in proton therapy using the fluka monte carlo code. *Physics in Medicine & Biology*, 52(12):3369, 2007.

[111] Dave W O Rogers. Fifty years of monte carlo simulations for medical physics. *Physics in Medicine and Biology*, 51(13):R287–R301, jul 2006.

[112] Frank Verhaegen and Jan Seuntjens. Monte carlo modelling of external radiotherapy photon beams. *Physics in medicine & biology*, 48(21):R107, 2003.

[113] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, Pedro Arce, Makoto Asai, D Axen, Swagato Banerjee, GJNI Barrand, et al. Geant4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[114] John Allison, Katsuya Amako, JEA Apostolakis, HAAH Araujo, P Arce Dubois, MAAM Asai, GABG Barrand, RACR Capra, SACS Chauvie, RACR Chytracek, et al. Geant4 developments and applications. *IEEE Transactions on nuclear science*, 53(1):270–278, 2006.

[115] J Allison, Katsuya Amako, John Apostolakis, Pedro Arce, M Asai, T Aso, E Bagli, A Bagulya, S Banerjee, GJNI Barrand, et al. Recent developments in geant4. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 835:186–225, 2016.

[116] Louis Archambault, Luc Beaulieu, Julien F. Carrier, Francesca Castrovillari, Stephane Chauvie, Franca Foppiano, Giovanni Ghiso, Susanna Guatelli, Sébastien Incerti, Ernesto Lamanna, Stig Larsson, M.C. Lopes, Luis Peralta, Maria G. Pia, Pedro Rodrigues, V.H. Tremblay, and Andrea Trindade. Overview of geant4 applications in medical physics. In *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515)*, volume 3, pages 1743–1745 Vol.3, 2003.

[117] Julien F Carrier, Louis Archambault, Luc Beaulieu, and R Roy. Validation of geant4, an object-oriented monte carlo toolkit, for simulations in medical physics. *Medical physics*, 31(3):484–492, 2004.

[118] Emily Poon and Frank Verhaegen. Accuracy of the photon and electron physics in geant4 for radiotherapy applications. *Medical physics*, 32(6Part1):1696–1711, 2005.

[119] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[120] Joscha Maier. *Artifact correction and real-time scatter estimation for X-ray computed tomography in industrial metrology*. dissertation, Ruperto-Carola University of Heidelberg, 2019.

[121] Murat H Sazli. A brief review of feed-forward neural networks. *Communications Faculty of Sciences University of Ankara Series A2-A3 Physical Sciences and Engineering*, 50(01), 2006.

[122] T. Han, J. Mikell, M. Salehpour, and F. Mourtada. Dosimetric comparison of acuros xb deterministic radiation transport method with monte carlo and model-based convolution methods in heterogeneous media. *Medical Physics*, 38(6):3722, 2011.

[123] Antonella Fogliata, Eugenio Vanetti, Dirk Albers, Carsten Brink, Alessandro Clivio, Tommy Knöös, Giorgia Nicolini, and Luca Cozzi. On the dosimetric behaviour of photon dose calculation algorithms in the presence of simple geometric heterogeneities: Comparison with monte carlo calculations. *Physics in Medicine and Biology*, 52(5):1363–1385, 2007.

[124] Edmond Sterpin, Milan Tomsej, Bert De Smedt, Nick Reynaert, and Stefaan Vynckier. Monte Carlo evaluation of the AAA treatment planning algorithm in a heterogeneous multilayer phantom and IMRT clinical treatments for an Elekta SL25 linear accelerator. *Medical Physics*, 34(5):1665–1677, 2007.

[125] Khaled Aljarrah, Greg C. Sharp, Toni Neicu, and Steve B. Jiang. Determination of the initial beam parameters in monte carlo linac simulation. *Medical Physics*, 33(4):850–858, 2006.

[126] Sigrun Saur Almberg, Jomar Frengen, Arve Kylling, and Tore Lindmo. Monte Carlo linear accelerator simulation of megavoltage photon beams: Independent determination of initial beam parameters. *Medical Physics*, 39(1):40–47, 2012.

[127] Miguel Antonio Cortés-Giraldo, José Manuel Quesada, María Isabel Gallardo, and Roberto Capote. An implementation to read and write IAEA phase-space files in GEANT4-based simulations. *International Journal of Radiation Biology*, 88(1-2):200–208, jan 2012.

[128] Maria F Belosi, Miguel Rodriguez, Antonella Fogliata, Luca Cozzi, Josep Sempau, Alessandro Clivio, Giorgia Nicolini, Eugenio Vanetti, Harald Krauss, Catherine Khamphan, Pascal Fenoglietto, Josep Puxeu, David Fedele, Pietro Mancosu, and Lorenzo Brualla. Monte Carlo simulation of TrueBeam flattening-filter-free beams using Varian phase-space files: Comparison with experimental data. *Medical Physics*, 41(5), 2014.

[129] Jun Deng, Steve B Jiang, Ajay Kapur, Jinsheng Li, Todd Pawlicki, and C-M Ma. Photon beam characterization and modelling for Monte Carlo treatment planning. *Physics in Medicine and Biology*, 45:411–427, 2000.

[130] Iwan Kawrakow and Blake R B Walters. Efficient photon beam dose calculations using dosxyznrc with beamnrc. *Medical Physics*, 33(8):3046–3056, 2006.

[131] Indra J Das, Chee-Wai Cheng, Ronald J Watts, Anders Ahnesjö, John Gibbons, X Allen Li, Jessica Lowenstein, Raj K Mitra, William E Simon, and Timothy C Zhu. Accelerator beam data commissioning equipment and procedures: report of the tg-106 of the therapy physics committee of the aapm. *Medical physics*, 35(9): 4186–4215, 2008.

[132] Jens Fleckenstein. *Methods to enhance the dosimetric accuracy of volumetric modulated arc therapy (VMAT) treatments in radiation therapy.* dissertation, Ruprecht-Karls-Universität Heidelberg, 2013.

[133] Daryoush Sheikh-Bagheri and Dave W.O. Rogers. Monte Carlo calculation of nine megavoltage photon beam spectra using the BEAM code. *Medical Physics*, 29(3): 391–402, 2002.

[134] Martin J Berger, Johnathan S Coursey, Max A Zucker, and J Chang. Stopping-power and range tables for electrons, protons, and helium ions, nist standard reference database 124, 2017. URL `https://physics.nist.gov/cgi-bin/Star/compos.pl?text`. Accessed: 25.07.2022.

[135] James L Bedford, Michael D R Thomas, and Gregory Smyth. Beam modeling and VMAT performance with the Agility 160-leaf multileaf collimator. *Journal of Applied Clinical Medical Physics*, 14(2):172–185, 2013.

[136] Elekta Limited. Agility$^{\text{TM}}$ and integrity$^{\text{TM}}$ r3.x, information for treatment planning systems, 05 2013. Document ID: 1504231 01.

[137] Yuji Nakaguchi, Takeshi Oono, Fujio Araki, and Masato Maruyama. [Physical characterizations for an integrated 160-leaf multi-leaf collimator with a new concept design]. *Nippon Hoshasen Gijutsu Gakkai Zasshi*, 69(7):778–783, 2013.

[138] Ryota Onizuka, Fujio Araki, and Takeshi Ohno. Monte Carlo dose verification of VMAT treatment plans using Elekta Agility 160-leaf MLC. *Physica Medica*, 51 (May):22–31, 2018.

[139] Gavin Cranmer-Sargison, Paul ZY Liu, Steven Weston, Natalka Suchowerska, and David I Thwaites. Small field dosimetric characterization of a new 160-leaf mlc. *Physics in Medicine & Biology*, 58(20):7343, 2013.

[140] Conzia Avigo. Assessment of an EPID based MLC quality assurance protocol using the commercial sofwtare Dosimetry Check. Master's thesis, Università di Pisa, 2015.

[141] Michael Snyder, Robert Halford, Cory Knill, Jeffrey N. Adams, Todd Bossenberger, Adrian Nalichowski, Ahmad Hammoud, and Jay Burmeister. Modeling the agility mlc in the monaco treatment planning system. *Journal of Applied Clinical Medical Physics*, 17(3):190–202, 2016.

[142] Jacek M Chojnowski, Lee M Taylor, Jonathan R Sykes, and David I Thwaites. Beam focal spot position determination for an elekta linac with the agility® head; practical guide with a ready-to-go procedure. *Journal of applied clinical medical physics*, 19 (4):44–47, 2018.

[143] Ulrike Kuck. Geant4 Modeling of the Medical Accelerator Head Agility and Subsequent Validation by Measurements. Master's thesis, Fakult¨at f¨ur Mathematik, Informatik und Naturwissenschaften der RWTH Aachen, 2014.

[144] Sara Gholampourkashi, Joanna E. Cygler, Jason Belec, Miro Vujicic, and Emily Heath. Monte carlo and analytic modeling of an elekta infinity linac with agility mlc: Investigating the significance of accurate model parameters for small radiation fields. *Journal of Applied Clinical Medical Physics*, 20(1):55–67, 2019.

[145] Michael Roche, Robert Crane, Marcus Powers, and Timothy Crabtree. Agility mlc transmission optimization in the monaco treatment planning system. *Journal of Applied Clinical Medical Physics*, 19(5):473–482, 2018.

[146] Arthur L Boyer and Shidong Li. Geometric analysis of light-field position of a multileaf collimator with curved ends. *Medical physics*, 24(5):757–762, 1997.

[147] Lori A Young, Fei Yang, Ning Cao, and Juergen Meyer. Rounded leaf end modeling in pinnacle vmat treatment planning for fixed jaw linacs. *Journal of Applied Clinical Medical Physics*, 17(6):149–162, 2016.

[148] Abdulaziz Alhazmi. *EPID-based 3D dosimetry for pre-treatment IMRT/VMAT Quality Assurance.* dissertation, Ludwig-Maximilians-Universität München, 2018.

[149] Susanne Schmid. Systematic Monte Carlo study on in vivo proton range verification with prompt gamma photons in proton therapy. Master's thesis, Ludwig-Maximilians-Universität München, 2015.

[150] Andreas Franz Resch. Development of a Monte Carlo simulation-based treatment planning tool for proton therapy. Master's thesis, Technische Universität München, 2015.

[151] Uwe Schneider, Eros Pedroni, and Antony Lomax. The calibration of CT hounsfield units for radiotherapy treatment planning. *Physics in Medicine and Biology*, 41(1): 111–124, jan 1996.

[152] Wilfried Schneider, Thomas Bortfeld, and Wolfgang Schlegel. Correlation between CT numbers and tissue parameters needed for monte carlo simulations of clinical dose distributions. *Physics in Medicine and Biology*, 45(2):459–478, jan 2000.

[153] Geant4 Collaboration. Book for application developers, 2021. URL `https://geant4-userdoc.web.cern.ch/UsersGuides/ForApplicationDeveloper/html/index.html#`. Accessed: 25.11.2022.

[154] Daryoush Sheikh-Bagheri and Dave WO Rogers. Sensitivity of megavoltage photon beam monte carlo simulations to electron beam and other parameters. *Medical physics*, 29(3):379–390, 2002.

[155] Omar Chibani and C-M Charlie Ma. On the discrepancies between monte carlo dose calculations and measurements for the varian photon beam. *Medical physics*, 34(4): 1206–1216, 2007.

[156] Lilie LW Wang and Konrad Leszczynski. Estimation of the focal spot size and shape for a medical linear accelerator by monte carlo simulation. *Medical physics*, 34(2): 485–488, 2007.

[157] DT Tai, ND Son, TTH Loan, and HD Tuan. A method for determination of parameters of the initial electron beam hitting the target in linac. In *Journal of Physics: Conference Series*, volume 851, page 012032. IOP Publishing, 2017.

[158] Edmond Sterpin, Yan Chen, Weiguo Lu, Thomas R Mackie, Gustavo H Olivera, and Stefaan Vynckier. On the relationships between electron spot size, focal spot size, and virtual source position in monte carlo simulations.

[159] Omar Chibani, Belal Moftah, and C-M Charlie Ma. On monte carlo modeling of megavoltage photon beams: a revisited study on the sensitivity of beam parameters. *Medical physics*, 38(1):188–201, 2011.

[160] Indrin J Chetty, Bruce Curran, Joanna E Cygler, John J DeMarco, Gary Ezzell, Bruce A Faddegon, Iwan Kawrakow, Paul J Keall, Helen Liu, C-M Charlie Ma, et al. Report of the aapm task group no. 105: Issues associated with clinical implementation of monte carlo-based photon and electron external beam treatment planning. *Medical physics*, 34(12):4818–4853, 2007.

[161] Chang Ming Charlie Ma, Indrin J Chetty, Jun Deng, Bruce Faddegon, Steve B Jiang, Jinsheng Li, Jan Seuntjens, Jeffrey V Siebers, and Erik Traneus. Beam modeling and beam model commissioning for monte carlo dose calculation-based radiation

therapy treatment planning: Report of aapm task group 157. *Medical physics*, 47 (1):e1–e18, 2020.

[162] PTW. microdiamond - synthetic diamond detectors. URL `https://www.ptwdosimetry.com/en/products/microdiamond/`. Accessed: 29.10.2022.

[163] Geant4 Collaboration. Guide for application developers using the geant4 toolk. URL `https://geant4-userdoc.web.cern.ch/UsersGuides/ForApplicationDeveloper/html/`. Accessed: 29.10.2022.

[164] Tarun K Podder, Greg Bednarz, Yan Yu, and James M Galvin. Physical characterization and comparison of two commercially available micro-mlcs. *Physica Medica*, 27(1):52–57, 2011.

[165] Mostafa Heydarian, Khadijeh Asnaashari, Mahmoud Allahverdi, and David A Jaffray. Dosimetric evaluation of a dedicated stereotactic linear accelerator using measurement and monte carlo simulation. *Medical Physics*, 35(9):3943–3954, 2008.

[166] Paolo Francescon, Stefania Cora, and Ninfa Satariano. Calculation of $k^{f_{clin},f_{msr}}_{Q_{clin},Q_{msr}}$ for several small detectors and for two linear accelerators using monte carlo simulations. *Medical Physics*, 38(12):6513–6527, nov 2011.

[167] L Garren, IG Knowles, T Sjöstrand, and T Trippe. Monte carlo particle numbering scheme. *The European Physical Journal. C, Particles and Fields.*, 15(1-4):205–207, 2000.

[168] Jarno Van de Walle, C Martens, Nick Reynaert, H Palmans, MARC Coghe, Wilfried De Neve, Carlos De Wagter, and Hubert Thierens. Monte carlo model of the elekta sliplus accelerator: validation of a new mlc component module in beam for a 6 mv beam. *Physics in Medicine & Biology*, 48(3):371, 2003.

[169] Martin B Tacke, Hanitra Szymanowski, Uwe Oelfke, Carsten Schulze, Susanne Nuss, Eugen Wehrwein, and Stefan Leidenberger. Assessment of a new multileaf collimator concept using geant4 monte carlo simulations. *Medical Physics*, 33(4):1125–1132, 2006.

[170] I. Antoniu Popescu, Christopher P. Shaw, Sergei F. Zavgorodni, and Wayne A. Beckham. Absolute dose calculations for Monte Carlo simulations of radiotherapy beams. *Physics in Medicine and Biology*, 50(14):3375–3392, 2005.

[171] Paolo Francescon, Carlo Cavedon, Sonia Reccanello, and Stefania Cora. Photon dose calculation of a three-dimensional treatment planning system compared to the monte carlo code beam. *Medical physics*, 27(7):1579–1587, 2000.

[172] Egor Borzov, Alexander Nevelsky, Raquel Bar-Deroma, and Itzhak Orion. Dosimetric characterization of elekta stereotactic cones. *Journal of Applied Clinical Medical Physics*, 19(1):194–203, 2018.

[173] Tanya Kairn, Scott B. Crowe, Christopher M. Poole, and Andrew L. Fielding. Effects of collimator backscatter in an elekta linac by monte carlo simulation. *Australasian Physical and Engineering Sciences in Medicine*, 32(3):129–135, 2009.

[174] Josep Sempau, Alberto Sánchez-Reyes, Francesc Salvat, H Oulad Ben Tahar, Steve B Jiang, and José M Fernández-Varea. Monte Carlo simulation of electron beams from an accelerator head using PENELOPE INSTITUTE OF PHYSICS PUBLISHING PHYSICS IN MEDICINE Monte Carlo simulation of electron beams from an accelerator head using PENELOPE. *AND BIOLOGY Phys. Med. Biol*, 46:1163–1186, 2001.

[175] Blake RB Walters, Ivan Kawrakow, and Dave WO Rogers. History by history statistical estimators in the BEAM code system. *Medical Physics*, 29(12):2745–2752, 2002.

[176] Sebastian Neppl. Picket fence and beyond - what cna 2d array measurments and the gamma index method tell about leaf misalignment? Master's thesis, Ludwig-Maximilians-Universität München, 2014.

[177] Maximilian Alber, Matthias Birkner, W. Laub, and Fridtjof Nüsslin. Hyperion — an integrated imrt planning tool. In Wolfgang Schlegel and Thomas Bortfeld, editors, *The Use of Computers in Radiation Therapy*, pages 46–48, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg. ISBN 978-3-642-59758-9.

[178] Jan Hofmaier, Steffi Kantz, Matthias Söhn, Oliver S Dohm, Stefan Bächle, Markus Alber, Katia Parodi, Claus Belka, and Maximilian Niyazi. Hippocampal sparing radiotherapy for glioblastoma patients: a planning study using volumetric modulated arc therapy. *Radiation Oncology*, 11(1):1–8, 2016.

[179] Daniel A Low, William B Harms, Sasa Mutic, and James A Purdy. A technique for the quantitative evaluation of dose distributions. *Medical Physics*, 25(5):656–661, 1998.

[180] Daniel A Low and James F Dempsey. Evaluation of the gamma dose distribution comparison method. *Medical physics*, 30(9):2455–2464, 2003.

[181] Jun Deng, Todd Pawlicki, Yan Chen, Jinsheng Li, Steve B Jiang, and CM Ma. The mlc tongue-and-groove effect on imrt dose distributions. *Physics in Medicine & Biology*, 46(4):1039, 2001.

[182] J Martins, S Neppl, A Alhazmi, J Fleckenstein, G Dedes, M Reiner, C Belka, K Parodi, and S Veloza. Analytical modeling of the parameters of a multileaf collimator for monte carlo dose simulations in radiotherapy. In *2020 Joint AAPM COMP Virtual Meeting - AAPM*, 2020.

[183] I Antoniu Popescu, Christopher P Shaw, Sergei F Zavgorodni, and Wayne A Beckham. Absolute dose calculations for monte carlo simulations of radiotherapy beams. *Physics in Medicine & Biology*, 50(14):3375, 2005.

[184] ICRU. Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (imrt). icru report 83. *Journal of the International Commission on Radiation Units and Measurements*, 10(1):106, 2010.

[185] Gary A Ezzell, Jay W Burmeister, Nesrin Dogan, Thomas J LoSasso, James G Mechalakos, Dimitris Mihailidis, Andrea Molineu, Jatinder R Palta, Chester R Ramsey, Bill J Salter, et al. Imrt commissioning: multiple institution planning and dosimetry comparisons, a report from aapm task group 119. *Medical physics*, 36(11):5359–5373, 2009.

[186] Roberto Capote, Robert Jeraj, Chang Ming Ma, Dave W O Rogers, Francisco Sánchez-Doblado, Josep Sempau, Jan Seuntjens, and Jeffrey V Siebers. Phase-space database for external beam radiotherapy. Summary report of a consultants' meeting. Technical report, International Atomic Energy Agency, 2006.

[187] Ivan Kawrakow. Accurate condensed history monte carlo simulation of electron transport. i. egsnrc, the new egs4 version. *Medical Physics*, 27(3):485–498, mar 2000.

[188] Francesc Salvat, José M Fernández-Varea, and Josep Sempau. PENELOPE-2008: A Code System for Monte Carlo Simulation of Electron and Photon Transport. In *the Workshop Proceedings, June*, 2008.

[189] Juliana Cristina Martins, Abdulaziz Alhazmi, Sebastian Neppl, Claus Belka, Michael Reiner, Katia Parodi, and Stella Veloza. Abstract ID: 144 An equipment-specific Geant4 model for the Elekta Agility collimator. *Physica Medica*, 42:30–31, oct 2017.

[190] E Tonkopi, M. R. McEwen, B R B Walters, and I Kawrakow. Influence of ion chamber response on in-air profile measurements in megavoltage photon beams. *Medical Physics*, 32(9):2918–2927, aug 2005.

[191] Loïc Grevillot, Thibault Frisson, D. Maneval, N. Zahra, J. N. Badel, and D. Sarrut. Simulation of a 6 MV Elekta Precise Linac photon beam using GATE/GEANT4. *Physics in Medicine and Biology*, 56(4), 2011.

[192] Tanya Kairn, David Cassidy, P. M. Sandford, and Andrew L. Fielding. Radiotherapy treatment verification using radiological thickness measured with an amorphous silicon electronic portal imaging device: Monte Carlo simulation and experiment. *Physics in Medicine and Biology*, 53(14):3903–3919, 2008.

[193] Rangoli Saxena. Optimization of linac beam modeling and clinical translation of a novel epid-based qa tool. Master's thesis, Ludwig-Maximilians-Universität München, 2018.

[194] Juliana Cristina Martins, Rangoli Saxena, Sebastian Neppl, Abdulaziz Alhazmi, Michael Reiner, Stella Veloza, Claus Belka, and Katia Parodi. Optimization of Phase Space files from clinical linear accelerators. *Physica Medica*, 64(45):54–68, aug 2019.

[195] Jens Fleckenstein, Lennart Jahnke, Frank Lohr, Frederik Wenz, and Jürgen Hesser. Development of a Geant4 based Monte Carlo Algorithm to evaluate the MONACO VMAT treatment accuracy. *Zeitschrift für Medizinische Physik*, 23(1):33–45, feb 2013.

[196] Edmond Sterpin, Yan Chen, Weiguo Lu, Thomas Rock Mackie, Gustavo H Olivera, and Stefaan Vynckier. On the relationships between electron spot size, focal spot size, and virtual source position in Monte Carlo simulations. *Medical Physics*, 38 3: 1579–1586, 2011.

[197] Juliana Cristina Martins, Joscha Maier, Chiara Gianoli, Sebastian Neppl, George Dedes, Abdulaziz Alhazmi, Stella Veloza, Michael Reiner, Claus Belka, Marc Kachelrieß, et al. Towards real-time epid-based 3d in vivo dosimetry for imrt with deep neural networks: A feasibility study. *Physica Medica*, 114:103148, 2023.

[198] Juliana Cristina Martins, Abdulaziz Alhazmi, Sebastian Neppl, Claus Belka, Michael Reiner, Katia Parodi, and Stella Veloza. Abstract id: 144 an equipment-specific geant4 model for the elekta agility collimator. *Physica Medica: European Journal of Medical Physics*, 42:30–31, 2017.

[199] Nikos Papanikolaou, Jerry J Battista, Arthur L Boyer, Constantin Kappas, Eric Klein, T Rock Mackie, Michael Sharpe, and Jake Van Dyk. Report of task group no. 65 of the radiation therapy committee of the american association of physicists in medicine: Tissue inhomogeneity corrections for megavoltage photon beams. *AAPM Task Group*, 65:1–142, 2004.

[200] Yulia Levakhina. *Three-Dimensional Digital Tomosynthesis: Iterative reconstruction, artifact reduction and alternative acquisition geometry*. Springer, 2014.

[201] Sebastian Meyer, Marco Pinto, Katia Parodi, and Chiara Gianoli. The impact of path estimates in iterative ion ct reconstructions for clinical-like cases. *Physics in Medicine & Biology*, 66(9):095007, 2021.

[202] Robert L Siddon. Fast calculation of the exact radiological path for a three-dimensional ct array. *Medical physics*, 12(2):252–255, 1985.

[203] Timothy C Zhu, Anders Ahnesjö, Kwok Leung Lam, X Allen Li, Chang-Ming Charlie Ma, Jatinder R Palta, Michael B Sharpe, Bruce Thomadsen, and Ramesh C Tailor. Report of aapm therapy physics committee task group 74: in-air output ratio, for megavoltage photon beams. *Medical physics*, 36(11):5261–5291, 2009.

[204] John H Hubbell and Stephen M Seltzer. Tables of x-ray mass attenuation coefficients and mass energy-absorption coefficients (version 1.4), 2004. URL `http://physics.nist.gov/xaamdi`. Accessed: 20.12.2020.

[205] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4.

[206] Daniel A Low, William B Harms, Sasa Mutic, and James A Purdy. A technique for the quantitative evaluation of dose distributions. *Medical physics*, 25(5):656–661, 1998.

[207] International Commission on Radiation Units and Measurements. *ICRU Report 83 Prescribing, Recording, and Reporting Photon-beam Intensity-modulated Radiation Therapy (IMRT)-Journal of the ICRU-Vol 10 No 1 2010.* Oxford University Press, 2010.

[208] Phillip Duncan-Gelder, Alicia Moggré, Andrew Cousins, Benjamin Wilder, and Steven Marsh. Accurate dosimetric measurement of large extended ssd fields for comparison to tps models. *Physica Medica*, 84:220–227, 2021.

[209] Baozhou Sun, Dharanipathy Rangaraj, Sunita Boddu, Murty Goddu, Deshan Yang, Geethpriya Palaniswaamy, Sridhar Yaddanapudi, Omar Wooten, and Sasa Mutic. Evaluation of the efficiency and effectiveness of independent dose calculation followed by machine log file analysis against conventional measurement based imrt qa. *Journal of applied clinical medical physics*, 13(5):140–154, 2012.

[210] Vicki Trier Taasti, Peter Klages, Katia Parodi, and Ludvig Paul Muren. Developments in deep learning based corrections of cone beam computed tomography to enable dose calculations for adaptive radiotherapy. *Physics and Imaging in Radiation Oncology*, 15:77–79, 2020.

[211] Christopher Kurz, George Dedes, Andreas Resch, Michael Reiner, Ute Ganswindt, Reinoud Nijhuis, Christian Thieke, Claus Belka, Katia Parodi, and Guillaume Landry. Comparing cone-beam ct intensity correction methods for dose recalculation in adaptive intensity-modulated photon and proton therapy for head and neck cancer. *Acta oncologica*, 54(9):1651–1657, 2015.

[212] David C Hansen, Guillaume Landry, Florian Kamp, Minglun Li, Claus Belka, Katia Parodi, and Christopher Kurz. Scatternet: a convolutional neural network for cone-beam ct intensity correction. *Medical physics*, 45(11):4916–4926, 2018.

[213] Guillaume Landry, David Hansen, Florian Kamp, Minglun Li, Ben Hoyle, Jochen Weller, Katia Parodi, Claus Belka, and Christopher Kurz. Comparing unet training with three different datasets to correct cbct images for prostate radiotherapy dose calculations. *Physics in Medicine & Biology*, 64(3):035011, 2019.

[214] Matteo Maspero, Antonetta C Houweling, Mark HF Savenije, Tristan CF van Heijst, Joost JC Verhoeff, Alexis NTJ Kotte, and Cornelis AT van den Berg. A single neural network for cone-beam computed tomography-based radiotherapy of head-and-neck, lung and breast cancer. *Physics and imaging in radiation oncology*, 14:24–31, 2020.

[215] Brad Warkentin, Stephen Steciw, Satyapal Rathee, and Biagio G Fallone. Dosimetric imrt verification with a flat-panel epid. *Medical physics*, 30(12):3143–3155, 2003.

[216] W Nijsten. van elmpt, m. jacobs, b. mijnheer, a. *Dekker, P. Lambin, and A. Minken," A global calibration model for a-Si EPIDs used for transit dosimetry," Med Phys*, 34:3872–3884, 2007.

[217] Emily Heath and Joao Seco. Dynamic beam delivery and 4d monte carlo. In *Monte Carlo Techniques in Radiation Therapy*, pages 111–126. CRC Press, 2016.

[218] Lennart Jahnke, Jens Fleckenstein, Frederik Wenz, and Jürgen Hesser. Gmc: A gpu implementation of a monte carlo dose calculation based on geant4. *Physics in Medicine and Biology*, 57(5):1217–1229, 2012.

[219] Otto Seiskari, Jukka Kommeri, and Tapio Niemi. Gpu in physics computation: Case geant4 navigation. pages 1–11, 2012. URL `http://arxiv.org/abs/1209.5235`.

[220] Reid W. Townson, Xun Jia, Zhen Tian, Yan Jiang Graves, Sergei Zavgorodni, and Steve B. Jiang. Gpu-based monte carlo radiotherapy dose calculation using phase-space sources. *Physics in Medicine and Biology*, 58(12):4341–4356, 2013.

[221] Julien Bert, Hector Perez-Ponce, Ziad El Bitar, Sébastien Jan, Yannick Boursier, Damien Vintache, Alain Bonissent, Christian Morel, David Brasse, and Dimitris Visvikis. Geant4-based monte carlo simulations on gpu for medical applications. *Physics in Medicine & Biology*, 58(16):5593, 2013.

[222] Robert E Drzymala, Radhe Mohan, Linda Brewster, James Chu, Michael Goitein, William Harms, and Marcia Urie. Dose-volume histograms. *International Journal of Radiation Oncology* Biology* Physics*, 21(1):71–78, 1991.

[223] Chee-Wai Cheng and Indra J Das. Treatment plan evaluation using dose–volume histogram (dvh) and spatial dose–volume histogram (zdvh). *International Journal of Radiation Oncology* Biology* Physics*, 43(5):1143–1150, 1999.

Acknowledgements

I have imagined this moment many times in the past few years, and what I could possibly say in this acknowledgments section. I don't think there are words that can describe the joy to finish a doctorate. More importantly, there might not be enough words to thank all the people that have somehow been part of it, but I will try my best to show my appreciation to them all. This might not be the shortest text, but these people deserve it.

I start thanking God (or whatever one might name it) for the strength. Having faith in a better tomorrow gives me purpose and helps me overcoming difficult moments.

I thank the Brazilian National Council for Scientific and Technological Development (Conselho Nacional de Desenvolvimento Científico e Tecnológico - CNPq, grant number 201383/2015-2) for the financial support between February 2016 and January 2020.

I would like to thank Prof. Dr. Jan Wilkens for the support, and for taking the time to review this thesis as a second referee. I am extremely grateful to Prof. Dr. Marc Kachelrieß and Dr. Joscha Maier for the fruitful discussions, their expertise and invaluable support with the deep learning part of this project.

I want to thank my beloved Henrique. We came together to Munich and we had to face so many challenges, the PhD itself being one of the hardest ones. So many weekends, nights and holidays spent working on my thesis, and you have been by my side on all of them, trying to keep me calm, bringing me water, preparing me a tea, cooking, cleaning... Countless gestures of love and support. Thanks to you, this day has become a reality. We did it!

For my parents, Rosane and Valdeir, who have supported my decision to move to a different continent, I owe the greatest acknowledgment. I also need to include Marcia here, which might not be a parent, but certainly nearly as important as one. Thank you for celebrating with me the achievements and the opportunities which came with this PhD project, even if it meant not being with each other as often as we wished. Thank you for accepting my absence over the years, for being loving and understanding despite the huge, sometimes seemly unbearable, distance. Here I also extend my appreciation to Denise and Roberto, which had to live far away from their beloved son Henrique. Thank you all for everything, not only during the PhD, but for providing us with so much over our entire lives. Our brothers Rafael and Fernando, and the lovely girls who joined the family, Daniely and Beatriz, I thank you for being with our parents on our absence. Thank you

also for being another source of support and positivity, for the visits, and for sharing the live with us!

To my loved grandparents, Dalila, Iolanda and Aluisio, I want to thank you for all the love, and apologize for my absence. Being away from you has been the hardest part. Our time together was and is always precious. I will treasure every moment that we have been together and keep them in my memories forever.

A great thanks to all my friends in Brazil (and around the world) for the support and words of encouragement. With you, the friendship is so natural that whenever we meet (virtually or personally), it feels like we've never been apart. Some names need to be mentioned, even if I take the risk of forgetting someone: Simone, Nicolle, Amanda, Aluisio, Taisi, Mari, Ju, Josi, Louise, Lucas, Juniar, Cadu, Bruna, Bárbara, Marta, Lana, Paulo and Mike. Thank you for being so great.

The most valuable gifts I have gotten from my PhD are the people I have met. I will try my best to express what you mean for me. Starting from the beginning, how lucky was I to meet Basti on my very first day! Thank you for the help on apartment hunting, moving, bureaucracies, for the Feuerzangenbowle nights, and for so much more, but especially thank you for your friendship. Silvia and Rita, thank you for being like sisters to me, you enrich my life, and being with you over the past few years filled me with optimism and strength to keep fighting. Dear Franz, we have been there for each other all the way, I cherish the honestly that our friendship is based on. Thank you for being kind, but also strict when necessary. Dear Matthias, thank you for being so kind, thank you for the jokes and the laughter, for the coffees and Krapfen, and for your friendship. My friend Thomas, thank you for your warm latino heart, it brings me closer to home every time (and I apologize for the Argentinian jokes!). Dear Katrin, it was an immense pleasure to share the office with you! Thank you for the motivating words, for the tears, for the laughter, countless coffees, cat snuggles, lunches, dinners, anyway, thank you for giving so much! Dear Chiara, thank you for your friendship, for your kindness, availability and for always being in a good mood! Thank you not only for your emotional support, but also for helping me with essential parts of this thesis. Dear Linda, Jennifer, Davide and Lukas, you also have been a source of support for all of us, PhD students. Thank you for the patience and the partnership. And for the little Mattia and Valentin, your parents are amazing, and I hope I can be a great aunt Ju for you.

Dear Dr. Stella Veloza, it was a pleasure working with you! Thank you for all the support at the start and at the end of my PhD, inside and outside of work. Dear Dr. George Dedes, thank you for the great help with Geant4 and cluster, your support was essential to the conclusion of this work. Dear Prof. Peter Thirolf, thank you for the kindness and for sharing all your knowledge. To my master student Rangoli, I have learned a lot with you, thank you for the great work. Basti, Ingrid, Max (both!), Flo, Jannis, Giorgia, Aziz, and so many others, thank you for the good times spent together, it has definitely made the time of my PhD lighter.

For the experimental part of this thesis, I am deeply grateful to the physicists at Klinikum Grosshadern, especially to Sebastian Neppl, for finding the time (at evenings!) to support me with the measurements I needed, always with a smile on his face and a good mood. I would also like to thanks Dr. Michael Reiner for the enlightening discussions, for support with so many questions, and for the great espressos shared in the clinic. I am also thankful to Prof. Dr. Claus Belka for providing access to the data and the linac installed at the Klinikum. I would like to extend my sincere thanks to Prof. Dr. Guillaume Laundry and Dr. Mark Podesta for the support with the gamma evaluation routine.

To my Brazilian friends in Munich: Andrea, thank you for opening your home for me during my first days in Munich, and for caring for me like a mother would. Natasha and Derk, what a gift from the universe to be neighbors with you! Thank you for the wines, the good food, the trips to Spina, simply for the good times we always have together. Camylla, Eliel, Sumaya and Renan, you bring a piece of Paraná to Munich, and Henrique and I are grateful for this.

I also would like to thank my friends and colleagues at Klinikum rechts der Isar. Thank you for receiving me so well and for understanding the stress I had during the final months of my thesis. Many thanks to Martin for his house in Mittenwald, truly the place where I was most productive in writing. I am especially thankful to Frauke for offering me a place on her office, where I could concentrate and write my thesis and manuscript, for the encouragement and support (both from her and from Liesl!), and simply for being a good company.

I want to give my deepest appreciation to all the teachers I had throughout my entire life. You have instigated in me the curiosity and the desire to learn. Special thanks to my dear Prof. Carias, who helped me so much, and sadly passed away during my PhD.

I would also like to acknowledge all the dedication and patience of our dear secretary Andrea Leinthaler, who helped so many students during her years at the chair. Thank you dear Andrea, you are deeply missed.

Finally, but the most importantly, I need to thank the person who has made it all possible. Dear Prof. Dr. Katia Parodi, I cannot describe how much I appreciate and admire you. Not only for being probably the most talented and dedicated person I have ever met, but especially for finding a way to be kind and for believing in people. You believed in me when I first applied and gave me a space in your group. You kept believing in me even when I felt hopeless. You were kind and understanding during the discouraging moments. You were a cheerleader and celebrated the small victories. You were a great supervisor when I needed guidance with the project. You are simply an incredible person. If I have here so much to say to so many wonderful people, it is because you gave me a chance. Il mio prossimo progetto è imparare l'italiano, quindi voglio concludere questa tesi con un piccolo accenno al futuro: Grazie mille cara Katia, con tutto il mio cuore.