Università degli Studi di Pavia

Dipartimento di Fisica "A. Volta"

Corso di Laurea Magistrale in Scienze Fisiche

Development and validation of a matRad-based independent dose calculation platform for experimental research in hadron therapy

Supervisor:
Ballarini Francesca

MSc thesis by:
Sara Petini

External supervisors:
Savazzi Simone
Pullia Marco

A.A. 2024-2025

A Mamma e Papà,
che sempre al mio fianco,
mi hanno accompagnata
in questo percorso.

Contents

Abstract

Hadron therapy is an alternative form of radiation therapy that uses hadrons, in particular protons, carbon ions and helium, to treat tumors.
In this context, the CNAO (*Centro Nazionale di Adroterapia Oncologica*) performs this type of oncological therapy using both types of particles. In addition to patient treatments, an intense research activity is carried out in the center. In many experiments, it is necessary to use plans generated by clinical treatment planning systems; however, this represents a strong limitation, since these plans require a significant amount of time to be produced and a continuous dialog between the research group and the medical physicists, resulting in a lack of autonomy for researchers.
This situation leads to the necessity of finding an alternative way to generate plans independently from professional systems. The idea is to make use of the open-source software matRad, which is developed to simulate a clinical treatment planning system in order to support study and research activities. Using this toolkit, it is possible to write personalized code to generate plans that can be used to perform experimental irradiations. Different trials are carried out, and different codes are developed in order to identify the best configuration to generate a uniform and flat dose distribution.
Once the best code is identified, it is necessary to validate the results by making a comparison with the plans generated by the clinical treatment planning system. To achieve this, the data obtained from the matRad output are properly exported and used to generate a plan with the professional software.
Finally, the dose profiles are compared and, taking into account some differences in the calculation modalities, it is possible to conclude that the outputs of the two programs are acceptable and represent a good path to be considered.
In conclusion, matRad can be considered an effective tool for independently producing treatment plans that can be used in experimental research in hadron therapy.

Introduction

This work focuses on improving a dose calculation code intended for research applications, exploiting the open-source toolkit *matRad*, which is used to simulate a treatment planning system.

In a hadron therapy center such as CNAO (*Centro Nazionale di Adroterapia Oncologica*), research activities in the fields of radiobiology, microdosimetry, and dosimeter characterization are carried out alongside clinical oncological treatments. These research areas play a fundamental role in deepening the understanding of the mechanisms that govern the response of cells, tissues, and the tumor microenvironment following irradiation with charged particles, and are therefore essential to improve treatment accuracy and effectiveness.

Experimental activities performed at the cellular level address several crucial topics, including tumor radioresistance, healthy tissue response, high-LET radiation effects, and the interactions among different therapeutic approaches.

A rapidly expanding area of research concerns the personalization of treatments, achieved through the identification of biological biomarkers capable of predicting individual radiosensitivity or describing the evolution of the cellular response over time. In addition, radiobiology investigates the late effects of radiation, such as the onset of secondary tumors, with particular attention to young patients.

Within this framework, an accurate characterization of the particle beams used in experimental studies is of paramount importance. A detailed knowledge of the three-dimensional dose distribution and its biological quality is fundamental not only for the validation of experimental results, but also to ensure a reliable translation of laboratory findings into clinical applications. To this end, research groups often rely on treatment planning systems (TPS) already employed in clinical practice; however, these software tools are certified for clinical use and are not always readily accessible for research purposes.

This limitation constitutes the starting point of the present work. Although the developed code cannot achieve the same level of precision as clinical TPS, it represents an independent tool that provides research groups with greater autonomy in simple plan generation. This approach accelerates the experimental workflow, allows for personalized configurations, and ensures consistency between physical planning and biological analyses, thereby contributing concretely to the integration of physics, biology, and clinical practice in hadron therapy research.

Chapter 1

Hadron therapy

Hadron therapy is a particular form of radiation therapy. The term "hadron" refers to "heavy particle"; they are made up of quarks, and a large number of particles fall under this classification.
Today, around 20 million new cancer cases are diagnosed worldwide each year [1]. At diagnosis, between 5% and 20% of patients present with metastatic disease, depending on tumor type and region [2]. Over the course of their illness, about 50% of all cancer patients will receive radiotherapy at least once [3, 4]. Among those treated with palliative intent for metastases, most still undergo conventional three-dimensional conformal radiotherapy [5], whereas advanced techniques such as IMRT (Intensity Modulated Radiation Therapy), VMAT (Volumetric Modulated Arc Therapy), SBRT (Stereostatic Body Radiation Therapy), or SRS (Stereostatic RadioSurgery) are increasingly applied in selected cases [6]. The first proposal for hadron therapy dates back to 1946, when Robert R. Wilson suggested its potential. The first treatments were carried out only a few years later, after the cyclotron was developed by Lawrence, Tobias, and Anger.[7] Since then, continuous research and technological advances have allowed this form of therapy to become increasingly available in clinics, offering new treatment opportunities and improving survival for cancer patients.[8]

1.1 Generalities

The term radiation therapy refers to all medical treatments making use of ionizing radiation to treat tumors. Its most common application is in oncology, specifically for the treatment and control of localized tumors, while it cannot be applied to metastatic disease. It is often combined with surgery, chemotherapy, or other thera-

pies to achieve the best possible local tumor control.
Radiation therapy is based on the physical principle of ionization. When ionizing radiation traverses biological matter, it deposits energy along its path through a series of interactions with atomic electrons and nuclei. The energy released during these interactions can break molecular bonds, leading to damage to DNA and other critical subcellular structures. Such damage may occur either directly, through ionization of the DNA molecule itself, or indirectly, via the production of reactive chemical species following the radiolysis of water. If the resulting cellular damage is severe or cannot be adequately repaired, it can ultimately result in loss of cellular function and cell death.
The ratio between the mean energy imparted by ionizing radiation and the mass of the irradiated portion of matter is defined as the absorbed dose, measured in Gray (Gy), where 1 Gy corresponds to 1 Joule per kilogram (1 Gy = 1 J/kg). In general, the higher the dose, the greater the radiation-induced damage.
There are two main strategies to deliver radiation dose to the tumor. One is *brachytherapy*, where radioactive sources are surgically implanted inside or near the tumor, producing a rapid dose falloff around the target volume. The second is *external beam radiotherapy*, where high-energy photon or particle beams are delivered from outside the patient, shaping the dose distribution to cover the tumor while minimizing exposure of surrounding healthy tissues.
When radiation is used to treat tumors, it is important to consider the heterogeneous nature of the irradiated region, where tissues with different radiosensitivities are present. In particular, it is essential to distinguish between the dose delivered to the tumor and the one received by healthy tissues, since it can lead to severe consequences such as the occurrence of new tumors.
The main goal of treatment planning is therefore to find a compromise between tumor control and normal tissue complications, delivering the highest possible dose to the tumor while sparing healthy tissues as much as possible.

1.2 Photon radiotherapy

Conventional radiotherapy employs highly energetic and penetrating external photon beams, calibrated to reach the tumor site. The imparted dose during the treatment is about 60-70 Gy divided in fractions of few Gy each.
To understand the effects of radiation, it is essential to recall the physical principles that describe photon interactions with matter. There are three main mechanisms

responsible for these interactions, depending on the photon energy and on the atomic number (Z) of the target material.
The *photoelectric effect* dominates at low photon energies and in high-Z materials. The *Compton effect* is predominant at intermediate energies and in low-Z materials. Finally, *pair production* becomes relevant for photon energies above 1.022 MeV and is more pronounced in high-Z materials.
The three mechanisms lead to the production of electrons, which, being charged particles, are responsible for the ionization and excitation of atoms and molecules in the traversed matter. The biological damage caused by electromagnetic radiation mainly arises from the indirect interaction of these secondary electrons with critical DNA structures, through the creation of free radicals. Direct interactions are also possible, but this type of damage generally results in single-strand breaks, which are easy to be repaired by the cell.

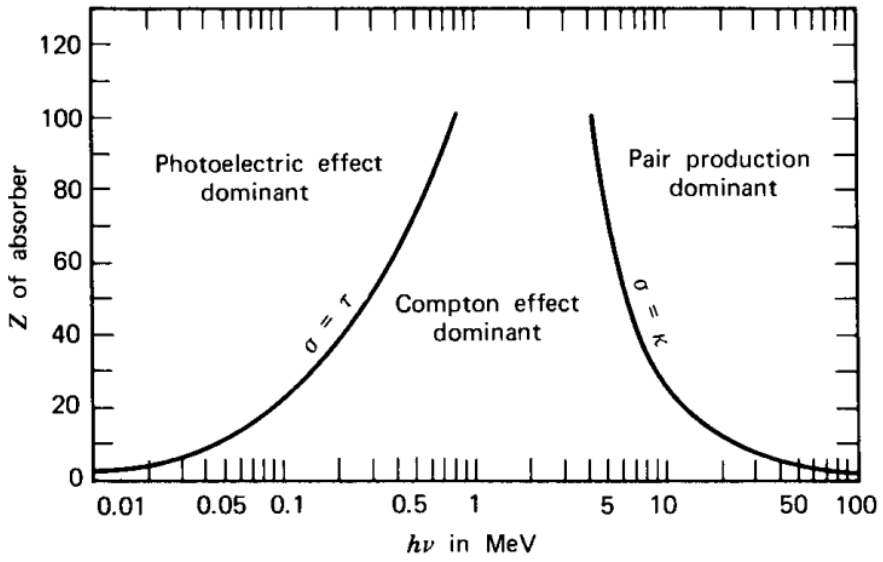

Figure 1.1: Cross sections of the photoelectric effect, Compton scattering, and pair production as a function of photon energy and the atomic number Z of the target material.

The cross sections of these processes, as a function of photon energy and the atomic number of the target material, are illustrated in Fig. **??**. To evaluate the overall effect of photon interactions, the individual contributions must be summed to obtain the total cross section, which represents the probability of any photon–matter interaction occurring:

$$\sigma_{\mathrm{tot}} = \sigma_{\mathrm{ph}} + \sigma_{\mathrm{c}} + \sigma_{\mathrm{pp}} \tag{1.1}$$

When a photon from the beam is absorbed through one of these processes, the beam intensity decreases. Assuming the initial intensity of the beam is known, the attenuation of this quantity as the beam traverses the target follows an exponential law:

$$I(x) = I_0\, e^{-\mu x} \tag{1.2}$$

where I_0 is the initial intensity, x is the thickness of the target material, and μ is the linear attenuation coefficient.
The coefficient μ depends on the photon energy and on the material properties, and it is directly proportional to the total interaction cross section.

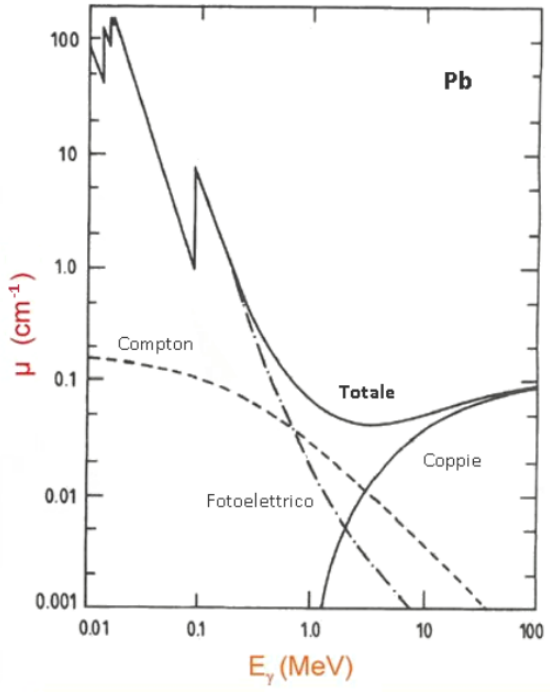

Figure 1.2: Linear attenuation coefficient μ as a function of photon energy for a lead target. The contributions of the photoelectric effect, Compton scattering, and pair production, as well as their sum leading to the total attenuation coefficient, are shown.

During radiotherapy with photons, the dose delivered to tissues decreases with depth as photons traverse the target. The majority of the dose is imparted in the initial segment of the path, where healthy tissues are located, and then attenuates with depth, where the tumor is typically situated.
The depth-dose curve for photons, shown in Fig. 1.3, highlights a fundamental limitation: the maximum dose deposition almost never coincides with the tumor depth. Therefore, the physical characteristics of photon beams alone cannot be

exploited optimally.

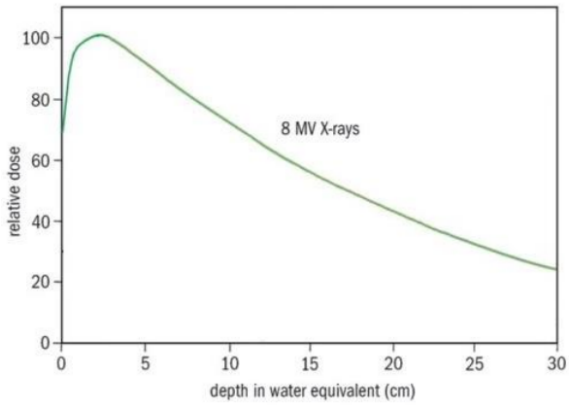

Figure 1.3: Depth-dose curve for photons.

To overcome this limitation and deliver a sufficiently high dose to the tumor while keeping the dose to healthy tissues within acceptable limits, patients are irradiated from multiple directions, each defined as a *treatment field*. These fields are planned with different entry angles and orientations so that the deposited doses overlap in the target volume, while the surrounding healthy tissues receive only a fraction of the total dose.

Another challenge in photon therapy arises from the physics of photon–matter interactions, which cause both energetic and spatial dispersion of the secondary electron beam. This effect, particularly evident in the plane orthogonal to the beam direction (*lateral straggling*), is mainly due to Compton scattering of primary photons and elastic scattering of secondary electrons. As a consequence, precise dose shaping and conformity to the target volume become non-trivial when treating deep-seated tumors, and in some critical cases, even impossible.

1.3 Hadrontherapy

Hadrontherapy, as already mentioned, is a highly precise form of radiotherapy that makes use of hadrons as projectiles. Hadrons are particles composed of quarks and antiquarks and are subject to the strong nuclear force. Among them, the particles most widely used in hadron therapy are protons and light ions such as He^{2+} and C^{6+}.

Unlike photons, which are neutral and interact weakly with matter, charged particles deposit their energy in a very different way. Energy loss mainly occurs through two mechanisms: *collisions* and *radiation*.

In collisions, the particle transfers part of its energy directly to the atoms and molecules of the medium, producing ionization or excitation. In radiation processes, instead, the particle is decelerated and slightly deflected from its path by the electromagnetic fields of the atomic nuclei it encounters. This deceleration causes the emission of X-rays with a continuous spectrum, a phenomenon known as *bremsstrahlung*, which lead to an additional energy loss.
The contribution of radiation losses depends on the inverse of the square of the projectile mass and, for ions, it is negligible even at extremely high energies. For this reason, this term can be disregarded and we can focus on the collisional contribution. The behavior of the energy lost through collisions is described by the Bethe-Bloch formula, reported below:

$$-\frac{dE}{dx} = k^2 N Z \frac{4\pi e^4 z^2}{m_e v^2} \left[\ln \frac{2 m_e v^2}{I} - \ln(1-\beta^2) - \beta^2 \right] \tag{1.3}$$

where:

N = atomic density of the target $[cm^{-2}]$

Z = atomic number of the target

z = charge number of the projectile

v = velocity of the projectile $[cm\, s^{-1}]$

I = mean ionization potential of the target [eV]

e = electron charge $[1.6 \cdot 10^{-19}\, C]$

m_e = electron rest mass $[9.11 \cdot 10^{-31}\, kg]$

k = Coulomb constant $[8.99 \cdot 10^{9}\, N\, m^2\, C^{-2}]$

$\beta = v/c$, where c is the speed of light in vacuum $[3 \cdot 10^{8}\, m\, s^{-1}]$

The quantity $-\frac{dE}{dx}$ is called the *stopping power* and represents the mean energy lost per unit path length by a charged particle due to collisions.
An example of the behavior of the stopping power as a function of the projectile energy and of the penetration depth in the target is shown in Fig. 1.4. The figure on the right shows the energy deposition as a function of the traversed thickness. It is worth emphasizing that the energy deposition increases sharply in the final part of the particle's path, that is, at low projectile energies. This particular behavior is known as the Bragg peak.

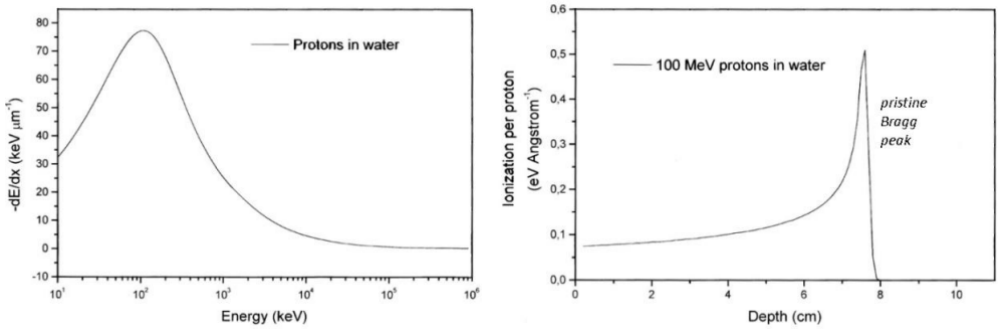

Figure 1.4: Left: stopping power of protons in water. Right: distribution of the ionization generated in water by 100 MeV protons.

Similarly to the stopping power, another quantity can be defined: the *linear energy transfer* (LET), commonly used in radiobiology to estimate the dose delivered to tissues. This quantity represents the linear transfer of energy from charged particles to matter.

If only collisions that involve a transferred energy lower than a certain threshold Δ are considered, the so-called *restricted LET* can be introduced:

$$L_\Delta = \left(\frac{dE}{dx}\right)_\Delta \tag{1.4}$$

If, on the other hand, all collisions depositing energy along the particle track are considered, without applying any threshold, the *total LET* (L) is used. Its value corresponds to the stopping power, $-\frac{dE}{dx}$. The higher the stopping power, the higher the LET of the particle, and consequently, the greater the dose deposited in the tissue.

In particular, it is possible to distinguish between high LET and low LET radiations. High LET radiations are characterized by a very dense energy deposition along their tracks, whereas low LET radiations deposit energy more sparsely. This difference in energy distribution has important implications for the type and complexity of DNA damage induced.

In Fig. 1.5, the comparison is shown between the dose deposition in tissue from photons, protons, and carbon ions.

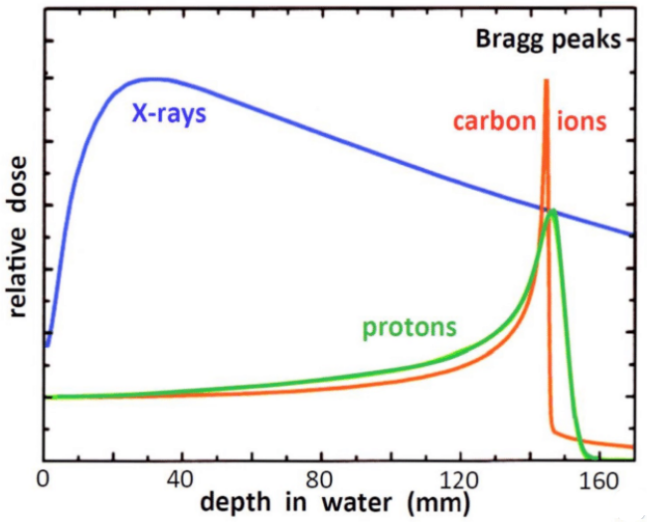

Figure 1.5: Comparison of deposited dose curves as a function of depth in tissue for photons, protons, and carbon ions.

It is immediately clear that the dose distribution produced by protons and carbon ions is much more advantageous compared to that of photons. In fact, they release only a small amount of energy in healthy tissues at the beginning of the particle path (proximal region), while the maximum energy is deposited within the tumor volume (distal region), where the Bragg peak occurs, followed by a rapid fall-off. This allows sparing of healthy tissues located beyond the tumor.
The Bragg peak is therefore particularly suitable for maximizing damage to tumor cells while minimizing the dose delivered to surrounding healthy tissues.
However, in the case of carbon ions, the fragmentation of the projectile at the end of the path must also be taken into account. As shown in Fig. 1.5, this effect produces a small tail in the distal region, which contributes to the dose deposited in healthy tissues beyond the tumor. This contribution is not negligible, especially when the tumor is located near organs at risk.
Another advantage of hadrons is their very small straggling, both in energy and in position, which allows them to maintain an almost linear trajectory up to their complete slowing down, thus ensuring ballistic precision during treatment.
The great potential of hadron therapy makes it possible to achieve an optimal dose distribution even in depth, delivering a high dose to the target while depositing only a minimal dose in the proximal tissues and virtually no dose in the distal healthy tissues.

1.3.1 Spread Out Bragg peak

The Bragg peak has a very narrow dimension, typically of the order of millimeters. A single Bragg peak is therefore not sufficient to cover the entire tumor volume, which makes it necessary to use the *Spread Out Bragg Peak* (SOBP). The SOBP is obtained by superimposing several Bragg peaks at different energies and different intensities, resulting in a dose profile like that shown in Fig. 1.6. In this way, it is possible to deliver a uniform dose to the target volume while irradiating with high precision tumor tissues.

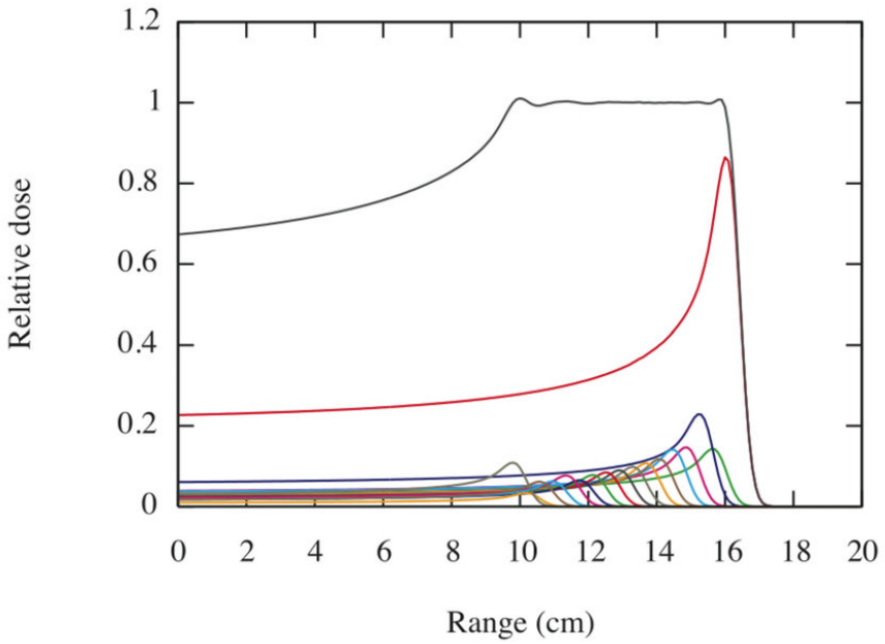

Figure 1.6: Spread Out Bragg Peak (SOBP): comparison of deposited dose profiles as a function of depth in tissue.

The main drawback of this technique is that the superposition of dose curves also increases the dose delivered in the proximal region, thus affecting healthy tissues. Nevertheless, the SOBP still allows for a significant *dose sparing* in the proximal region compared to photon irradiation. As shown in Fig. 1.7, in order to deliver 100% of the planned dose to the target volume, the photon dose required in the proximal region is much higher than the corresponding proton dose.
The SOBP does not present drawbacks in the distal region, where the dose profile maintains the characteristic sharp fall-off at the end of the target volume.

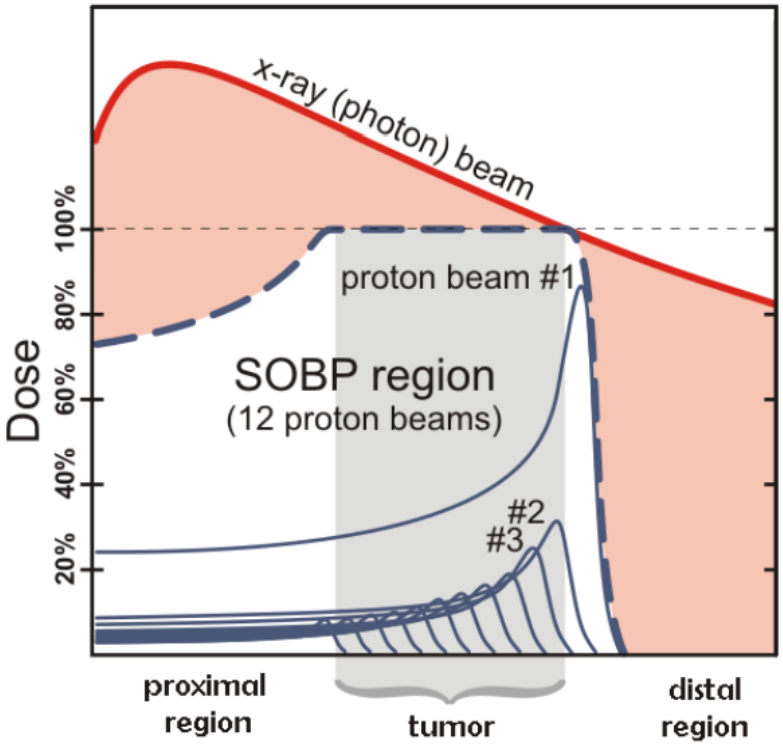

Figure 1.7: Comparison of the depth dose distribution between SOBP and photons.

1.4 Biological radiation effects

So far, the discussion has focused on the physics of ionizing radiation and how they deposit energy in matter. However, when considering the biological effects, it is essential to take into account not only the absorbed dose but also the *quality* of the radiation. In fact, different types of particles that deliver the same dose can induce completely different biological outcomes.

Beyond the absorbed dose curves, it is important to investigate how different particles interact with the cells of the irradiated tissues from a radiobiological perspective and how these interactions affect the progression of their life cycle.

The biological effects of radiation are mainly dependent on ionization and excitation processes in the molecules of the tissue. In a cell, these processes trigger a complex cascade of events that can alter molecular composition and lead to modifications of biological functions, such as mutations, transformations, or cell death [9].

The sensitivity of cells to radiation is not uniform. The most vulnerable cellular component is the nucleus, which contains DNA, the molecule that carries all the information necessary for cell reproduction and for maintaining essential biological functions. In addition to DNA, RNA is another crucial molecule, responsible for

transcribing the genetic information encoded in DNA.
DNA has a double-helix structure with a diameter of a few nanometers (typically 2-3 nm). It is a polymer whose functional units are the nucleotides, each consisting of a phosphate group, a sugar (DNA deoxyribose, RNA deoxyribose), and one of four nitrogenous bases: adenine (A), guanine (G), cytosine (C) and thymine (T). In RNA, thymine is replaced by uracil (U).
DNA is the key element in cell reproduction. A single modification in the nucleotide sequence, if it occurs in regions involved in transcription , can lead to severe alterations in the progeny.
Cells do have several DNA repair mechanisms, but these are not 100% efficient and errors may still occur. In particular, mutations affecting tumor suppressor genes or proto-oncogenes, which regulate cell growth, may lead to tumor development, it is no coincidence that cancer cells are characterized by uncontrolled proliferation. The therapeutic goal is therefore to damage the DNA of these cells in order to block their growth and induce cell death.
The biological effects induced by conventional radiotherapy and by hadron therapy differ significantly, both in the type of interaction with the cell and in their biological consequences.

1.4.1 Direct and Indirect Interaction

When ionizing radiation interacts with biological matter, it can damage DNA either directly or indirectly, leading to various types of lesions, which can be classified as follows:

- *Single base damage*: damage affecting a single DNA base, primarily induced by low LET radiations such as X-rays or electrons;
- *Single strand break (SSB)*: a break occurring in only one of the two strands of the DNA molecule;
- *Double strand break (DSB)*: a complete rupture of the DNA double helix, typically caused by high LET radiations.

Among these different types of damage, *single base damage* is generally the easiest for the cell to repair, whereas *double strand breaks* are the most difficult to deal with. Therefore, in the context of tumor radiotherapy, inducing DSBs is particularly desirable, as they are more likely to lead to effective cell killing.

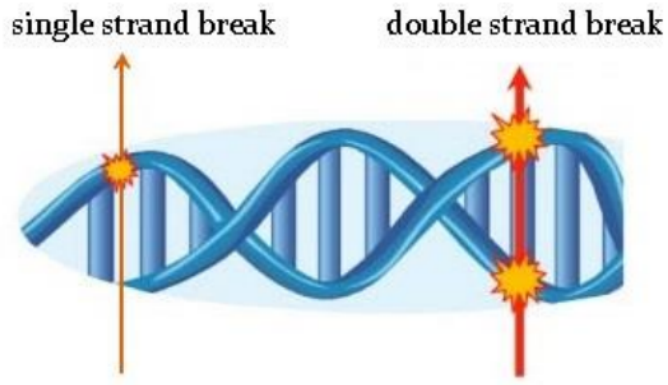

Figure 1.8: Representation of single strand break and double strand break in the DNA chain.

As previously mentioned, low-LET radiations primarily affect DNA bases through an indirect mechanism. Because cells are largely composed of water, most of the energy deposited in cellular tissue leads to the ionization and excitation of water molecules, producing highly reactive free radicals (e.g., $OH^{\cdot -}$ and $H^{\cdot +}$). These radicals subsequently interact with DNA bases, causing damage. In this case, the interaction is chemical rather than direct: the radiation does not act on the DNA itself, but the damage is mediated by reactive intermediates.
In contrast, single and double strand breaks result from a direct interaction between ionizing radiation and the DNA double helix. The ionization produced along the particle track directly induces the rupture of the DNA strand.
The extent of direct interactions depends on the LET of the incident particle. High LET particles, such as protons and carbon ions, produce dense ionization tracks that lead to more clustered DNA damage. Low LET particles, such as photons, primarily induce indirect or more dispersed damage, often resulting in SSBs.
Because of their high LET, light ions are therefore excellent candidates for effective radiotherapy.

1.4.2 RBE: Relative Biological Effectiveness

The biological effects of radiation on the cells of an organism do not depend solely on the LET of the impinging radiation, but on several other factors, including the delivered dose, the type of particle, the kind of cell irradiated, and the specific phase of the cell cycle at the time of exposure.
The efficacy of a given type of radiation in inducing cell death can be evaluated by measuring its effect on a defined cell population. The results are typically represented

through cell survival curves, where the logarithm of the surviving cell fraction is plotted as a function of the administered dose.
An example is shown in Fig. 1.9. The initial shoulder, observed mainly with low-LET radiation, reflects the ability of cells to repair DNA damage induced by irradiation. Beyond this region, the curve becomes linear, indicating an exponential decrease in the surviving fraction: the damage becomes too complex to be repaired, ultimately leading to cell death.

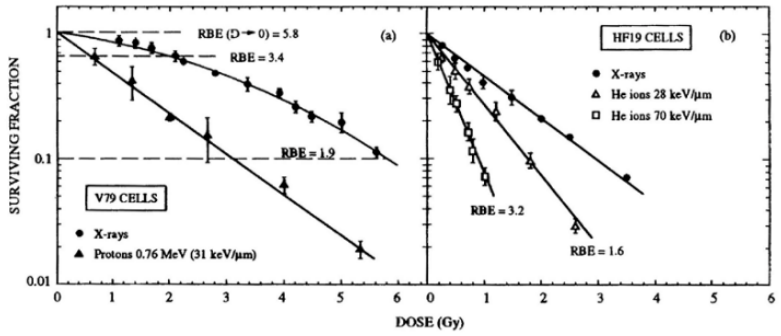

Figure 1.9: Cell surviving curves after the irradiation with X-rays, protons and He ions.

To take into account all the factors that play a role in the effectiveness of different radiation qualities, radiobiology makes use of a quantity called *relative biological effectiveness* (RBE), defined by the following formula:

$$RBE = \frac{D_{ref}}{D} \tag{1.5}$$

where D is the dose delivered by the radiation field under investigation, and D_{ref} is the dose imparted by a reference photon field (generally X-rays) that produces the same biological effect as D. In most cases, the considered iso-effects are either the surviving fraction of cells or the extent of cell death.

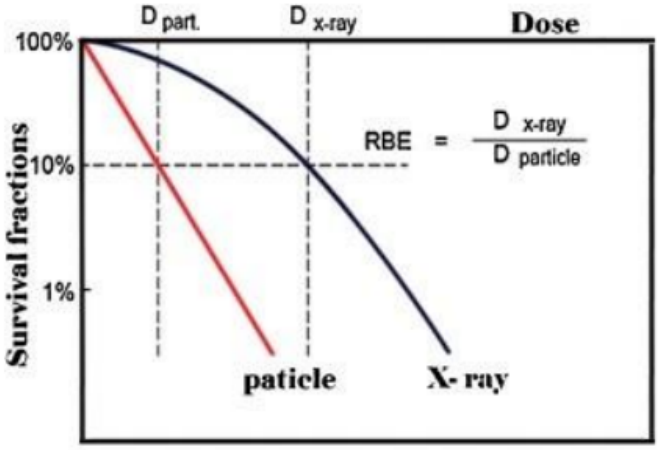

Figure 1.10: Cell survival curves and definition of RBE.

The steeper the slope of the survival curve, the higher the relative biological effectiveness of the particle. This implies that such radiation will be more effective in inducing cell death compared to photons. Consequently, to achieve the same biological effect, a lower dose is sufficient with respect to that required by the reference photon field. Light ions typically exhibit an RBE that reaches its maximum around the Bragg peak, while being considerably lower in the proximal region of the depth-dose curve. Thanks to this variation in RBE along the dose deposition profile, hadrons are able to deliver only a modest dose at the entrance region, while releasing a much higher and more biologically effective dose corresponding with the target volume.
As a result, hadron therapy proves significantly more advantageous than conventional radiotherapy, not only in terms of dose distribution but also from the perspective of biological effectiveness.

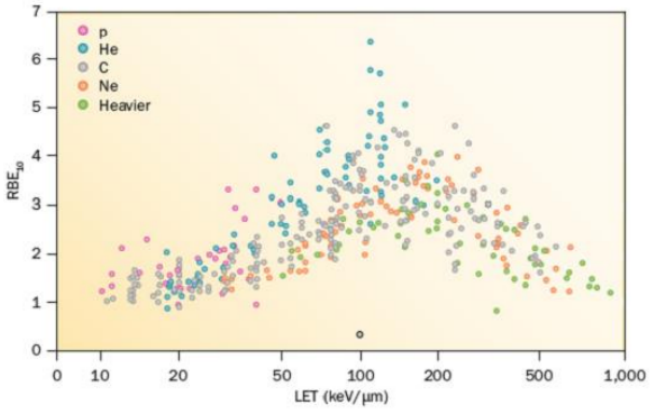

Figure 1.11: RBE as a function of LET, for a surviving fraction of 10%, for different particles.

1.4.3 OER: Oxygen Enhancement Ratio

The presence of oxygen inside cells during treatment is a crucial factor, as it strongly influences the biological response to radiation. Oxygen improves cellular radiosensitivity: in normoxic conditions, the same radiation dose produces a higher biological effect compared to irradiation of the same tissue under hypoxic conditions.
To quantify the magnitude of this effect, the *Oxygen Enhancement Ratio (OER)* is defined as:

$$OER = \frac{D}{D_{ox}} \tag{1.6}$$

where D is the dose required to produce a given biological effect in hypoxic conditions, and D_{ox} is the dose that produces the same effect under conditions of full oxygenation.

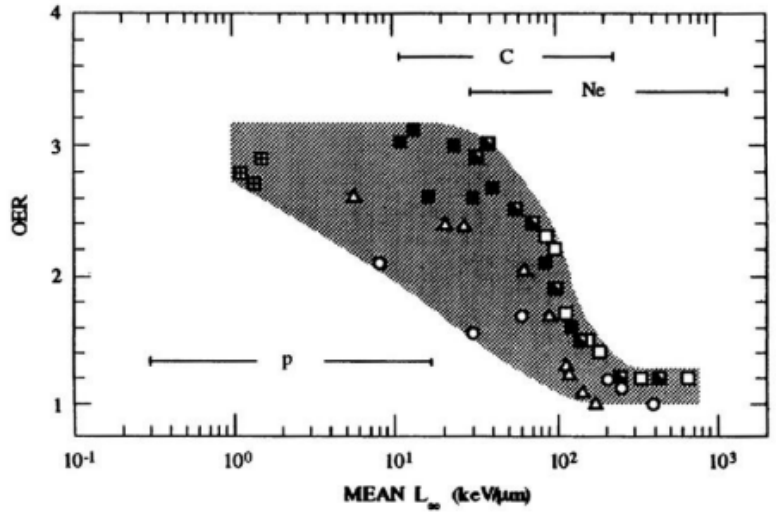

Figure 1.12: OER values as a function of LET at 10% survival for human cells.

For low LET radiation, such as photons, the oxygen effect is particularly pronounced. This makes the treatment of hypoxic tumors challenging, as the doses typically used in conventional X-ray therapies may not be sufficient; considerably higher doses would be required to achieve comparable biological effectiveness.
In contrast, for high-LET particles, the oxygen effect is much less significant. The required therapeutic dose is therefore almost independent of the oxygen concentration within the cells, allowing effective treatment of radio resistant tumors that cannot be adequately managed with conventional radiation therapy.

Chapter 2

CNAO

The CNAO center is one of the three hadron therapy facilities currently active in Italy and one of the six worldwide where both protons and carbon ions are employed. In 2024, approximately 542 patients were treated at the center: 264 with carbon ions and the remaining 278 with protons.[10]
However, CNAO is not only a clinical facility but also a research institute. In parallel with therapeutic activities, it conducts an extensive research program, with experiments spanning all the scientific areas involved in its operations. This dual vocation guarantees both the continuous development of innovative projects and the advancement of cutting-edge technologies for diagnostics and therapy.

2.1 The high-tech facility

The CNAO center is located in Pavia, in close proximity to major hospital facilities such as the San Matteo Hospital and the Mondino Foundation, as well as to the scientific hub of the University of Pavia. This strategic location ensures strong synergy from both clinical and research perspectives.
The building comprises several functional areas: the outpatient clinic for specialist consultations, the imaging and nuclear medicine departments, the section dedicated to treatment delivery (including patient preparation and immobilization rooms, as well as three treatment rooms), the administrative department with staff offices, and the technological areas housing the equipment essential for the operation of the facility.
At the heart of the center lies the accelerator — a synchrotron with a diameter of 25 meters — which allows the acceleration of protons to energies between 60 MeV and 250 MeV, and carbon ions to energies between 120 MeV/nucleon and

400 MeV/nucleon [10]. Once the desired energy is reached, the beam is extracted and transported to one of the three treatment rooms, where it is precisely delivered to the patient. In this context, patient positioning is also an essential part of treatment, which is performed with sub-millimetric accuracy through three independent monitoring systems.
The beam is transported from the synchrotron ring to the treatment rooms via four branches of the high-energy transport line (three horizontal and one vertical).
The patient's dose monitoring is carried out at the end of each beamline through a *Dose Delivery System (DDS)*, which provides real-time measurements of the number of particles delivered and the exact position of the beam.

Figure 2.1: Left: entrance of the CNAO center; right: model of the CNAO complex.

The synchrotron hall and the treatment rooms are located inside an underground bunker, taking advantage of the natural shielding provided by the soil. This is complemented by thick concrete walls built around treatment areas to ensure compliance with the radio protection limits established by law.
In addition to the clinical treatment rooms, the center also includes an experimental area equipped with a dedicated beamline, used for research activities in various scientific and technological fields.

2.2 The Accelerator and the Transport Lines

The prototype of the CNAO accelerator comes from the European project PIMMS (Proton Ion Medical Machine Study), which focuses on the design and planning of a synchrotron ad hoc for use in the medical field for tumor treatment. The same prototype has been replicated in the construction of MedAustron, the twin center of CNAO rising in Vienna.

The high technology of CNAO is made of different accelerating machines and transfer lines for the particles beam. In sequence the main components are:

- ion sources;
- LEBT line (*Low-Energy Beam Transport*);
- RFQ (Radio Frequency Quadrupole);
- LINAC (LINear ACcelerator);
- MEBT line (Medium-Energy Beam Transport);
- Synchrotron Ring;
- HEBT line (High-Energy Beam Transport);

The scheme of the acceleration chain is represented in Fig. 2.2.

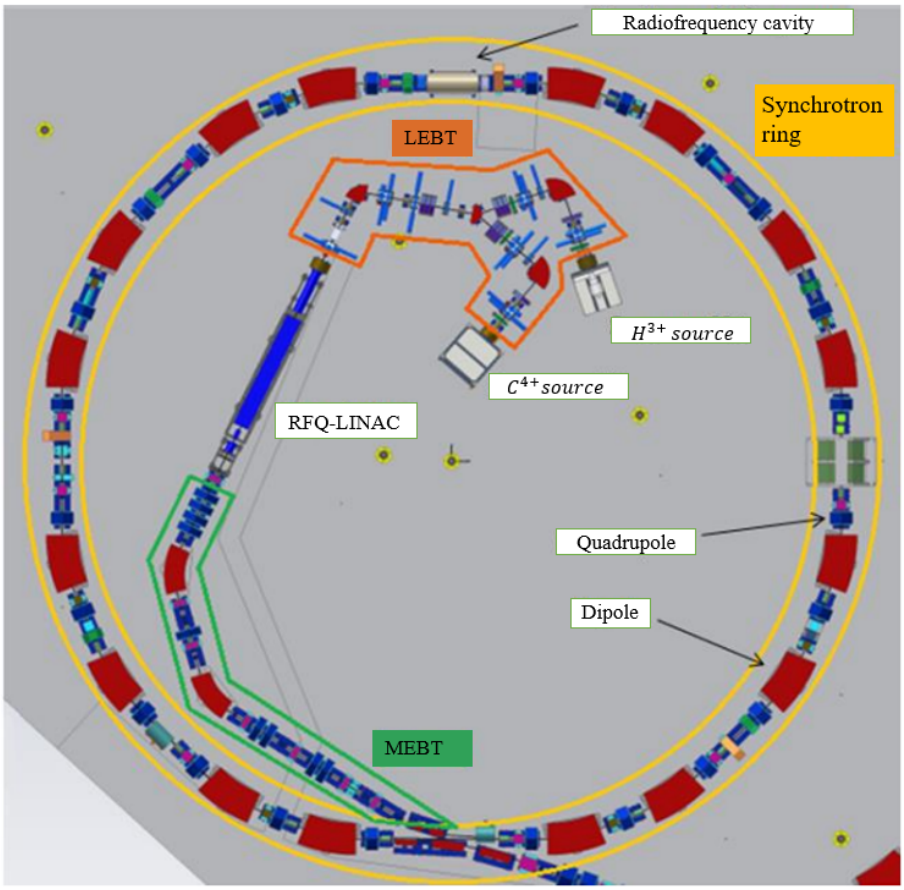

Figure 2.2: Scheme of the injector and of the principal accelerator present at CNAO.

The beam originates from one of the two ECR (*Electron Cyclotron Resonance*) sources. By simply varying the type of gas used, together with the RF settings and the high-voltage parameters, each source can produce different ion species, thus allowing the selection of the kind of particle to be accelerated. In case of malfunction or unavailability of one source, the other can compensate for its absence, ensuring continuity of operation.
Once the source is operating, the injected gas is heated until it reaches the plasma state.
At CNAO, two types of ions are currently accelerated: carbon and protons, produced respectively from CO_2 and H_2, from which C^{4+} and H_{3+} ions are extracted.
Once produced by the sources and after a partial ionization, the particles enter the LEBT (*Low Energy Beam Transport*), the low-energy line that connects the sources to the Radio Frequency Quadrupole (RFQ). This section of the accelerator is responsible for separating the desired ion species using a spectrometer, selecting the beam component to be accelerated through a chopper, and adapting the beam characteristics for optimal injection into the RFQ.
The Radio Frequency Quadrupole is located immediately upstream of the LINAC in the acceleration chain. Its function is to pre-accelerate the beam from 8 keV/nucleon to 400 keV/nucleon, preparing it for injection into the LINAC, the Radio Frequency Linear Accelerator, where the particles reach an energy of 7 MeV/nucleon. This represents the main acceleration stage before the synchrotron ring. The LINAC can accelerate all particles with a charge-to-mass ratio $q/A \geq 1/3$, where q is the ion charge and A is the mass number.
Downstream of this component, the beam passes through a *stripping foil*, which removes the remaining electrons from the ions, allowing complete ionization of the particles before their injection into the synchrotron.
At this point, the MEBT line transports the beam up to the entrance of the synchrotron. Its role is to control the emittance — a quantity describing the size and divergence of the beam cross-section — to adapt the optical functions, optimize the dispersion, and vary in a controlled manner the intensity of the current injected into the synchrotron, typically around 0.7 mA for protons and 0.15 mA for carbon ions.
At this stage of the sequence, if required by the treatment plan, it is possible to reduce the beam intensity to 10%, 20%, or 50% of its nominal value by inserting a *degrader*, a grid with different levels of transparency that allows only a defined fraction of particles to pass through. This adjustment is essential whenever low doses need to be delivered to tumor tissues.

The main component of the acceleration chain is the 25 m-diameter synchrotron, a circular accelerator in which the particles travel along a closed circular orbit. The particles enter the synchrotron with the energy reached in the LINAC and, by cyclically circulating within the ring, are further accelerated to the desired energy - from a minimum of 60 MeV for protons to a maximum of 400 MeV/nucleon for carbon ions. The electromagnetic field within the accelerator varies periodically in synchrony with the beam to keep the particles on the nominal orbit as their velocity increases. [11]

Along the synchrotron ring, several magnets are installed: 16 dipoles for beam deflection and 24 quadrupoles for focusing, arranged to form two symmetrical achromatic arcs connected by two straight, non-dispersive sections. In one of these sections, the beam is injected from the MEBT line, while in the other, the radio frequency cavity and the resonance sextupole are located.

The injection into the ring takes place by creating a *bump*, a local distortion of the orbit that brings the beam trajectory closer to the electrostatic septum used for injection. At each turn, the distortion is progressively reduced to zero, so that the beam gradually fills the entire transverse section of the synchrotron (*multi-turn injection*).

In the opposite non-dispersive section, the radiofrequency cavity is located, where the beam is accelerated at every turn.

The extraction process represents the most delicate phase of beam delivery, since the beam transported to the treatment rooms must exhibit a high degree of uniformity in both spatial and temporal distributions. The beam is extracted from the synchrotron ring in discrete packages, referred to as *spills*, each with a typical duration of approximately 1 s. Although the circulating beam can be considered as a bunch with a temporal length shorter than 1 μs, it is gradually extracted through a *slow extraction* mechanism in order to spread the beam intensity over time and obtain a quasi-continuous spill.

In the original extraction scheme, slow extraction was achieved by means of a controlled betatron excitation provided by the *betatron core*. The betatron core induces an additional accelerating potential that slowly drives the beam toward an unstable region of the transverse phase space generated by the resonance sextupole. As a result, the particles begin to oscillate around the nominal closed orbit with increasing betatron amplitude. Once the resonance separatrix is crossed, the oscillation amplitude grows until the particles intercept the electrostatic septum, which deflects their trajectory away from the circulating orbit and directs them toward the High

Energy Beam Transport (HEBT) line.
At CNAO, this extraction mechanism has been replaced by the *Radio Frequency Knock-Out* (RFKO) technique. In RFKO extraction, a transverse radio-frequency excitation is applied to the circulating beam at frequencies close to the betatron tune, inducing a controlled diffusion of particle amplitudes in the transverse phase space. This process progressively drives particles toward the extraction resonance in a smooth and controllable manner. By appropriately modulating the RF frequency and amplitude, the extraction rate can be precisely regulated, resulting in a stable and uniform spill. Compared to pure betatron excitation, RFKO provides enhanced flexibility and improved spill quality, which are crucial for hadrontherapy applications requiring accurate dose delivery and rapid beam intensity modulation.
After the extraction from the ring, the beam is transported to the treatment rooms through the HEBT (High Energy Beam Transport) line. One of the main components of the HEBT is the *chopper*, consisting of four fast dipoles in series, which allows the beam to be stopped whenever patient irradiation is not required. Under normal conditions, the magnets of the chopper are turned off, and the beam hits a dedicated obstacle (*dump*) designed to absorb it completely. When the magnets are activated, the beam is deflected in order to bypass the dump and continue toward the treatment room. The response time of the chopper is less than 200 μs, ensuring rapid switching of the beam on and off, corresponding to the delivery of less than 2.5% of the prescribed dose for a single spot.

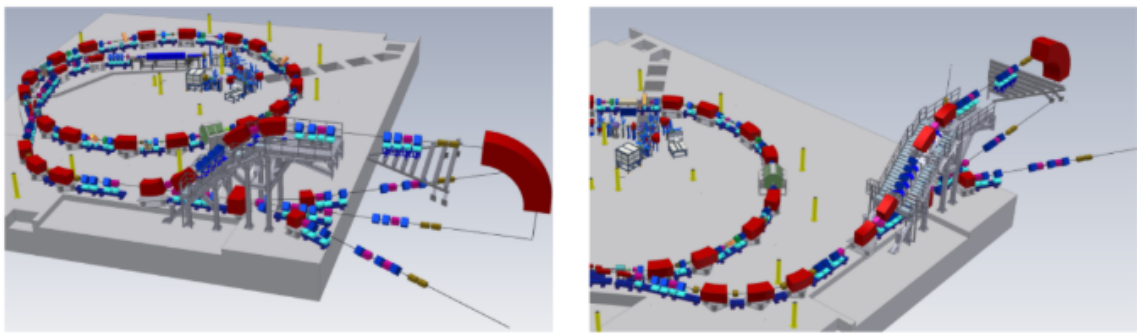

Figure 2.3: On the left: a complete model of the accelerator and of the transport line. On the right: a focus on the HEBT, showing the three horizontal lines and the vertical one transporting the beam to the treatment rooms.

After the chopper, the HEBT splits into four fixed beamlines: three horizontal and one vertical, which transport the beam to the treatment rooms where it is delivered to the patient. The lateral rooms are equipped only with horizontal beamline, while

the central one allows both vertical and horizontal irradiation. To realize the vertical beamline, four magnets are used to raise the beam to about 6 m above the synchrotron plane, and a large bending magnet at the end of the line deflects the beam by 90° onto the patient, as illustrated in Fig. 2.3.

The beam is directed into the desired line by means of *switching magnets*, which deflect it toward the appropriate path. The first switching magnet selects between the vertical and horizontal beamline: when activated, the beam is deflected toward the vertical line and directed to Room 2. When deactivated, the beam continues straight toward the second switching magnet, which then determines the distribution among the three horizontal beamlines.

Each of the four branches of the HEBT is equipped with a safety system called a *beam stopper*, which consists of a tungsten block designed to stop the beam in the event that it remains active when irradiation is not required. Under normal conditions, all beam stoppers are fully inserted, completely covering the transverse section of the beamline. When the beam is required in one of the treatment rooms, the corresponding beam stopper is lifted, allowing the beam to pass through.

The entire acceleration line operates under vacuum in order to minimize particle losses caused by interactions with air molecules.

It should be mentioned that, at CNAO, protons are accelerated to energies between 60 MeV and 250 MeV, with a maximum number of particles per spill of approximately 10^{10} ($\sim$ 2 nA). For carbon ions, energies between 120 and 400 MeV/nucleon are reached, with about 4×10^8 particles per spill ($\sim$ 0.4 nA). The extraction energies of the particles have been designed for clinical applications in tumor treatment: the beam can cover a range (i.e., the distance required to completely stop the particles) of up to 27 cm in water, which closely approximates the density of biological tissues, with energy steps of 1 mm. This configuration allows the treatment of deep-seated tumors. [11]

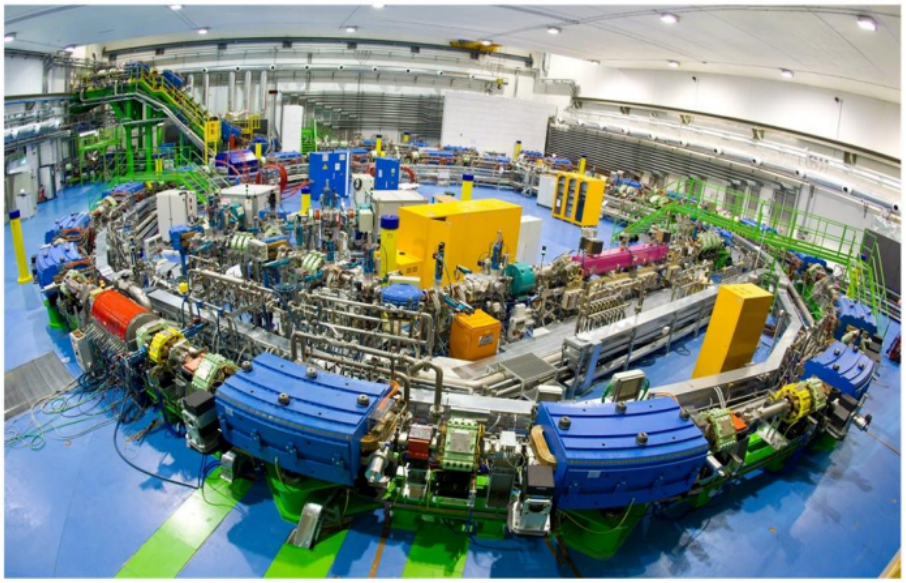

Figure 2.4: Picture of the synchrotron hall

2.3 The Research at CNAO

Since its foundation, CNAO has been conceived not only as a therapeutic center but also as a research facility. Alongside patient treatment, an extensive research activity is carried out across all the scientific fields involved in the center, with the goal of continuously improving the therapeutic effectiveness. Research at CNAO is mainly focused on three areas: clinics, radiobiology, physics and technology.

2.3.1 Clinical research

Studies in the field of clinical research aim to evaluate treatment outcomes for different types of tumors, evaluating the efficacy of hadron therapy both as a standalone treatment and in combination with other therapeutic modalities such as surgery or hormonal therapy.
The clinical studies conducted at CNAO are of both observational and interventional nature, allowing researchers to address various clinical questions and to identify the therapeutic approaches that are most effective for cancer treatment.
Observational research has no impact on the patient's treatment pathway. The CNAO Foundation collects clinical, physical, dosimetric, and psychophysical well-being data from patients who provide informed consent to participate in research activities. All data are stored in a dedicated registry, managed by an information system known as *REGAL (Registry Trial)*. The analysis of these data allows researchers to

identify prognostic and predictive factors of treatment response, enabling an accurate assessment of outcomes, including toxicity effects, even in the long term.
Interventional research is defined as any research activity performed on an individual in order to understand the causes, evolution, and treatment modalities of a disease, with the ultimate goal of improving the quality of life of both the individual under study and the general population.
Clinical studies are therefore conducted to achieve various objectives, ranging from prevention to diagnosis and therapy.
Clinical research is governed by well-defined regulations that must be followed to ensure the safety of the individuals involved, the protection of personal data, and the respect of human rights.
Several clinical observational studies are currently ongoing on the treatments delivered so far, aiming to highlight the state of the art in hadron therapy and to establish the basis for future interventional studies.
Clinical research at CNAO mainly focuses on phase II studies (which are designed to investigate the effects of a specific therapy on patients) in order to demonstrate the efficacy and safety profile of hadron therapy.
Some of the activities currently ongoing or soon to be activated include studies on pathologies of the gynecological district, such as melanomas in the mucosa of the lower female genital tract, and on lateral pelvic recurrences of gynecological tumors occurring after previous treatments. Other studies aim to assess the potential enhancement of the antitumor immune response following carbon ion radiation therapy on a single metastatic lesion in oligometastatic patients with stable disease undergoing immunotherapy.
CNAO also carries out national and international collaborations with various oncological institutions and medical centers that refer their patients for treatment.

2.3.2 Radiobiological research

The radiobiological research activity has mainly focused on the in-depth analysis of the biological effects of heavy ions and protons, with the objective of improving clinical effectiveness and promoting the personalization of oncological therapies.
Through *in vitro* studies, both in 2D and 3D cultures, as well as *in vivo* experiments, the radiobiology group investigates the cellular and molecular mechanisms responsible for the response to treatment, particularly those underlying tumor radioresistance.
Among the main research topics are studies on magnetic hyperthermia combined

with hadron therapy applied to 3D cellular cultures. The aim is to evaluate the effectiveness of the combined action of radiation and magnetic hyperthermia on biological and synthetic scaffolds containing pancreatic cancer cells and magnetic nanoparticles.

Research activities in this context encompass a wide range of studies aimed at understanding and improving the radiobiological response of different cancer types to various ionizing radiation modalities. These efforts include investigations into the effects of high- and low-LET radiation on cells with distinct genetic backgrounds, the development of advanced models for *in vitro* and *in silico* analyses, and the optimization of innovative therapeutic strategies, including combined treatments and novel dosimetry approaches. Collectively, these studies aim to enhance treatment personalization, improve radiotherapy efficacy, and provide a comprehensive understanding of radiation-induced biological effects across multiple scales.

More recent studies concern the biological characterization of BNCT (Boron Neutron Capture Therapy) at CNAO. From a radiobiological perspective, the characterization involves a preliminary experimental phase aimed at determining the RBE (Relative Biological Effectiveness) values of the neutron beam, which are specific for each AB-BNCT system.

Both *in vivo* and *in vitro* studies will be carried out to verify the safety and biological effectiveness of the system.

In *in vitro* experiments, the cytotoxicity of tumor and normal cells will be evaluated with and without boron administration, using different boron concentrations. The observed effects will be compared with those obtained by irradiating the same cellular models with photons, in order to calculate the RBE. Genotoxicity will be assessed through the quantification of micronuclei formation.

In vivo effects will be investigated in mice irradiated with neutron beams, both in the presence and absence of boron. Also in this case, the biological responses will be compared with those induced by photon irradiation to evaluate the corresponding RBE values.

This characterization, besides ensuring the safety of BNCT treatments at CNAO, will also enable the comparison of the radiobiological parameters of AB-BNCT systems developed worldwide.

2.3.3 Technical research

Research in this field is focused, among the other topics, on the design and development of a new particle accelerator to be installed in South-East Europe for both clinical and research applications. Many projects are being carried out in collaboration with international partners, including CERN, while others are coordinated directly by CNAO. Among these it is possible to find several technical activities, the most relevant of which concerns the development of a carbon-ion gantry to be coupled with the synchrotron.
Within this framework, Work Package 8 (Superconducting Magnet Design) includes the study of possible superconducting magnet solutions suitable for this purpose. Figure 2.5 shows a simplified scheme of the magnetic system for the gantry.

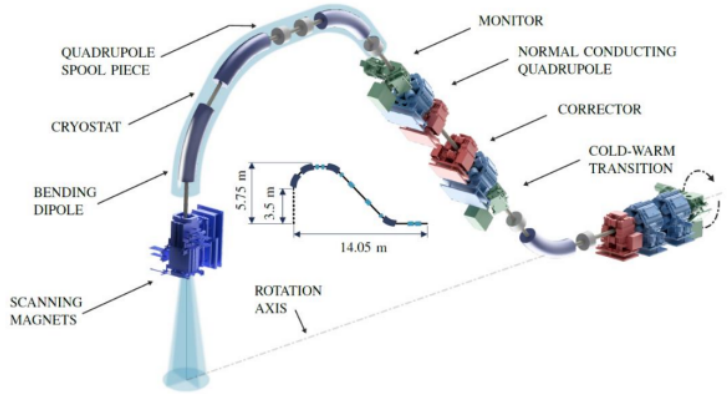

Figure 2.5: Simplified scheme of the optical and magnetic elements of the iso centric head.

CNAO is the project leader for the design and construction of a demonstrator of the scanning magnets for the gantry. In this context, a detailed study of the magnet design has been carried out.
With respect to beam optics, detailed studies conducted in synergy with mechanical analyses have led to a deeper understanding of orbital errors associated with systematic and random misalignment. These analyses are essential for improving the precision and reliability of beam transport in future accelerator facilities.
Regarding the CNAO upgrade, one of the main aspects concerns the use and evolution of the experimental room. This area is regularly employed for research activities, often in collaboration with external groups. Within this framework, several projects are currently underway for the implementation of new devices for irradiation experiments. In particular, the activities focus on the development of positioning systems,

low-intensity beam monitors, and a joint collaboration with GSI for the development of the next-generation *Dose Delivery System (DDS 4.0)*.
In order to allow sample movement during irradiation — a feature that can be particularly useful for certain experimental setups — a remote positioning system has been developed. This system includes a high-precision positioning stage operated from above, as shown in Fig. 2.6.

Figure 2.6: Remote positioning system for sample manipulation during irradiation.

Since some experiments and beam measurements must be performed in vacuum, a dedicated vacuum chamber has been developed, similar to the extension modules used when the irradiation position is modified within the experimental room. The new chamber, shown in Fig. 2.7, features several lateral ports of different sizes, compatible with the standard beam monitors employed at CNAO.

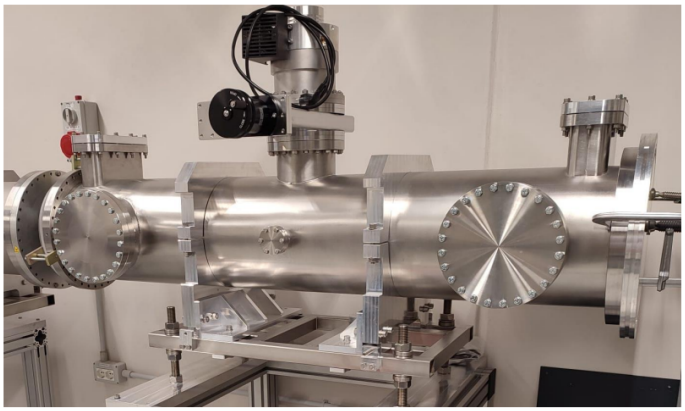

Figure 2.7: Vacuum chamber with lateral access ports for diagnostic devices.

In some experimental configurations, beams of much lower intensity than those used for clinical treatments are required. However, at such low intensities, the detectors of the standard *Dose Delivery System* are unable to perform reliable measurements. To overcome this limitation, a collaboration has been established with INFN and the University of Rome for the development of a dedicated low-intensity beam monitor. The device is based on scintillating fibers coupled to solid-state photomultipliers. Its geometry, illustrated in Fig. 2.8, is mechanically compatible with the CNAO nozzle, allowing the monitor to replace one of the two DDS boxes when required.

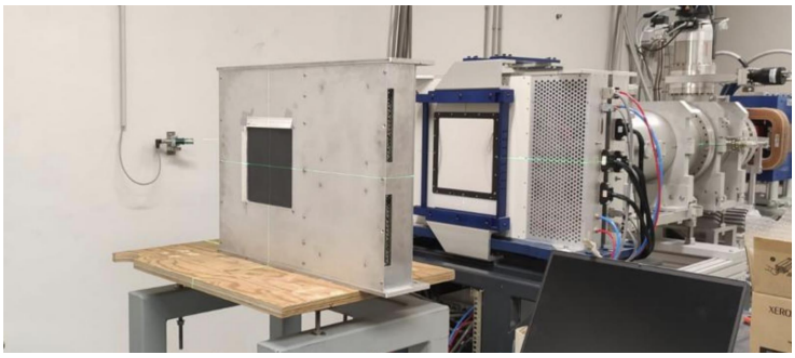

Figure 2.8: Low intensity beam monitor system during tests in the experimental room

In 2024, a new version of the *Dose Delivery System*, named *DDS 4.0*, was developed. The prototype, shown in Fig. 2.9, was designed to allow extensive laboratory testing before clinical implementation. The installation of the DDS 4.0 on the clinical beamline will represent an upgrade rather than a full replacement: the ionization chambers currently in use will be maintained, while the external communication modules and the power supply components will be replaced with the new system.

Figure 2.9: prototype locker for the DDS 4.0

The research activities carried out at CNAO—spanning clinical, radiobiological, physical, and technological domains—provide the scientific and methodological framework within which modern treatment planning must be understood. The continuous development of new irradiation techniques, advanced beam delivery systems, and innovative experimental tools naturally translates into the need for increasingly accurate and flexible computational platforms capable of supporting clinical decision making. In this context, the evolution of Treatment Planning Systems represents the bridge between research and clinical implementation: it is through these software tools that experimental knowledge, physical models, and technological advancements are integrated into optimized and patient-specific treatment strategies. For this reason, the following chapter introduces the principles, structure, and functionalities of Treatment Planning Systems, highlighting their role as the operational interface between CNAO's research activity and its clinical practice.

Chapter 3

Treatment Planning System

A Treatment Planning System (TPS) is a computer-based platform used to determine the optimal beam arrangements, energies, field size and fluence patterns to produce a safe and effective dose distribution for the treatment of oncological patients.
These software tools utilize anatomical and clinical data from the patient to calculate the optimal dose distribution, aiming to deliver the maximum dose to tumor tissues while minimizing exposure of surrounding organs at risk (OARs).
TPS are based on complex dose calculation algorithms and include functionalities for the segmentation of target volumes and OARs. Continuous technological development has enabled increasingly complex and personalized treatments, such as Intensity-Modulated Radiation Therapy (IMRT), Volumetric Modulated Arc Therapy (VMAT) and proton therapy, significantly improving clinical outcomes and patient safety.
Modern treatment planning systems are generally divided into three main components:
Beam Model: a computerized representation of the treatment beam, defined by energy distribution, collimators, filters, machine-specific geometry and beam modifiers. Its accuracy directly affects the ability of the TPS to reproduce the real behavior of the radiation beam.
Dose Calculation Engine: responsible for applying the beam model to a given patient or phantom geometry and accurately calculating dose deposition. Dose engines range from analytical algorithms to full Monte Carlo simulations of particle transport. The selection of the dose calculation engine is always a trade-off between computational accuracy and processing time.
Optimization Engine: determines the optimal arrangement of fields and beam modulation to generate the treatment plan. In conventional 3D radiotherapy, the user planner acts as the optimizer, iteratively adjusting parameters to reach an acceptable

plan. In contrast, IMRT and VMAT rely on inverse planning, where a computerized optimization algorithm automatically produces complex fluence modulation patterns starting from predefined clinical objectives and constraints.

3.1 Dose Calculation Algorithms

Accurate dose calculation within Treatment Planning Systems is essential to ensure an optimal and homogeneous dose distribution to the target volume while minimizing irradiation of surrounding healthy tissues. The precision of dose calculation depends on multiple factors, including patient anatomy, tissue heterogeneities, beam characteristics and the physical modeling of radiation–matter interactions.
Several types of dose calculation engines exist, each with different levels of complexity, computational cost and accuracy.
Pencil Beam algorithms are computationally fast and sufficiently accurate in homogeneous media. However, they become less reliable in the presence of significant anatomical heterogeneities, since lateral particle scattering and density variations are only approximated or neglected.
Convolution/Superposition algorithms (kernel-based) model dose deposition by convolving an energy deposition kernel representing the scattered radiation distribution with the incident fluence after ray-tracing corrections through the patient geometry. These methods account for density variations and typically provide clinically acceptable accuracy while maintaining reasonable calculation times.
Boltzmann Transport Equation solvers offer high accuracy by explicitly modeling particle transport and interactions in discrete energy and angular bins. While they approach the precision of Monte Carlo simulations, their level of discretization may still introduce some approximation errors.
The Monte Carlo method represents the gold standard for dose calculation, as it simulates particle transport at a fundamental level, fully accounting for photon–electron interactions and tissue heterogeneities. However, widespread clinical adoption has historically been limited by high computational cost and statistical noise, although improvements in hardware and variance reduction techniques are progressively increasing clinical feasibility.
The final accuracy of dose calculation also depends on several practical and physical factors, such as proper beam commissioning (including source modeling, attenuation characteristics, and fluence distribution), the spatial resolution of the dose calculation grid, and the presence of tissue interfaces or materials with significantly different

densities.
Any inaccuracies in dose computation may lead to either insufficient target coverage or excessive irradiation of organs at risk, potentially compromising both the effectiveness and safety of the treatment.

3.1.1 Kernel-Based Algorithms

Kernel-based algorithms use precomputed dose kernels and ray-tracing techniques to model energy deposition from radiation interactions within the patient. Dose is calculated by summing and appropriately scaling the contribution of kernels according to the local energy fluence.
A kernel represents the spatial distribution of deposited energy originating from a primary photon interaction site. Energy spreads laterally because secondary charged particles and scattered photons transport energy away from the interaction point. Kernels are usually precalculated using Monte Carlo simulations and are generally assumed to be radially symmetric.

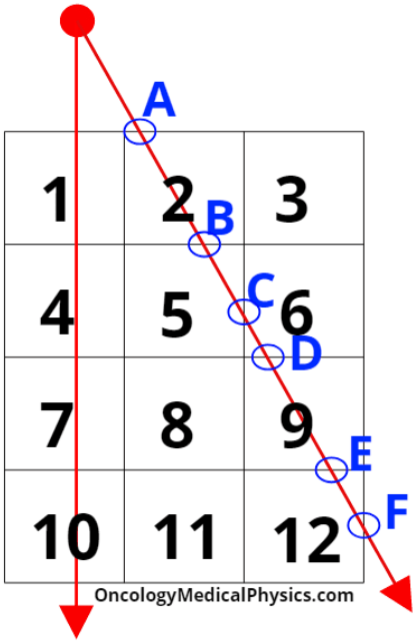

Figure 3.1: Schematic representation of ray-tracing through a voxelized patient model.

Ray-tracing algorithms are used in kernel-based models to compute the transport of primary energy fluence from the radiation source through the patient/phantom dataset. The process consists of:

- generating a ray from the source through the beam aperture (collimators, MLC, compensators, etc.) into the voxelized geometry;
- identifying the intersection points between the ray and each voxel boundary (blue points in Fig. 3.1);
- calculating the path length of the ray within each voxel: the longer the traversed segment, the greater the local fluence contribution.

Proper ray sampling is required to balance computational accuracy and calculation speed.

3.1.2 Pencil Beam Algorithms

Pencil Beam dose calculation is the simplest and fastest kernel-based method. It models energy deposition using one-dimensional line kernels and often neglects modifications in dose deposition caused by medium density variations. For this reason, Pencil Beam algorithms may perform poorly in highly heterogeneous regions such as the thoracic cavity (e.g., lung treatments).
The dose calculation proceeds as follows:

1. the treatment aperture is projected through the patient and subdivided into small beamlets (individual "pencils");
2. for each beamlet, a kernel-based dose distribution is calculated;
3. individual beamlet contributions are summed to obtain the final dose distribution.

This approach is fast and efficient for homogeneous anatomical regions but is typically unsuitable for complex geometries or advanced treatment techniques where high precision is required.

3.1.3 Kernel-Based Algorithms: Convolution/Superposition

Convolution/superposition algorithms compute dose using point kernels rather than line kernels, allowing local scaling according to tissue density. Multiple kernels are typically employed to represent different energy components of the clinical beam and to model beam hardening. The absorbed dose is obtained by convolving the kernel with the TERMA, which is derived from the attenuated photon fluence and the mass attenuation coefficient.

The method involves primary source modeling, collimator modeling through mask and aperture functions, inclusion of extra focal radiation, and ray tracing of photon fluence through the voxelized patient geometry. Fluence attenuation depends on material density and beam energy spectrum and is commonly handled using precomputed Fluence Attenuation Tables (FAT). The absorbed dose is calculated via convolution of TERMA with the point kernel; kernel tilting may be applied to better align the kernel with the beam direction, providing only modest accuracy improvements.
Tissue heterogeneities significantly affect scatter and dose deposition. These effects are approximated through superposition, where kernels are scaled using the radiologic distance between interaction and deposition points. Convolution can be implemented using direct summation, which is accurate but computationally intensive; Fast Fourier Transform (FFT) methods, which are computationally efficient but limited in heterogeneous media; or Collapsed Cone Convolution (CCC), which accelerates calculations by reducing the three-dimensional kernel to a finite set of ray-traced cones. Finally, electron contamination is added to account for surface dose contributions from electrons scattered in the treatment head of MV photon beams.

3.1.4 Boltzmann Transport Dose Computation

The Boltzmann Transport Equation describes the macroscopic behavior of radiation as it travels through matter.
To perform practical dose calculations using the Boltzmann Transport Equation, certain assumptions are generally made. Particles are assumed to interact only with the medium and not with each other, and the total number of particles emitted from the source is equal to the sum of those transported through the medium plus those absorbed.
The calculation process typically begins with the transport of the primary fluence from the source through the patient or phantom geometry, which is often performed using ray-tracing techniques. Once the primary fluence is established, the contribution of scattered photons is computed to account for energy deposited at secondary interaction sites. In addition, the fluence of scattered electrons is evaluated to include the dose contributions from secondary charged particles. Finally, the absorbed dose is determined from the total energy fluence, taking into account material-specific energy deposition coefficients.
Boltzmann-based dose computation methods are particularly valuable for heterogeneous media, providing a high level of accuracy that lies between simpler kernel-based

algorithms and full Monte Carlo simulations. This approach allows for a detailed modeling of particle transport and interactions, making it suitable for complex treatment planning scenarios.

3.1.5 Monte Carlo Dose Computation

The Monte Carlo (MC) method is a numerical approach that simulates physical processes using random number generation. In radiation therapy, it is employed to compute dose distributions by modeling the interactions of a large number of particles, such as photons, electrons, or protons, as they traverse a medium. Monte Carlo inherently involves statistical noise inversely proportional to the number of simulated particles, and achieving a calculation accuracy of approximately $\pm 1\%$ per voxel typically requires about 10^4 simulated particle histories. Due to its detailed atomic-level simulation of particle transport, Monte Carlo is the most accurate method for dose calculation, but also the most computationally intensive.

The simulation begins by generating a particle, whose initial direction is determined probabilistically based on the source characteristics. The distance to the particle next interaction is randomly assigned according to the linear attenuation coefficient of the medium it traverses. The particle is then transported to this interaction site using ray tracing. The type of interaction is probabilistically determined according to known cross-sections, and the interaction is simulated, which may involve energy deposition, scattering, or the release of secondary particles that are tracked in turn. This process continues until the particle energy falls below a defined threshold, at which point the remaining energy is deposited locally as dose.

One of the main challenges of Monte Carlo calculations is the required computation time. To address this, several acceleration techniques, collectively called variance reduction methods, have been developed.

A full Monte Carlo simulation would track electrons from the accelerator's bending magnets through all treatment head components—including the target, flattening filter, scattering foils, jaws—and through the patient or phantom. To save computation time, interactions in the treatment head are often pre-calculated and stored in phase space files. These files can become very large, potentially creating bottlenecks, which can be mitigated by simplifying the modeled interactions, applying variance reduction techniques, or using the phase space file to define a virtual source.

3.2 Optimization Algorithms

Treatment planning involves a complex balancing act between competing clinical objectives. For example, the target volume should ideally receive the full prescription dose, while surrounding healthy tissues and organs at risk (OARs) should receive as low a dose as reasonably achievable. The objective of plan optimization is therefore to determine the optimal compromise between these intrinsically conflicting goals while respecting treatment delivery constraints.

3.2.1 Optimization Parameters

One of the main elements used in optimization is the *cost function*, also referred to as the *objective function*. This parameter quantifies the extent to which a treatment plan satisfies its competing dosimetric objectives. The purpose of the optimizer is to minimize the cost function by iteratively modifying the dose distribution within the limits of clinical deliverability.
Another widely used tool for treatment plan evaluation is the Dose Volume Histogram (DVH), which graphically represents the relationship between the volume of a given anatomical structure and the radiation dose it receives. Essentially, it describes what percentage of a volume receives at least a certain dose value. DVHs allow verification of adequate target coverage while ensuring compliance with dose constraints for organs at risk.

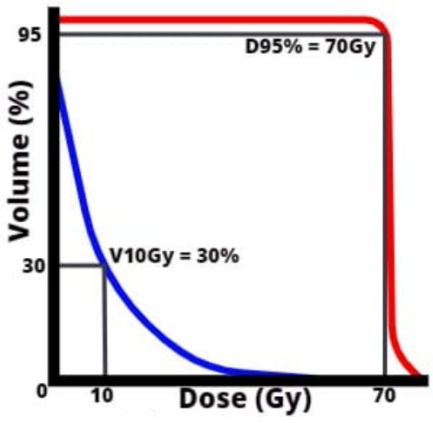

Figure 3.2: Example of DVH

The Equivalent Uniform Dose (EUD) is also employed as a metric during optimization. It is defined as the dose which, if uniformly delivered to the entire structure, would result in the same biological effect (in terms of cell survival) as the actual

heterogeneous dose distribution. EUD enables direct control of dose distributions by taking into consideration a specific dose value, d, and the structural characteristics of the organ, described by parameter α:

$$EUD(d) = \left[\frac{1}{N}\sum(d_i)^{\alpha}\right]^{1/\alpha} \tag{3.1}$$

The value of α depends on the radiobiological behavior of the organ. For parallel organs, α is typically set to 1, reducing the EUD to the mean dose. For serial organs, α may be set to a high value (e.g., 10), approximating a maximum dose constraint. Target volumes can also be constrained using EUD, commonly with negative α values to promote dose homogeneity. Serial organs are those that are particularly sensitive to radiation, since damage to a single portion of the organ can impact the entire structure, while parallel organs are organized in such a way that damage to one region does not necessarily compromise the function of the rest of the organ.

3.2.2 Common Optimization Algorithms

Several different optimization algorithms are used in treatment planning.
One of the most established methods is *Fluence Map Optimization* (FMO), which relies on gradient-based techniques to determine the optimal fluence distribution.
The advantages of FMO are mainly related to its rapid computation time compared to other optimization strategies. Furthermore, this kind of problem is typically convex: when the cost function and associated delivery constraints define a convex solution space, any local minimum is also the global minimum. Therefore, FMO is guaranteed to converge to the optimal fluence solution.
However, one important limitation is that machine-specific delivery constraints are not directly considered during the optimization process. As a result, the final dose distribution may not represent the best achievable plan that the machine could deliver.
A different strategy is *Direct Machine Parameter Optimization* (DMPO), also known as *Direct Aperture Optimization*. In this method, machine parameters such as MLC leaf positions and segment weights are included directly as optimization variables. In VMAT implementations, additional delivery parameters—including gantry rotation, couch motion and collimator angles—are also incorporated from the start. This strategy leads to treatment plans that are inherently deliverable and can achieve dose distributions closer to the true clinical optimum compared to those produced

by FMO.

However, the inclusion of machine constraints makes the optimization problem nonlinear and nonconvex. In such cases, multiple locally optimal solutions may exist that are worse than the global optimum. Gradient-based methods therefore become less reliable, as the search process may become trapped in suboptimal local minima, where small variations in machine parameters degrade the solution—even though a more substantial change could lead to a better configuration.

To overcome this issue, various mathematical strategies have been adopted in DMPO-based approaches. One of these is *simulated annealing*, which introduces controlled randomness during the optimization, enabling the algorithm to escape local minima by occasionally accepting worse intermediate solutions. Another class of techniques is based on *column generation*, where the overall optimization problem is decomposed into smaller subproblems, such as optimizing individual aperture shapes, while segment weights are updated globally. Approaches based on gradients continue to be used to guide the optimization trajectory, although they are less effective when the problem landscape is highly nonconvex. Finally, *genetic algorithms* provide an alternative, based on an evolutionary model in which a population of candidate solutions is iteratively refined through mutation and selection, progressively improving the plan quality until acceptable convergence is reached.

3.2.3 Multi-Criteria Optimization (MCO)

Traditional plan optimization approaches—such as FMO and DMPO—rely on manually assigned weights for DVH-based objectives in order to attempt to obtain an optimal plan. Unfortunately, these objectives and weights are often not directly associated with the true clinical goals. This makes conventional optimization a time-consuming trial-and-error process. Multi-Criteria Optimization (MCO) attempts to overcome these issues by linking clinical intent more directly to the optimization process.

A first approach within MCO is *Prioritized Optimization*, in which clinical objectives are ranked based on their importance. For example: a maximum dose to the spinal cord of $D_{max} < 45\,\mathrm{Gy}$, a target coverage for the PTV of $D_{95\%} = 70\,\mathrm{Gy}$, and a constraint for the larynx of $V_{58\mathrm{Gy}} < 27\%$.

In prioritized optimization, the optimizer attempts to satisfy higher-priority goals without compromising them while working to improve objectives with lower priority. This method is particularly useful when the user can clearly pre-define which trade-

offs are clinically acceptable. However, a main limitation is that it may ignore large potential improvements in lower-priority objectives even when only a minimal trade-off in a higher-priority goal would be required—resulting in solutions that may not fully exploit the achievable plan quality.

Another essential method in MCO is *Pareto Optimality*. A plan is defined as Pareto-optimal when no objective can be improved without simultaneously worsening another objective. In other words, every Pareto-optimal plan represents the best achievable compromise without making an undesired trade-off. The complete set of such solutions is called the *Pareto surface*.

A Pareto surface is typically generated by Direct Machine Parameter Optimization, producing a range of optimal treatment plans that lie along this surface. The clinically preferred plan is assumed to be located somewhere on the Pareto-optimal set. Therefore, MCO workflows center on generating, exploring, and selecting among plans along the Pareto surface.

The interactive evaluation of trade-offs among Pareto-optimal plans is defined as *navigation over the Pareto surface*. Two common strategies are used for this purpose: the *library method* and the *graphical user interface (slider) method*.

In the library method, a limited subset of Pareto-optimal plans is pre-computed and presented to the planner, who can select the plan that best fulfills the clinical needs. If necessary, interpolation between selected plans can be performed to generate additional deliverable options.

In the graphical user interface method, a set of sliders is used to continuously explore the trade-offs between competing clinical criteria. In this way, the planner can directly visualize how improving the dose to one structure impacts the dose to others and dynamically choose the most clinically appropriate solution.

3.3 Commercial and academic TPS

Various TPS are available, both commercial and for academic or research purposes. In general, commercial software can be considered as "black boxes," since their algorithms are not always accessible to the user. Conversely, academic or open-source systems, despite being highly flexible and customizable, are not certified for clinical use and therefore cannot be employed to generate treatment plans for actual patients.

The most widespread commercial TPS are produced by leading companies in the radiotherapy sector, such as Varian Medical Systems with *Eclipse*, Elekta with *Monaco*, and RaySearch Laboratories with *RayStation*. These platforms are widely

used in hospitals and oncology centers worldwide for the planning of conventional treatments, IMRT, Volumetric Modulated Arc Therapy (VMAT), and more recently, Image-Guided Radiation Therapy (IGRT) and hadron therapy. Each system is characterized by its specific dose calculation methods, including the use of advanced *Monte Carlo* algorithms.

On the other hand, in the field of research and development, non-commercial open-source software such as PLUNC, FoCa, matRad, CERR, and TOPAS are widely used. These systems offer a high degree of customization, allowing direct access to the calculation modules. Such tools are particularly employed in academic contexts and research centers to validate new algorithms for dosimetric analysis, optimization, and dose calculation.

Although they cannot be used in clinical settings, these platforms represent a fundamental resource for innovation in radiation therapy, bridging the gap between clinical practice and research by enabling the testing and development of novel methodologies and computational approaches.

In the following paragraph some of the most used TPS software are explained in detail.

3.4 Commercial TPS

3.4.1 Eclipse

The *Eclipse Treatment Planning System* is a software platform developed by Varian Medical Systems to assist clinicians in planning radiation therapy for patients with malignant or benign diseases. It utilizes medical imaging data to create a virtual patient model, enabling the simulation and optimization of radiation dose distributions for both external beam and internal radiation treatments. The system is designed to ensure effective tumor targeting while minimizing exposure to surrounding healthy tissues.

Eclipse also includes dedicated modules for proton therapy and employs advanced optimization algorithms that enhance treatment quality and delivery efficiency.

The most recent versions of Eclipse incorporate several advanced tools, such as knowledge-based planning through the *RapidPlan* module, multicriteria optimization (MCO) with interactive isodose dragging, and the possibility of custom scripting via the Eclipse Scripting API (ESAPI). For photon dose calculation, it employs the *Acuros XB* algorithm, which provides accurate modeling even for small fields.

Its main strengths include its widespread adoption, which facilitates user training and the standardization of clinical protocols. The strong integration with the Varian software suite allows a seamless workflow, reducing potential interface errors. Moreover, Eclipse offers several advanced functionalities — such as RapidPlan, MCO, scripting, and dedicated proton therapy models — that enable the creation of sophisticated treatment plans, automation of processes, data extraction, and workflow customization.

Despite its advantages, some limitations must also be acknowledged. First, the analytical dose calculation algorithms show reduced accuracy in the presence of strong tissue inhomogeneities. This includes suboptimal modeling of large air gaps and the simplified handling of aperture edge scattering effects. The increasing availability of Monte Carlo–based dose calculation tools is expected to mitigate several of these limitations, particularly those related to scattering and secondary particle transport. Another limitation is the current inability of the Eclipse TPS to compute monitor units for passive scattering proton beam fields. However, this constraint is not unique to Eclipse, as similar limitations are still present in several other treatment planning systems used for proton therapy.

In conclusion, Varian Eclipse is a versatile and robust TPS that has significantly improved over the past years. Nevertheless, an important limitation remains its focus on proton dose calculation only, without support for carbon ion therapy—an emerging modality whose integration is increasingly demanded in modern treatment planning environments.[12]

3.4.2 RayStation

RayStation, developed by RaySearch Laboratories, is one of the most advanced and versatile treatment planning platforms currently available, characterized by a strong orientation toward automation and multimodality support for photons, protons, and ions. It also integrates cutting-edge functionalities such as multi-criteria optimization (MCO), adaptive planning, and machine learning.

The robust management of beam angles, patient data, and photon/proton combinations, together with seamless integration with the Oncology Information System (OIS) *RayCare*, makes RayStation an extremely flexible and efficient system.

Among the main strengths of this software are its high flexibility, supported by a wide range of delivery techniques, treatment machines, and functional modules. Automation and integration of artificial intelligence - particularly the implementation

of machine learning and deep learning algorithms for patient geometry segmentation - improve plan reproducibility and significantly reduce computation times. The system also provides extensive support for data extraction and personalized workflow configuration.
However, to fully exploit the advanced features of RayStation, high-performance hardware and adequate clinical support are required, which may lead to increased operational costs. Additionally, some of the most recent functionalities have not yet been fully validated through clinical publications. Finally, in multimodal contexts, the management of software licenses for the various modules can become rather complex.[13]

3.4.3 Monaco

Monaco, developed by Elekta, is a treatment planning system that utilizes a Monte Carlo-based dose engine, which is considered the gold standard for dose calculation accuracy, particularly in the presence of significant anatomical heterogeneities and when advanced treatment techniques are employed.
The latest version consolidates the platform's capability to support VMAT, IMRT, and 3D dose calculations within a single environment. Calculation times have also been optimized, achieving speeds up to four times faster than previous versions.
The main advantages of Monaco include the high accuracy provided by the Monte Carlo dose engine and the versatility supported by a wide range of treatment techniques.
However, despite the improved calculation speed, Monte Carlo simulations still require substantial hardware resources and precise commissioning. Additionally, the user interface and learning curve are considered less intuitive compared to other systems. Customization and automation options may also be somewhat less flexible than in other modern TPS platforms. [14]

3.5 Non-commercial TPS

Non-commercial TPS represent a fundamental resource for research and academic development in radiation therapy. Unlike commercial systems, these software tools provide full access to algorithms and to the internal structure of the code, enabling experimentation, validation of new methods, and customization of treatment plans according to specific research needs. Their flexibility is particularly valuable for the

advancement of innovative techniques such as experimental IMRT, adaptive planning, proton therapy, and heavy-ion therapy.
Here, some of the most widely used software are presented.

3.5.1 PLUNC

PLUNC (Planning for Linear and Non-Linear Conformal treatments) is one of the earliest TPS developed in an academic context, originally designed for 3D conformal treatment planning. It enables the simulation of photon, electron, and neutron beams through pencil-beam based algorithms and supports multi-beam optimization. PLUNC has been extensively used for methodological studies focused on field configuration and dose distribution optimization.

3.5.2 FoCa

FoCa (Forward and Convolution Algorithm) is a TPS developed for academic purposes, centered around convolution-based dose calculation and advanced beam modeling. It supports photon dose computation using convolution and transport algorithms, allowing accurate simulations in heterogeneous media. FoCa is widely employed in medical physics laboratories for benchmarking and evaluating innovative dose calculation methods.

3.5.3 CERR

CERR (Computational Environment for Radiotherapy Research) is an open-source MATLAB-based environment designed for visualization and analysis of radiotherapy data. Although it is not a fully independent TPS, it allows importing patient and dose information from commercial TPS, simulating custom treatment plans, comparing dose distributions, and performing DVH-based assessments for research purposes.

3.5.4 TOPAS

TOPAS (TOol for PArticle Simulation) is an open-source simulation framework based on Geant4, specifically developed for Monte Carlo simulations in proton and heavy-ion therapy. It enables highly accurate dose calculations in complex geometries, detailed modeling of clinical beamline components and materials, validation of novel optimization algorithms, and comparison with experimental measurements. TOPAS

is widely utilized in advanced proton therapy research, beam commissioning studies, and out-of-field dose evaluations.

However, most currently available open-source solutions focus on only specific steps of the radiation treatment planning workflow, such as image processing, plan evaluation, or provide dose calculation and optimization capabilities limited to a single radiation modality.

In the next chapter, the last non-commercial TPS considered in this work, *matRad*, is described. It is an open-source, multi modality software framework for radiation treatment planning.

Chapter 4

MatRad

MatRad is an open-source dose calculation and optimization toolkit for treatment planning in intensity-modulated photon, proton, and carbon ion therapy. It has been developed for educational and research purposes and is entirely written in `MATLAB`. MatRad provides a high degree of accessibility and flexibility, the software package comprises, besides the source code itself, anonymized examples of patient data, as well as physical and biological base data.
In particular, matRad provides a comprehensive set of functionalities that cover the entire radiotherapy treatment planning workflow. These include ray tracing, dose calculation for photons, protons, and carbon ions, inverse planning based on both physical dose and biological effect, multileaf collimator (MLC) sequencing, as well as tools for basic treatment plan visualization and evaluation.

4.1 Base Data

The matRad source code includes base data modeling generic radiation devices for photons, protons, and carbon ions. Base data for each radiation modality are stored as a single file in MATLAB's native format and are loaded dynamically in matRad depending on the current treatment planning scenario.
Base data for photon dose calculation algorithms are obtained from the clinically approved photon dose calculation engine PDC++, and describe decomposed lateral scatter kernels and depth-dose components for a 6 MeV linear accelerator for multiple source-to-surface distances (SSDs).
For protons, analytical calculations are used to approximate integrated proton depth–dose profiles in water for 114 proton beam energies ranging from 31.72 MeV to 236.1 MeV, corresponding to Bragg-peak positions from 7 mm to 347 mm. Lateral

beam broadening due to multiple Coulomb scattering is calculated using the Molière formalism as described by [15] and is added quadratically to the initial beam widths. Additional broadening of the initial beam width in air is not considered, unless the corresponding lookup tables are provided.
Both depth–dose profiles and lateral beam properties are stored for each initial beam energy as a function of water-equivalent depth.
Regarding carbon ions, integrated depth–dose profiles are based on FLUKA Monte Carlo simulations describing interactions of carbon ions and their fragments with a generic beamline including a 3 mm ripple filter. These Monte Carlo simulations were performed in water.
In total, 121 carbon ion depth–dose profiles linearly spaced from 115.23 MeV/u to 398.84 MeV/u, corresponding to Bragg-peak positions from 32 mm to 294 mm, are included in the base data file. Lateral spread in water is also taken into account.
To enable the planning of biological treatment, each depth-dose profile of monoenergetic carbon ions is complemented by dose-averaged α and β profiles as a function of depth in water, describing the properties of biological beam according to the Linear Quadratic Model (LQM). The biological base data were obtained from the 4th version of the Local Effect Model (LEM IV). The validity of these biological data is limited to therapeutic doses below 10 Gy(RBE). For the calculation of dose-averaged α and β profiles, synergetistic effects originating from mixed radiation fields are considered.
Integrated depth–dose curves for protons and carbon ions are given in $\mathrm{MeV} \cdot \mathrm{cm}^2 \cdot \mathrm{g}^{-1}$ and converted within the particle dose calculation to $\mathrm{Gy} \cdot \mathrm{mm}^2$ per 10^6 primaries. Since the lateral components of the dose calculation exhibit a unit of mm^2, the weight of each pencil beam directly corresponds to the number of particles used (10^6), while the resulting dose is expressed in Gy.
It may be useful to add custom base data, for example to simulate additional particle species or alternative beamlines. This can be implemented by creating a customized machine structure as described in Chapter 5.

4.1.1 Patient data

The matRad software toolkit comes with five anonymized open-source patient datasets for treatment planning which are stored in a custom Matlab file format (*.mat). Besides a cubic box phantom, four additional cases (prostate, liver, head and neck, and TG119 phantom) stemming from the Common Optimization for Radiation

Therapy (CORT) dataset, are included in the public base code. And there is also the possibility to customize the phantom according to one's research necessities.
The CORT dataset is an open-source dataset intended to be used by researchers when developing and benchmarking radiation treatment planning algorithms. Each dataset contains computed tomography (CT) images (*ct*) as well as segmentation and predefined treatment planning parameters (*cst*).
For dose calculation the original CT image is down sampled to a lower resolution, that can be selected by the user through the command *ctResolution = []*. The initial volume is thus divided into cubes called *voxels*. These elements are required for dose distribution computation and are useful for implementing objective functions.

4.1.2 Interfaces (DICOM import)

matRad features a DICOM (Digital Imaging and Communications in Medicine) import module allowing for the conversion of DICOM data (CT and RT structure set) into matRad's custom patient data format (ct, cst).
The matRad import functionality requires the MATLAB Image Processing Toolbox since it relies on built-in functions to import DICOM data. The DICOM import can either be initiated via the graphical user interface (GUI) or by manually calling the corresponding function.
During DICOM import, the grid resolution for the resulting CT cube (ct) can be adapted. By default, the original CT resolution is used; resolution adaptation applies trilinear interpolation. It is possible to define separately the dose grid resolution, which determines the dimensions of the dose matrix to be computed during dose calculation, and then matRad will interpolate it onto the CT grid.
In order to accurately account for tissue heterogeneities during radiotherapy treatment planning, Hounsfield Units of the CT scans are linked to electron densities (photons) or relative stopping powers in water (particles) by means of Hounsfield Lookup Tables (HULT). If no HULT is provided, a default lookup table is used instead.

4.2 Workflow

matRad models the radiation therapy treatment planning workflow by sequential function calls which can either be triggered using MATLAB's command line prompt or matRad's GUI.
The first step of the workflow is to load patient data consisting of a CT dataset

(*ct*) and segmentation (*cst*). Loading data can either be achieved via the DICOM or binary import interface, or by directly importing patient data stored in matRad's custom format. Next, treatment plan parameters are specified within a MATLAB structure (*pln*) to determine the radiation modality, beam angles, and number of fractions, among other settings. Afterwards, the beam geometry (*stf*), describing the placement of bixels/beamlets for photons or the individual pristine pencil beams for charged particles, is generated.
Once the beam geometry is defined, the dose calculation function can be called, returning a dose influence structure (*dij*) that contains the dose contribution from each bixel/pencil beam. Then photon and charged particle fluence distributions can be optimized according to user-defined objectives and constraints. In the case of treatment planning for photons, a multileaf collimator (MLC) sequencing algorithm can be applied to translate continuously optimized fluences into deliverable segments. Resulting dose distributions are saved in the variable *resultGUI* and can be analyzed in the GUI, which also supports DVH computation and additional dose statistics.

4.3 Definition of irradiation geometry

matRad adopts the LPS (Left, Posterior, Superior) coordinate system and models beam delivery by projecting the target volume into the beam's-eye-view (BEV) for each beam angle. Only the portion of the patient geometry contained within this projection contributes to fluence optimization.
The CT volume is therefore transformed into the corresponding BEV coordinate frame, where the target is projected onto the isocenter plane, perpendicular to the beam central axis. A fluence grid is then generated by sampling equidistant rays across the projected target area. The lateral spacing is controlled through the parameter *pln.bixelWidth*.
For photon beams, each ray corresponds to a discrete rectangular fluence element (bixel).
For proton and heavy-ion beams, additional spots are placed along the beam direction to ensure full target coverage in depth, with longitudinal spot spacing set via *pln.longitudinalSpotSpacing*.
The resulting beam steering and sampling information is stored in the MATLAB structure *stf*, which fully defines the irradiation geometry used for dose calculation.

4.3.1 The LPS coordinate system

In the LPS reference frame, the x-axis increases from right to left, the y-axis from anterior to posterior, and the z-axis from inferior to superior, as illustrated in Figure 4.1.

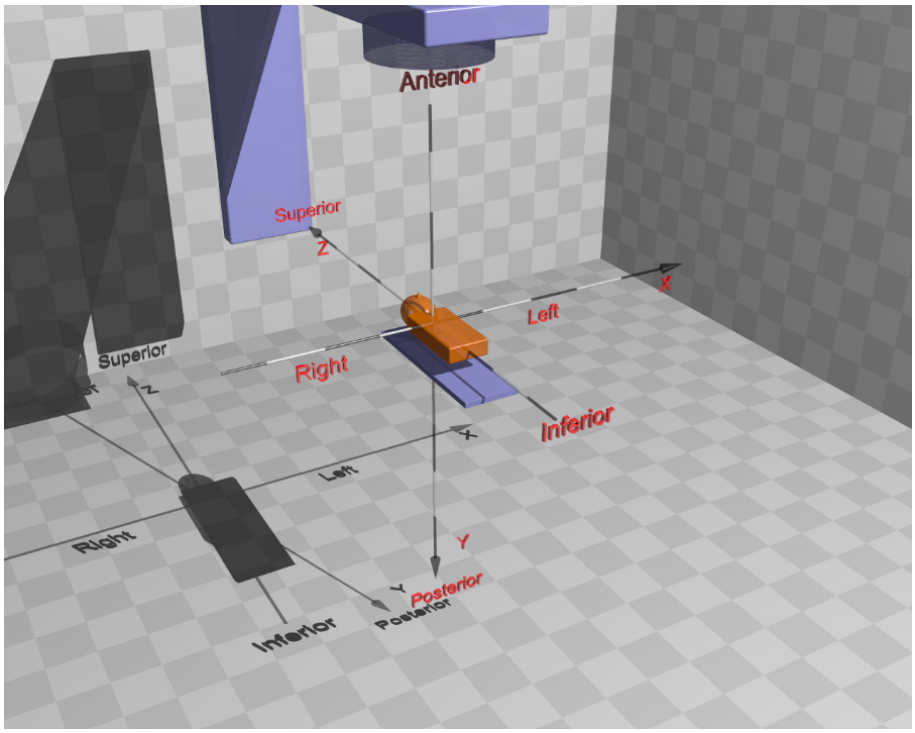

Figure 4.1: The LPS coordinate system

All spatial coordinates in the *stf* structure follow this convention. It is important to remember that CT and dose volumes are stored as 3D MATLAB arrays where the first dimension maps to y, the second to x, and the third to z, preserving the correct spatial orientation.

4.3.1.1 Coordinates of voxels

Voxel coordinates are defined such that the center of the most right–anterior–inferior voxel is located at $(\text{res.x}, \text{res.y}, \text{res.z})$. Consequently, the corner of that voxel closest to the origin is positioned at $(\text{res.x}/2, \text{res.y}/2, \text{res.z}/2)$ rather than $(0, 0, 0)$.

4.3.1.2 Gantry and couch rotation

Gantry rotation is modeled as a clockwise rotation around the z-axis, while couch rotation is defined as a counter-clockwise rotation around the y-axis, ensuring

consistent and clinically representative beam orientation handling.

4.3.2 Set treatment plan parameters

The first thing to set are the treatment plan parameters. We define them in the first section of the code. However, they can also be set in the matRad GUI.

```
%load HEAD_AND_NECK
load TG119.mat
%load PROSTATE.mat
%load LIVER.mat
%load BOXPHANTOM.mat

% meta information for treatment plan

pln.radiationMode   = 'photons';     % either photons / protons / carbon
pln.machine         = 'Generic';

pln.numOfFractions  = 30;

% beam geometry settings
pln.propStf.bixelWidth      = 5; % [mm] / also corresponds to lateral spot spacing for particles
pln.propStf.gantryAngles    = [0:72:359]; % [?]
pln.propStf.couchAngles     = [0 0 0 0 0]; % [?]
pln.propStf.numOfBeams      = numel(pln.propStf.gantryAngles);
pln.propStf.isoCenter       = ones(pln.propStf.numOfBeams,1) * matRad_getIsoCenter(cst,ct,0);

% dose calculation settings
pln.propDoseCalc.doseGrid.resolution.x = 5; % [mm]
pln.propDoseCalc.doseGrid.resolution.y = 5; % [mm]
pln.propDoseCalc.doseGrid.resolution.z = 5; % [mm]

% optimization settings
pln.propOpt.optimizer       = 'IPOPT';
pln.propOpt.bioOptimization = 'none'; % none: physical optimization;            const_RBExD; constant RBE of 1.1;
                                      % LEMIV_effect: effect-based optimization; LEMIV_RBExD: optimization of RBE-weighted dose
pln.propOpt.runDAO          = false;  % 1/true: run DAO, 0/false: don't / will be ignored for particles
pln.propOpt.runSequencing   = false;  % 1/true: run sequencing, 0/false: don't / will be ignored for particles and also triggered by runDAO below
```

Figure 4.2: Screenshot of the parameters section

You can import patient data or create your own phantom thanks to the function *matLab_PhantomBuilder.m*. Anyway, all the settings are specified in the *pln struct*.

4.3.3 The *pln* struct

The pln struct holds the information about the radiation treatment plan, it appears as in Figure 4.3.

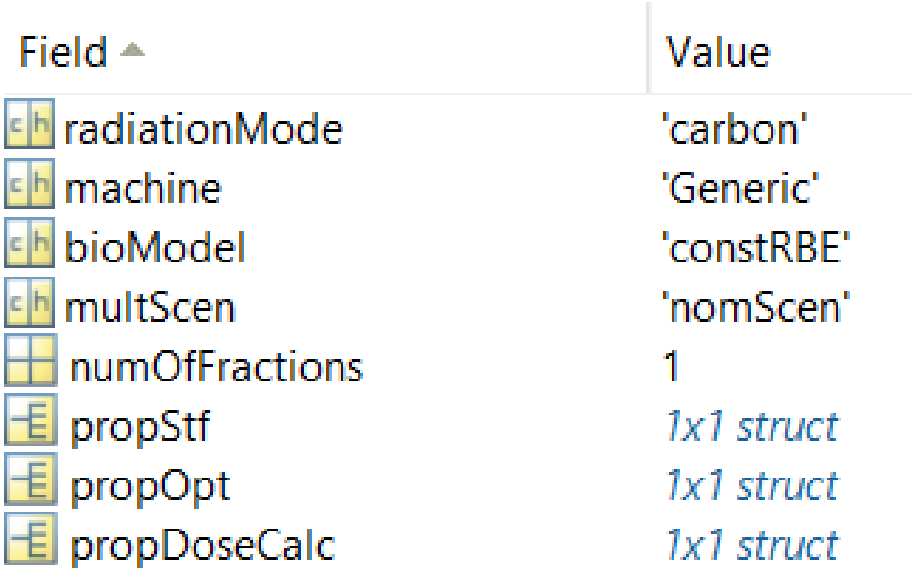

Figure 4.3: The pln struct

It contains different fields: *pln.radiationMode* which specifies the radiation modality, can be either *photons,protons* or *carbon*.

The *pln* struct defines the main parameters of the treatment plan, containing all information required for dose calculation and optimization. The field *pln.radiationMode* specifies the radiation modality, which can be either *photons*, *protons*, or *carbon*. In order to load the appropriate base data, it is necessary to define the treatment machine in the field *pln.machine*. The code automatically searches for the corresponding file under the naming convention *pln.radiationMode* _ *pln.machine.mat*. The number of treatment fractions is defined by *pln.numOfFractions*, while the biological model used for optimization is indicated by *pln.bioModel*. The field *pln.multScen* refers to the scenario model; in the present case, a nominal scenario is employed, as motion effects are not considered.

Meta information within *pln* is organized into dedicated substructures. The steering parameters are contained in *pln.propStf* Figure 4.4, whereas the optimization-related parameters are stored in *pln.propOpt* Figure 4.5. The field *pln.propDoseCalc* holds the configuration of the dose calculation grid.

Within *pln.propStf*, geometric properties of the treatment plan are specified, such as gantry and couch angles (following matRad's coordinate conventions), the bixel width for photons, or the lateral spot spacing for particle beams. The longitudinal spacing between spots (*pln.propStf.longitudinalSpotSpacing*) determines the number of energies required to adequately cover the target volume. The number of beam

directions is set in *pln.propStf.numOfBeams*, while the isocenter position, defined in voxel coordinates within the CT cube, can be either manually specified or automatically calculated as the center of gravity of the target volume.

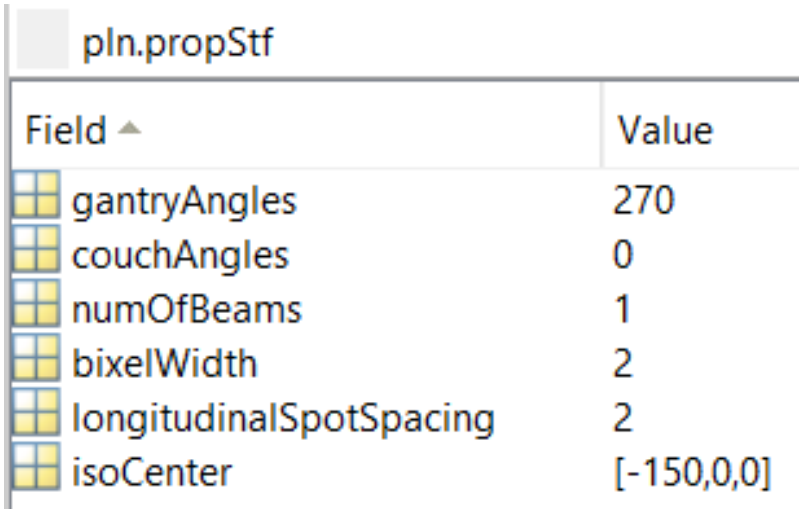

Figure 4.4: The pln.propOpt struct

The field *pln.propOpt.bioOptimization* determines the type of biological optimization to be performed. When set to *none*, the optimization is based on the physical dose; *effect* corresponds to an optimization based on biological effect, while *RBExDose* enables optimization on the RBE-weighted dose. Additional flags such as *pln.propOpt.runDAO* and *pln.propOpt.runSequencing* allow the user to activate Direct Aperture Optimization and sequencing algorithms, respectively.

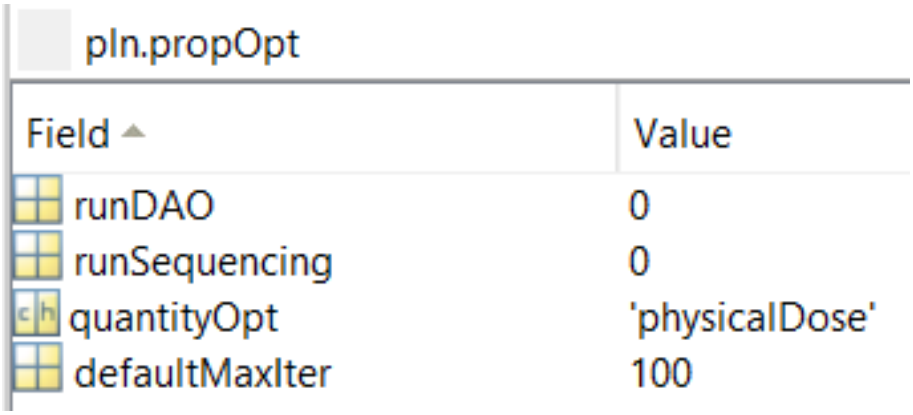

Figure 4.5: The pln.propOpt struct

The dose calculation parameters are managed in *pln.propDoseCalc*. Here, the dose

grid resolution is typically set to (3 mm, 3 mm, 3 mm), and the lateral cutoff value is defined by default as 0.995.
The optimization process can be limited by setting a maximum number of iterations through *pln.propOpt.defaultMaxIter*, which is commonly set to 500. Finally, in the case of Monte Carlo simulations, the number of simulated particles per pencil beam is specified in *pln.propMC.MCsquare_defaultHistories*, which by default is 10^6.

4.3.4 The *stf* struct

The steering information is contained in the *stf* struct, which is generated by calling the function *matRad_generateStf*.
The *stf* struct operates using both the LPS coordinate system and a beam's eye view (BEV) coordinate system, with all geometric quantities expressed in millimeters.
The irradiation geometry follows a ray–bixel (or spot) concept, schematically illustrated in Figure 4.6:

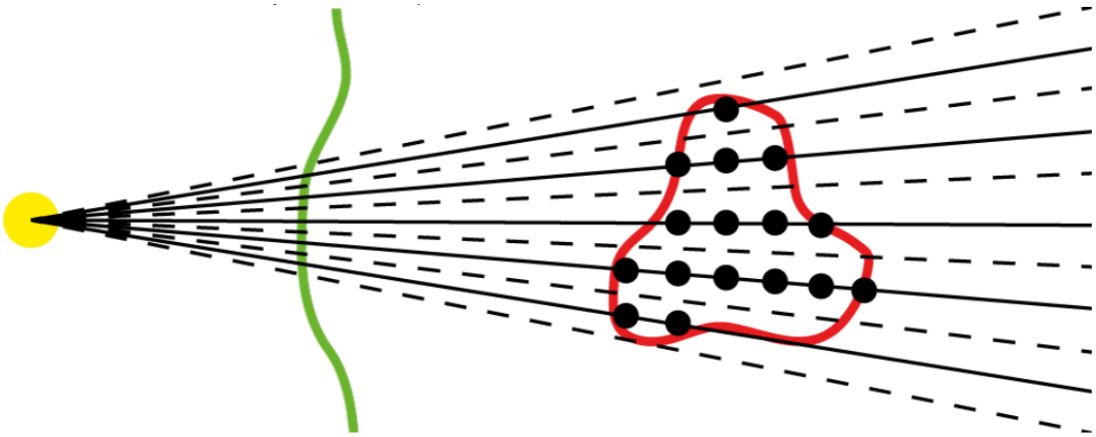

Figure 4.6: Schematic visualization of the ray and bixel concept

From the virtual radiation source (yellow), equidistant rays (black points) are projected toward the patient (green), covering the entire target volume (red). For clarity, only a 2D cross-section of the 3D ray cone is shown.
The parameter *bixelWidth* defines the lateral spacing of rays in the isocenter plane, i.e., the plane orthogonal to the beam direction and containing the isocenter. In photon therapy, each ray corresponds to a discrete fluence element (bixel), delimited by the dashed black lines. The union of all bixels ensures coverage of the target cross-section.

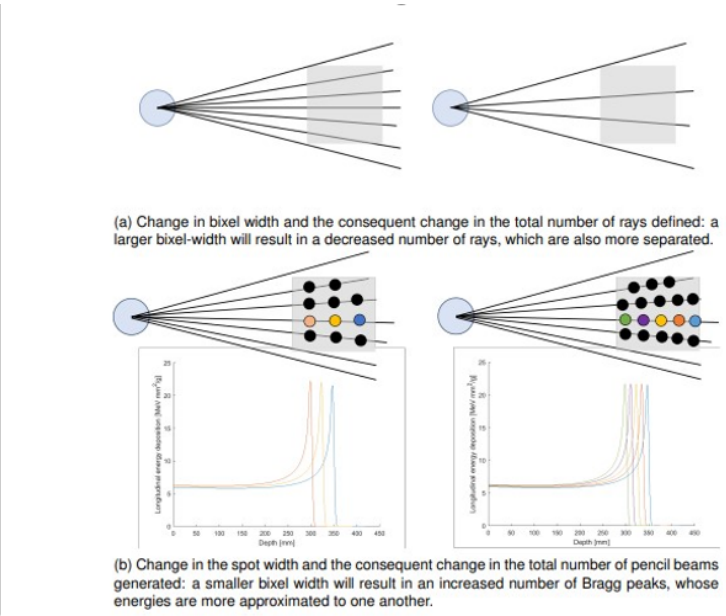

Figure 4.7: Effect of varying bixel width (a) and longitudinal spot spacing (b)

In particle therapy (IMPT), an additional degree of freedom is present: the particle range, controlled by the beam energy. During *stf* generation, the depth extent of the target along each ray is evaluated, and individual dose spots are placed accordingly. The longitudinal sampling is defined by the parameter *pln.propStf.longitudinalSpotSpacing*. A smaller spacing results in a higher number of spots (and energies) per ray, whereas a larger spacing reduces the number of required energies, as shown in Figure 4.7.

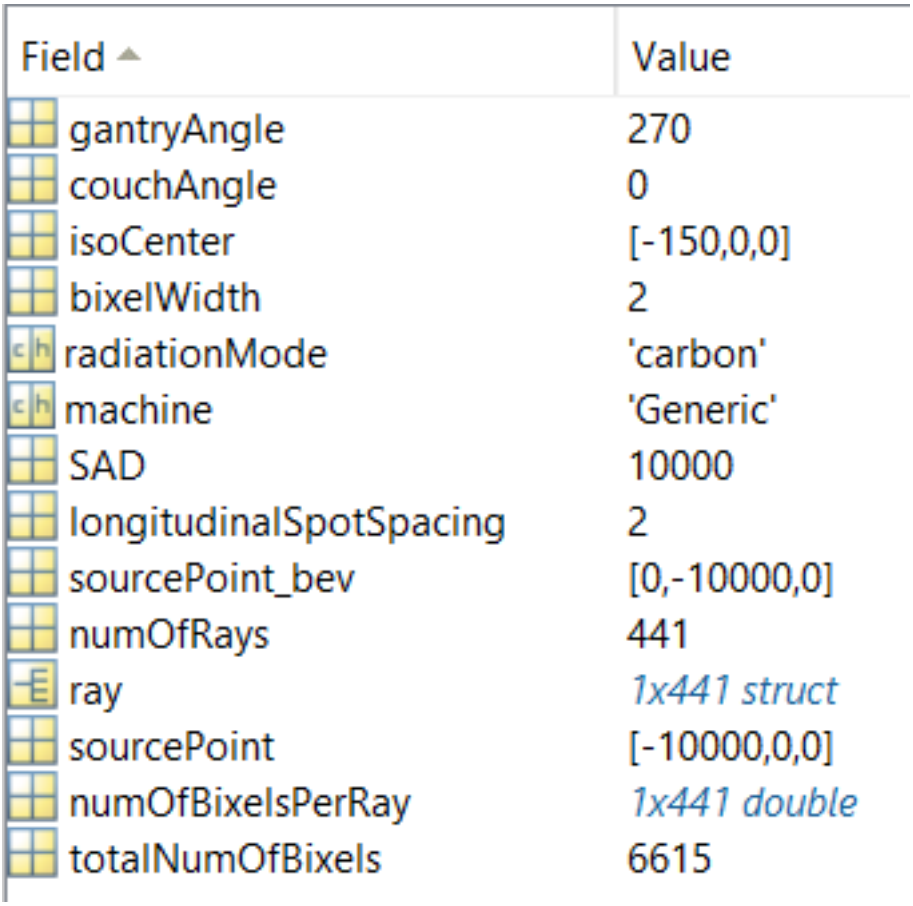

Figure 4.8: The stf struct

The *stf*, shown in Figure 4.8 structure therefore stores all geometric information needed to model the irradiation setup, for each treatment beam, its orientation (gantry and couch rotations, defined within $0° - 359°$), the LPS coordinates of the virtual source, and the spatial distribution of bixels/spots. Detailed ray information, including their origin and direction vectors, is stored in the sub-structure *stf.ray* . For photon beams, each ray contains a single bixel. In contrast, proton and carbon ion beams include multiple spots positioned along the beam axis to guarantee full volumetric coverage. The total number of bixels/spots per beam is finally recorded within the *stf* structure and is subsequently used during dose calculation and—if applicable—MLC sequencing or particle fluence optimization.

4.3.4.1 stf.ray substructure

The *stf.ray* substructure, shown in Figure 4.9, stores detailed information for each individual ray (photons) or pencil beam (particles), including spatial coordinates and energy sampling parameters used during dose computation.

stf.ray

Fields	rayPos_bev	targetPoint_bev	rayPos	targetPoint	energy	rangeShifter	focusIx	numParticlesPerMU	minMU	maxMU
1	[20,0,20]	[20,0,20]	[0,20,20]	[0,20,20]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
2	[20,0,18]	[20,0,18]	[0,20,18]	[0,20,18]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
3	[20,0,16]	[20,0,16]	[0,20,16]	[0,20,16]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
4	[20,0,14]	[20,0,14]	[0,20,14]	[0,20,14]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
5	[20,0,12]	[20,0,12]	[0,20,12]	[0,20,12]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
6	[20,0,10]	[20,0,10]	[0,20,10]	[0,20,10]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
7	[20,0,8]	[20,0,8]	[0,20,8]	[0,20,8]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
8	[20,0,6]	[20,0,6]	[0,20,6]	[0,20,6]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
9	[20,0,4]	[20,0,4]	[0,20,4]	[0,20,4]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
10	[20,0,2]	[20,0,2]	[0,20,2]	[0,20,2]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
11	[20,0,0]	[20,0,0]	[0,20,0]	[0,20,0]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double
12	[20,0,-2]	[20,0,-2]	[0,20,-2]	[0,20,-2]	1x15 double	1x15 struct	1x15 double	1x30 double	1x30 double	1x30 double

Figure 4.9: The stf.ray struct

The main fields of this structure are the following:

- **rayPos_bev**: coordinates of the ray interception point with the isocenter plane in beam's eye view (BEV) coordinates. Stored as (y, x, z), following MATLAB array index conventions.
- **targetPoint_bev**: the ray's projected target point in BEV coordinates. It defines an extended line passing from the virtual source through *rayPos_bev* and beyond the patient, enabling ray tracing during dose calculation.
- **rayPos**: same ray position but expressed in the global LPS coordinate system.
- **targetPoint**: target point converted into the LPS coordinate system for geometry consistency across the planning workflow.
- **energy**: For photon beams, a single nominal energy value corresponding to the LINAC beam nominal energy (e.g., 6 MV). For particle beams, a vector containing all individual energy layers (Bragg peak depths) of dose spots placed along the ray. Thus, the vector length equals the number of spots per ray.
- **numOfBixelPerRay**: number of fluence elements (photons) or dose spots (particles) defined along the ray. This directly matches the length of the *energy* field for particle therapy cases.

4.4 Dose calculation

The dose calculation engine in matRad is organized as a nested loop structure: the outer loop cycles through the treatment beams, while the inner loop processes the

rays associated with each beam. For charged particles, a further level of iteration accounts for the individual pencil beams distributed along each ray. For each ray, geometric distances from the virtual source and the radiological depth along the beam path are evaluated using an exact ray-tracing algorithm, which ensures an accurate modeling of tissue heterogeneities.
In inverse planning, the physical dose distribution is expressed in matrix form as:

$$d_i = \sum_j D_{ij}\, w_j \quad \Longrightarrow \quad \vec{d} = D\,\vec{w} \tag{4.1}$$

where D_{ij} quantifies the dose deposited in voxel i by fluence element j (photon bixel or ion pencil beam), and w_j denotes its fluence weight or number of primary particles. Since the matrix is typically large but sparsely populated, it is stored in MATLAB's sparse format to reduce memory usage. If no fluence optimization is required, storing the matrix can be avoided.
Photon dose computation relies on an analytical model based on the singular-value decomposition of pencil-beam kernels. For actively scanned protons and carbon ions, matRad uses a conventional pencil-beam approach, where the dose in each voxel is decomposed into a depth-dependent component and a lateral-scattering contribution. The lateral spread is modeled through a depth-dependent Gaussian profile—single Gaussian for carbon ions, double Gaussian for protons to better describe the low-dose halo.
The physical dose engine for protons and carbon ions follows the same computational structure. When biological optimization is activated for carbon ions, matRad additionally generates depth-dependent α and $\sqrt{\beta}$ matrices, which allow integrating biological effectiveness directly within the optimization process.

4.4.1 Dose influence matrix calculation

The computation of dose influence data proceeds in two main stages. First, the irradiation geometry is fully defined by generating the steering information (*stf*) corresponding to the selected beam configuration. In a second step, the geometric and radiological quantities derived from ray tracing are used to build the dose influence matrix required for inverse planning.

4.4.1.1 Dose objectives

Customized dose objectives can be specified using:

```
DoseObjectives.matRad_SquaredDeviation(penalty, dRef)
```

where *penalty* defines the optimization weight and *dRef* the prescribed dose for the target or organ at risk.

4.4.1.2 Dose influence data

Once the steering information is available, matRad evaluates the dose deposited by each fluence element in all voxels by combining ray tracing with pre-tabulated dosimetric data. For protons and carbon ions, the depth-dose component is obtained from the base datasets, while lateral scatter is modeled using the tabulated depth-dependent Gaussian width $\sigma(z)$.
The contribution delivered by fluence element j to voxel i is stored in the dose influence matrix D_{ij}, enabling efficient forward dose computation during iterative optimization, as illustrated in Fig. 4.10.

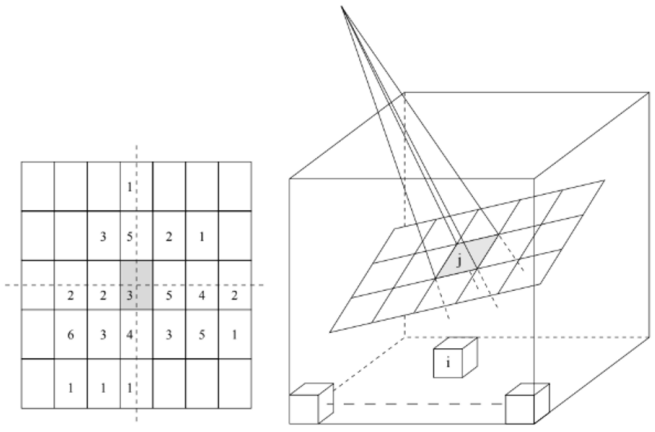

Figure 4.10: Schematic representation of the D_{ij} matrix, which maps voxel indices i and fluence elements j to the corresponding dose contributions.

For carbon ions, matRad also computes depth-dependent α and β values derived from the Local Effect Model IV (LEM IV), enabling the determination of RBE-weighted quantities throughout the volume. All physical and radiobiological data used in this process are contained in the `protons_Generic` and `carbon_Generic` base files, which include integrated depth-dose curves, lateral beam widths and (for carbon)

the α and β tables. Depth is expressed in millimeters, α in [1/Gy], β in [1/Gy^2], and the integrated depth-dose in MeV cm^2/(g $\cdot$ primary). During dose calculation, the fluence weights correspond to the number of primary particles (in units of 10^6), and the resulting physical dose is reported in gray (Gy).

4.5 Fluence optimization

In matRad, fluence optimization constitutes the core of the inverse treatment planning process. It aims to determine the optimal fluence map — or equivalently, the optimal particle spot weights — that produce a dose distribution matching the prescribed clinical objectives while satisfying all physical and biological constraints. This optimization is formulated as a constrained minimization problem, where the difference between the calculated dose and the prescribed dose is minimized according to user-defined penalty functions.
The optimization problem can be expressed as:

$$\min_{\vec{w}} \; f(\vec{w}) = \sum_{s=1}^{N_S} \sum_{i \in s} p_s \; (d_i(\vec{w}) - d_{s,\text{ref}})^2 \tag{4.2}$$

where $\vec{w}$ represents the fluence or spot weights, p_s is the penalty factor assigned to the structure s, $d_i(\vec{w})$ denotes the dose in voxel i as a function of the weights, and $d_{s,\text{ref}}$ is the prescribed reference dose for that structure. The objective function thus penalizes deviations from the clinical dose goals within target volumes and organs at risk.
The dose d_i for each voxel is evaluated using the dose influence matrix formalism introduced previously, such that the optimization problem remains linear with respect to the fluence weights. For photon plans, the fluence variables correspond to the individual bixels composing each beam. In the case of charged particle therapy, they represent the intensities of the individual pencil beams.
MatRad provides several optimization algorithms that can be selected depending on the desired trade-off between computational speed and convergence accuracy. The default optimization method is a gradient-based iterative solver, which exploits the sparse structure of the dose influence matrix to efficiently compute the gradient of the objective function with respect to the fluence weights. Optional regularization terms can be introduced to enforce smoothness in the fluence distribution and avoid unphysical intensity modulation.
When dealing with biological optimization — particularly in carbon ion therapy

— the optimization process is performed on the biologically effective dose, using the α and β matrices derived from the Local Effect Model IV. In this case, the objective function accounts for the radiobiological effectiveness (RBE) of each voxel, making the relationship between physical dose and biological response inherently nonlinear. The solver iteratively recalculates the RBE-weighted dose distribution until convergence is achieved.

All optimization parameters, including penalty weights, prescribed doses, and structure-specific constraints, are defined within the *pln* structure under the field *pln.propOpt*. This allows a flexible and modular configuration of the optimization objectives, enabling simultaneous control of target coverage, dose homogeneity, and organ-at-risk sparing. The optimization results are stored in the MATLAB workspace as the variable *resultGUI*, which can be visualized and analyzed through the graphical user interface.

The fluence optimization step concludes the computational core of the inverse planning workflow in matRad. The resulting fluence maps or spot intensities are subsequently used for dose recalculation, plan evaluation, and, in the case of photon treatments, multileaf collimator (MLC) sequencing to derive a clinically deliverable treatment plan.

4.6 Graphical User Interface - GUI

matRad provides a graphical user interface (GUI) that allows users to visualize and explore computed tomography (CT) data together with superimposed transparent dose distributions, as illustrated in Fig. 4.11.

In addition to visualization, the GUI enables the user to execute the complete sequential treatment planning workflow, from patient data loading to optimization and plan evaluation. All GUI data elements are automatically synchronized with MATLAB's base workspace, ensuring consistency between graphical and command-line operations.

This bidirectional synchronization provides a flexible environment for algorithm development and testing, allowing users to seamlessly switch between the command line and the GUI while maintaining full access to the same data structures and computational results.

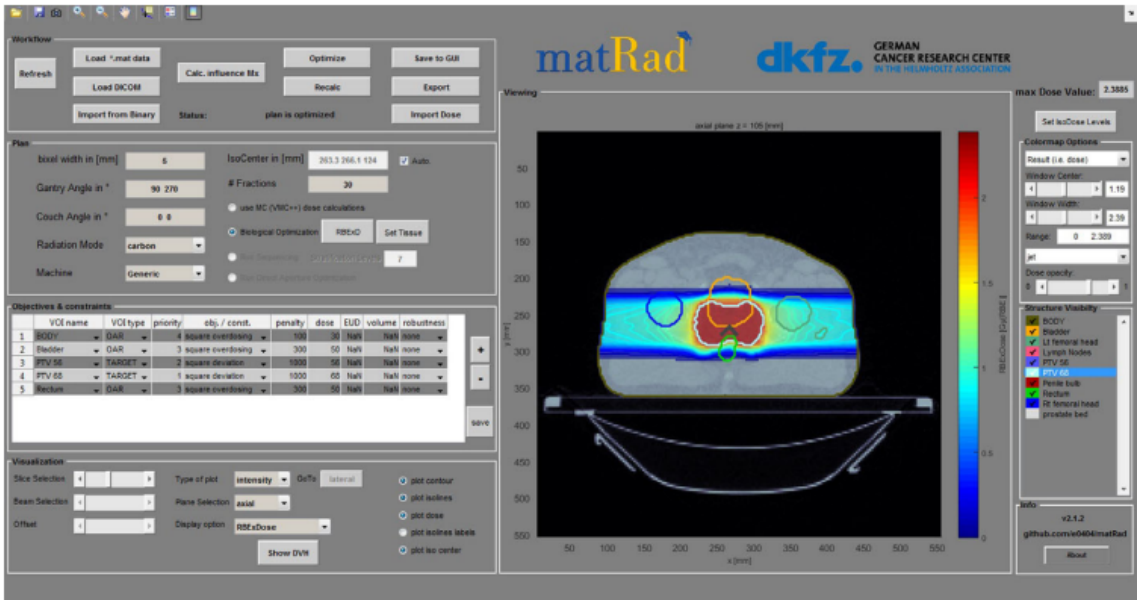

Figure 4.11: The graphical user interface of matRad is grouped into several functional blocks, whereas each block addresses certain aspects of treatment planning, and also presents a visualization tool where both anatomical structures and dose distributions are dsyaplaied

4.7 Comparison of dose calculation

To assess the accuracy of matRad's dose computation, fluence optimization and dose calculation were first performed with a clinically approved reference system, and the resulting treatment plans were subsequently recalculated in matRad.
The comparison shows a very good agreement between matRad and the clinical platform, with global γ-analysis pass rates of $\geq 99.67\%$ across all radiation modalities. Moreover, dose calculation and optimization for realistic patient scenarios can be completed within a few minutes, demonstrating both the reliability and the computational efficiency of the framework.[16][17]

Chapter 5

Code analysis and treatment plan definition

5.1 Rationale

In this chapter, the code developed within the framework of this thesis will be analyzed in detail. As previously stated, the main goal is to develop a calculation tool capable of performing dose calculations with protons and carbon ions on a defined target, with the aim of employing it in research applications for treatment plan generation in irradiation experiments.
The experimental context requires a high degree of flexibility in the definition of both geometry and beam parameters. Furthermore, there is the necessity to rely on a tool that allows independence from the clinical treatment planning systems currently used at CNAO.
This work is embedded in the broader framework of research activities carried out at CNAO, with the objective of developing a calculation chain consistent with the beam models currently implemented at the center and experimentally validated. In this context, particular attention has been devoted to the construction of the machine files, to the analysis of calculation times, and to the resolution of the various problems encountered throughout the entire development process.
Finally, a detailed description of the implemented code and its workflow is provided, followed by a benchmark comparison between the obtained results and those calculated using the conventional treatment planning system adopted at CNAO (RayStation).

5.2 Code description

5.2.1 The Geometry definition

As discussed in the previous chapters, matRad provides various phantom samples for dose calculation, mainly designed for clinical simulations or optimization algorithm tests. However, since our goal is to generate treatment plans for different irradiation setup, such phantoms are not adequate to describe in a realistic way our experimental apparatus.

For this reason, it is necessary to define a personalized phantom, capable of reproducing a homogeneous environment compatible with the typical conditions of in vitro irradiation experiments or other situations.

The base of our model is made by a parallelepiped-shaped water phantom with dimensions [300, 300, 100] mm in LPS coordinates (x, y, and z, respectively). This volume represents the medium surrounding our flask and it is referred to as *waterBox*. Inside this structure, a second smaller water phantom is inserted, also having a parallelepiped shape with variable dimensions according to the case of study. As a reference, the dimensions are [60, 40, 40] mm, but they can be easily varied according to the necessity.

These structures are designed respectively as OAR (Organ At Risk) and target, even if in this context the denomination OAR for the waterBox has a purely formal function inside the data structure of matRad. The dose objective is imposed exclusively on the target volume, that represents the volume of interest for cellular irradiation experiments.

When a structure is defined in matRad, in addition to the dimensions it is necessary to specify its position inside the CT structure by an *offset* parameter, that describes the translation of the geometrical center of each structure with respect to the center of the CT volume itself. Furthermore, for each volume it is necessary to define the material of which it is made, by assigning a Hounsfield Unit value (HU). In our case, since both the Water Box and the target are made of water, the HU value is set to zero.

The dose objective is another parameter to be linked to the structures of interest, in our case to the target, and can be defined as follows:

```
objective = struct(DoseObjectives.matRad_SquaredDeviation
   (100, 2));
```

Here, 100 represents the penalty value, hence the number of voxels that does not have to satisfy the objective, while 2 is the prescribed dose in Gy. However, the dose objective can be also defined together with the *cst* structure along with different dose constraints for the various structures.
Together with the CT, its resolution must also be specified; in our case, it is set to [1, 1, 1] mm, indicating the voxel size of the CT grid.
Once the dimensions, the offset and the structure materials are defined, the CT and the structure containing all the information on volumes are generated using the function:

```
[ct,cst] = builder.getctcst();
```

This passage concludes the geometry definition and gives the necessary elements for the treatment plan configuration.

5.2.2 Treatment plan parameters definition

Once the geometry has been defined, the treatment plan parameters must be specified. This phase includes the choice of the type of radiation to be used, the corresponding machine file, and the other geometrical and physical parameters of the beam. In this work, both proton and carbon ion beams have been used, employing machine files specifically reconstructed to reproduce the beam characteristics currently used at CNAO, it will be made explicit later in section 5.3.
Among the other parameters, the gantry angle should be highlighted; in our case, it is set to 270° in order to obtain a beam entering from the left along the X axis in the LPS coordinate system.
A further parameter is the *isocenter*, namely the reference point around which all components of the radiotherapy system rotate. In the experimental context of this work, the isocenter is positioned on the proximal face of the *waterBox*.
Other fundamental aspects are the longitudinal and lateral spot spacing, which directly influences both the spatial resolution of the dose distribution and the total number of spots used in the calculation. To maintain coherence with the configuration adopted at CNAO, for protons the longitudinal spot spacing has been fixed to 4 mm, while the lateral spacing, hence the bixel width, is set to 3 mm. For carbon ions, both parameters have been set to 2 mm, reflecting the higher lateral precision typical of this irradiation technique.

A further important parameter is the choice of the quantity to be optimized. In this work, it has been decided to use the physical dose, since the objective is to study the energy distribution deposited inside the target volume. Although matRad also allows the optimization of radiobiological quantities, such aspects are outside the scope of this work.
An additional relevant parameter, the *dose grid resolution* defines the resolution of the grid used for dose calculation. This does not necessarily need to coincide with the CT resolution, as matRad subsequently interpolates the dose distribution onto the CT grid. As will be discussed later, the choice of this parameter can significantly impact the overall calculation time.
Once all the parameters have been defined, the beam geometry and the structure containing all the steering information (*stf*) are generated using the following function:

```
stf = matRad_generateStf(ct,cst,pln);
```

5.2.3 Dose Calculation and Optimization

Once the geometry definition of the problem has been completed and all the treatment plan parameters have been defined, it is possible to proceed with the calculation of the dose distribution inside the considered volume. Using matRad, this step is performed through the function:

```
matRad_calcDoseInfluence(ct, cst, pln, stf)
```

This function represents one of the fundamental passages of the entire calculation chain. It has the role of determining how each element of the beam, hence each *bixel*, contributes to the dose deposition over each voxel composing the CT.
MatRad makes available different dose calculation engines, which differ in terms of physical accuracy and computational complexity. In particular, the software allows the use of Monte Carlo engines, such as MCsquare, and analytical engines based on pencil beam models, such as HongPB. Monte Carlo engines provide an extremely accurate description of particle transport and physical interactions in the medium, and are particularly indicated in the presence of complex heterogeneities or geometries. However, they require significantly longer calculation times and intensive use of computational resources.
In the context of this work, which is focused on the study of *in vitro* irradiation over

homogeneous and regular geometries, it has been decided to use an analytical pencil beam model, according to what has been discussed in the paragraph *Dose calculation* in Chapter 4. This choice could represents a reasonable compromise between physical accuracy and computational efficiency. In a homogeneous water phantom, such as the one used to simulate cellular flasks, pencil beam models provide a reliable description of the dose distribution, while allowing a significant reduction of calculation time.
The pencil beam model used in matRad describes the beam as a superposition of different elementary contributions, each characterized by a longitudinal and a lateral dose distribution. The longitudinal component is described through depth–dose curves, which reproduce the Bragg peak, while the lateral component is modelled using one or two Gaussian distributions. The parameters of these distributions depend on the beam energy and are contained in the machine file. The combination of these components allows the reconstruction of the dose distribution deposited inside the three-dimensional volume.
The main result of the calculation is the dose influence matrix, referred to as d_{ij}, where each element represents the dose deposited in the i-th voxel following the activation of the j-th bixel. In other words, the dose influence matrix encodes, in a discrete way, the linear relationship between the bixel weights and the resulting dose distribution inside the irradiated volume. This matrix represents the core of the treatment planning problem and is the key element on which the subsequent optimization is based.
When required, the dose influence matrix is used to perform the fluence optimization, which matRad carries out through the function:

```
matRad_fluenceOptimization(ct, cst, pln)
```

The aim of the optimization is to determine the optimal bixel weights in order to satisfy the dose objectives defined by the user. These objectives are expressed as objective functions and constraints associated with the different structures, such as the target and possible OARs. In the case of this work, the optimization is focused on the target volume and employs a squared dose deviation objective, which penalizes the quadratic deviation of the dose from the prescribed value.
From a mathematical point of view, this problem is formulated as a convex optimization problem, in which the cost function depends on the dose influence matrix and on the bixel weights. matRad solves this problem through iterative algorithms, progressively updating the bixel weights until the desired dose distribution is reached

within a given tolerance with respect to the imposed objectives. The linearity of the dose model, typical of pencil beam approaches, allows the computational problem to remain manageable even in the presence of a large number of bixels.
The results of the dose calculation and fluence optimization are finally stored within the matRad GUI. They can be analyzed and visualized through the graphical interface of the software, which allows the inspection of both the geometry and the dose distribution. This phase is particularly useful for verifying the quality of the treatment plan and for comparing the obtained results with those produced by other optimization tools.

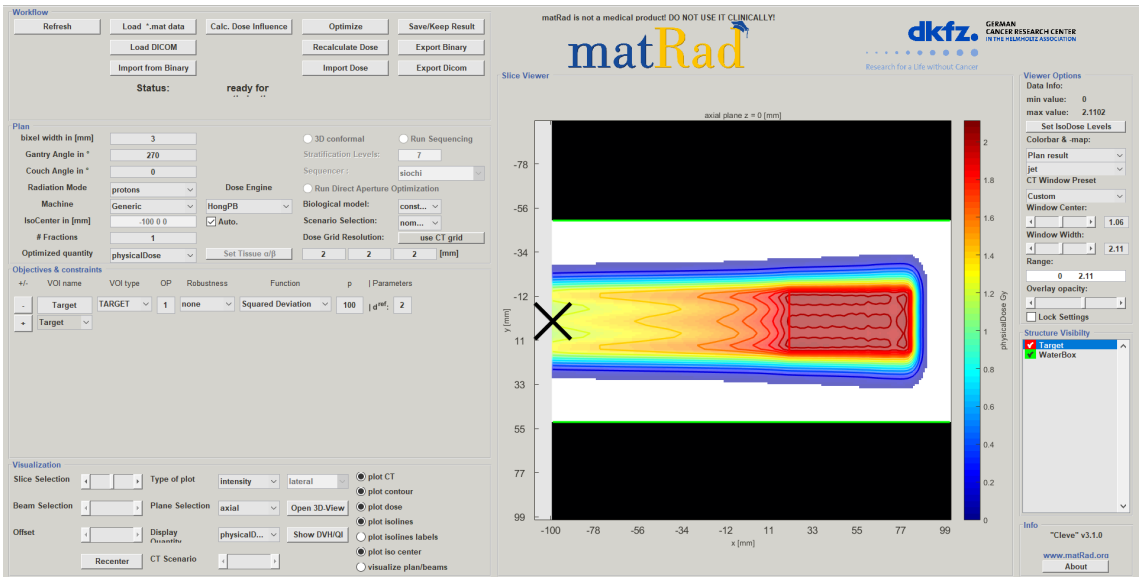

Figure 5.1: Visualization of the GUI when irradiating with protons a [60,30,30] water target (red) and a [200, 100, 100] waterBox (green) with a dose objective of 2 Gy

5.3 Machine file construction

A crucial part of this work was the construction of the machine files, in order to substitute the beam data already present in matRad with the description of CNAO beams.
The data related to the beam are defined inside files called machine files; they contain all the necessary information to describe the physical characteristics of the beam, such as the available energies, the position of the Bragg peak, the depth–dose curves, and the modeling of the lateral beam widening. As already mentioned, these data are already present as default in matRad and are essential to perform a correct dose simulation.
To develop a code able to perform calculations coherent with those carried out at CNAO it was therefore necessary to rewrite these machine files. Two different files have been created, one for protons and one for carbon ions, containing the specific characteristics of the beams used at CNAO. The starting data are obtained from previous experimental measurements and from Monte Carlo simulations performed with the code FROG, whose beam models have been experimentally verified and adopted at the center.
Through a dedicated script, it was possible to build machine files in a coherent way with respect to the structure already present in matRad. The script constructs a structure composed of different fields, each of them filled appropriately on an energy-by-energy basis. During this process, several precautions were required. First of all, it was necessary to subtract 3.1 mm from the parameter called *depth*, which contains all the position describing the Bragg peak and, together with the Z parameter, describes the depth–dose curves for each energy level, some examples are shown in Fig. 5.2. This modification was necessary because the measurements were performed using a ripple filter, a beam-modifying device designed to increase the energy spread by introducing controlled energy modulations, thereby broadening the Bragg peak and achieving a more uniform dose distribution along the penetration depth.

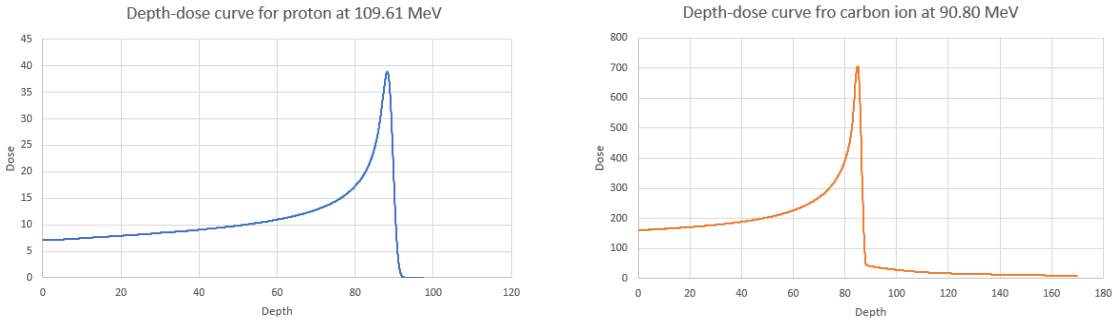

Figure 5.2: Example of depth-dose curves for both protons (blue) and Carbon ions (orange) at different energies

Another fundamental element in the machine file definition is the modeling of the lateral beam widening. The model used in the FROG data is a triple Gaussian model, which allows an accurate representation of the transversal dose distribution. In this model, the beam is described by three different Gaussian components: a central component (core), describing the main part of the beam that dominates the high-dose region; an intermediate component (halo), mainly associated with multiple Coulomb scattering; and a more external component (tail), which takes into account the effects due to nuclear interactions and secondary particles. This description is particularly suitable for Monte Carlo simulations, but it implies a significant increase in computational complexity.

On the contrary, matRad uses a simplified lateral model, based on one or at most two Gaussian functions, depending on the particle type. In particular, for protons a double Gaussian model is used, while for carbon ions a single Gaussian model is generally adopted. This choice reflects a compromise between physical accuracy and computational efficiency, consistent with the objective of matRad to provide a flexible and relatively fast tool for treatment plan generation.[15, 18, 19]

As a consequence, in order to use the beam data from FROG within matRad, it was necessary to convert the triple Gaussian lateral model into a single or double Gaussian representation. Such a conversion requires a recalibration of the beam parameters in order to preserve, as much as possible, the characteristics of the original beam distribution.

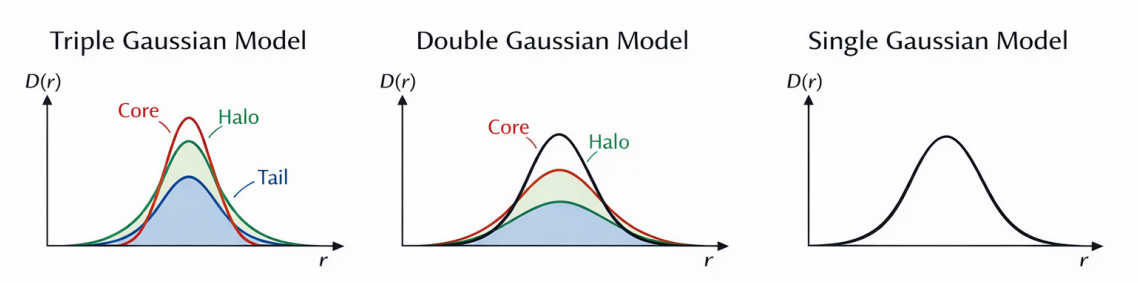

Figure 5.3: Schematic representation of the difference between single, double and triple Gaussian Model

From a mathematical point of view, the equivalent width σ_{eq} is calculated by combining the three components of the initial model through a weighted mean square, in which the weights reflect the relative contribution of each component to the total dose. In this way, the equivalent Gaussian remains consistent with the original model and can be used for the conversion to the single or double Gaussian representation. This procedure has been implemented in the script used to generate the machine files and from the following formula the new weights have been found:

$$\sigma_{\text{eq}} = \sqrt{\frac{\sigma_1^2 w_1 + \sigma_2^2 w_2 + \sigma_3^2 w_3}{w_1 + w_2 + w_3}} \tag{5.1}$$

In matRad, the lateral kernel—namely the function describing the lateral diffusion of the dose—is defined through the Gaussian parameters contained in the machine file. This kernel is used by the dose calculation engine, in combination with the depth–dose curves, to model the transverse spread of the dose within each voxel. The conversion of the lateral beam models is performed at the machine file level, rather than during the dose calculation itself, ensuring that the entire calculation chain operates in a consistent and well-controlled manner.

This process of adaptation of the beam model represents a crucial step of this thesis work, as it allows the integration of high-quality beam data derived from experimentally validated Monte Carlo simulations into a lighter and faster calculation framework such as matRad. At the same time, it highlights how the choice of the lateral beam model influences not only the dosimetric accuracy, but also the calculation time and the numerical stability of the dose influence matrix, aspects that will be discussed in detail in the following sections.

What has been described so far refers to the general workflow of matRad for dose calculation and optimization. However, in order to adapt the framework to the specific requirements of this work, it was necessary to develop and implement several

custom codes, each designed with distinct features and objectives. These additional scripts extend the standard matRad functionalities, allowing greater flexibility and control over the optimization process and enabling the achievement of more accurate and reliable results tailored to the problem under investigation.

5.4 Written codes

During this work, five distinct MATLAB codes were developed, each designed for a specific dose calculation modality. The codes, named *layerOptimization.m*, *ForwardDoseCalc.m*, *stfNew.m*, *stfNewNoEnergy.m*, and *stfNewForwardDoseCalc.m*, constitute a flexible suite that enables the exploration of different dose calculation approaches while maintaining strict consistency with the beam models used at CNAO. Taken together, these tools provide a comprehensive environment for the simulation and optimization of dose distributions in proton and carbon ion therapy. They allow both forward dose calculation and layer-by-layer optimization, while offering a high degree of control over beam geometry and bixel weights.

5.4.1 ForwardDoseCalc.m

The code *ForwardDoseCalc.m* was developed to perform forward dose calculations without any optimization procedure. In this approach, bixel weights are manually assigned by the user, allowing direct evaluation of the dosimetric system's response for a given beam configuration.
In this work, weights were selected to produce a SOBP with a selected depth and thickness in water, simulating a uniform dose distribution within the target volume. This type of calculation is particularly useful for verifying the correctness of both the machine files and the dose kernels.
Moreover, *ForwardDoseCalc.m* allows the analysis of different weight choices and enables the exact reproduction of measurements performed at CNAO, without interference from optimization algorithms, which could otherwise obscure potential discrepancies. The simulated plans are used for cellular irradiations; this justifies the fact that the curves are irregular and not perfectly flat.
Figure 5.4 shows the results of calculations performed with the same configuration for both protons and carbon ions.

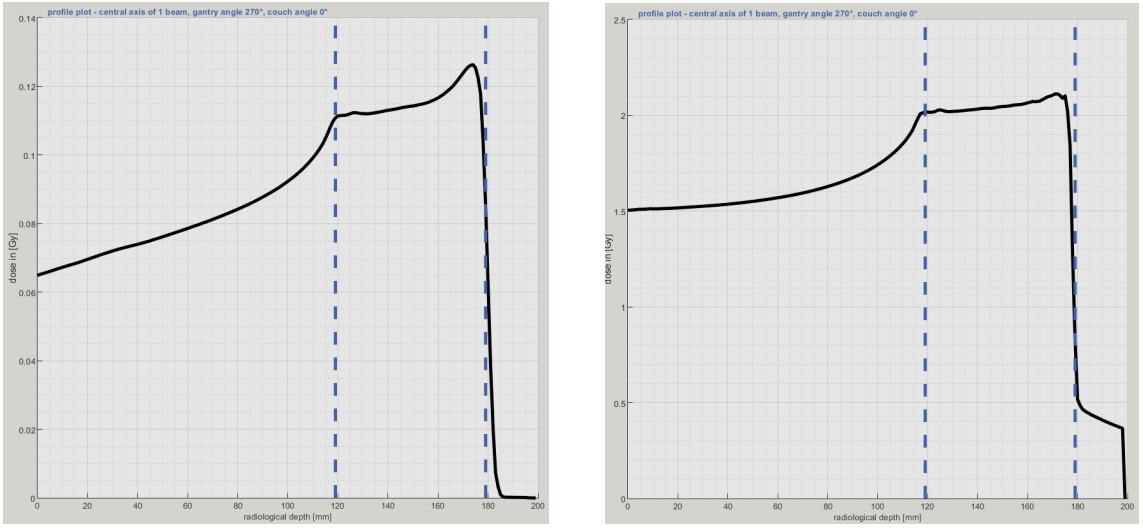

Figure 5.4: SOBP between 120 and 180 mm, obtained with the dose calculation over a target of dimension [60,40,40] mm, using *ForwardDoseCalc.m* with protons (left) and carbon ions (right).

5.4.2 layerOptimization.m

By default, matRad performs a spot-by-spot optimization, assigning a different weight to each bixel. However, to maintain consistency with CNAO's irradiation methods, where a SOBP is generated by assigning the same weight to all spots within a slice, i.e. bixels with the same energy, a specialized approach is required.
The code *layerOptimization.m* was developed to satisfy this requirement. Initially, it calculates the dose using matRad's pencil beam engine, based on the kernels described in previous chapters. From the resulting dose influence matrix, the code performs an optimization under the constraint that all bixels at a given energy—i.e., all bixels within the same energy layer—share the same weight.
This approach significantly reduces the dimensionality of the optimization problem: rather than optimizing each individual bixel, each energy layer is treated as a single optimization variable. Physically, this corresponds to considering each energy layer responsible for dose deposition in a specific depth region, essentially forming a transversal slice of the irradiated volume.
Layer-by-layer optimization is particularly suited to this context, allowing controlled modulation of the dose along the depth direction while maintaining high lateral dose uniformity within each layer. Although less flexible than fully voxel-based optimization, this method offers increased numerical robustness and substantially

reduced computation time.
The results of the dose calculation and optimization are shown in Fig. 5.5.

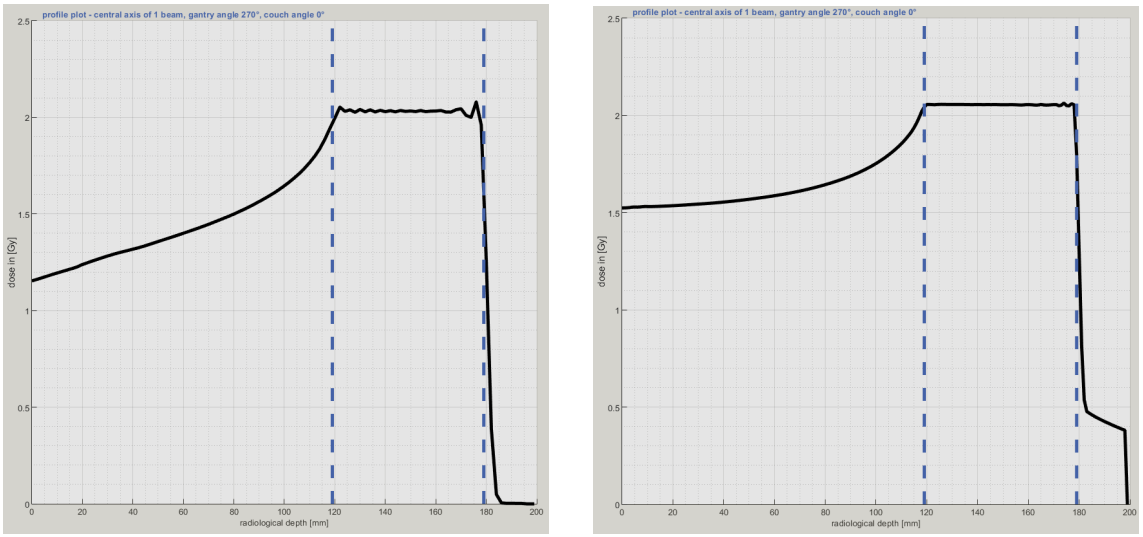

Figure 5.5: SOBP between 120 and 180 mm, obtained with the dose calculation and optimization over a target of dimension [60,40,40] mm using *layerOptimization.m* with protons (left) and carbon ions (right).

5.4.3 stfNew.m

The code *stfNew.m* allows the modification and customized construction of the STF structure in matRad, which defines the spatial arrangement of spots and their associated energies. Specifically, the code permits manual assignment of lateral spot coordinates in the y and z directions, while the x coordinate is fixed at either 0 or 1000, according to the chosen coordinate system. The interaction depth is determined by the specified energies, respecting the weight constraints imposed by *layerOptimization.m*.
This method was used to design an STF capable of generating a cubic irradiated volume of variable dimensions, with reference dimensions $60 \times 40 \times 40$ mm^3, and with an energy range tuned to achieve a SOBP between 120 and 180 mm in water, adjustable as needed. *stfNew.m* thus provides direct control over beam geometry and target coverage, overcoming matRad's limitation where for some geometries do not assign the same number of energies to each voxel.
Direct manipulation of the STF structure is crucial for advanced beam modeling, allowing more faithful reproduction of what occurs in clinical systems and enabling the study of sophisticated treatment scenarios.

A variant, *NewStfNoEnergy.m*, performs the same calculations as *stfNew.m*, manually setting spot coordinates but leaving energy selection to matRad based on the machine file. This variant has been observed to be more precise and efficient during optimization. Figures 5.6 and 5.7 illustrate dose calculations and optimizations for both protons and carbon ions using the two codes. The SOBP is visibly flatter and more uniform with *NewStfNoEnergy.m*, particularly for carbon ions.

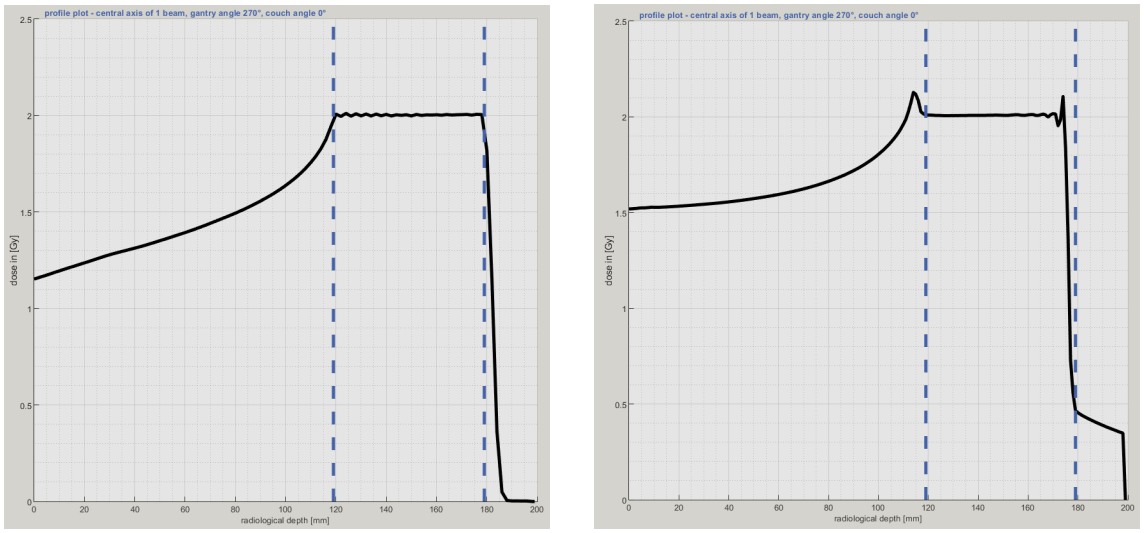

Figure 5.6: SOBP between 120 and 180 mm obtained with *NewStf.m* over a target of dimension [60,40,40] mm using protons (left) and carbon ions (right).

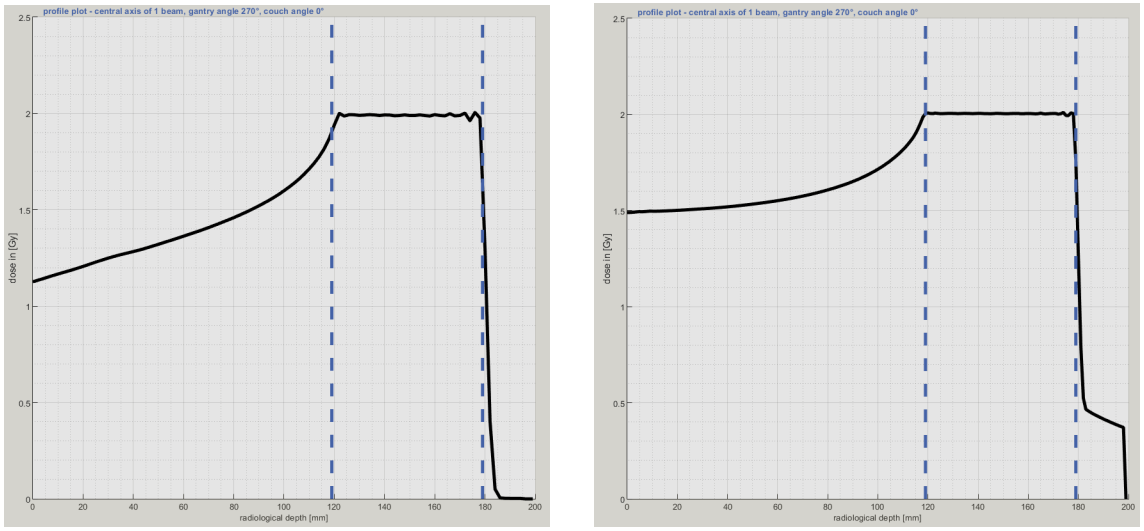

Figure 5.7: SOBP between 120 and 180 mm obtained with *NewStfNoEnergy.m* over a target of dimension [60,40,40] mm using protons (left) and carbon ions (right).

5.4.4 stfNewForwardDoseCalc.m

The code *stfNewForwardDoseCalc.m* combines the functionalities of the previous codes, integrating personalized STF creation with forward dose calculation and manual bixel weight assignment. This configuration represents the most flexible tool in the suite, allowing simultaneous control over spot positions, energies, and weights. This setup enables detailed studies of how local variations in spot positions and bixel weights influence the resulting dose distribution, which is particularly valuable when reproducing experimental measurements or comparing results with professional TPS optimizations.
Figure 5.8 shows the calculation results for protons and carbon ions with the same configuration.

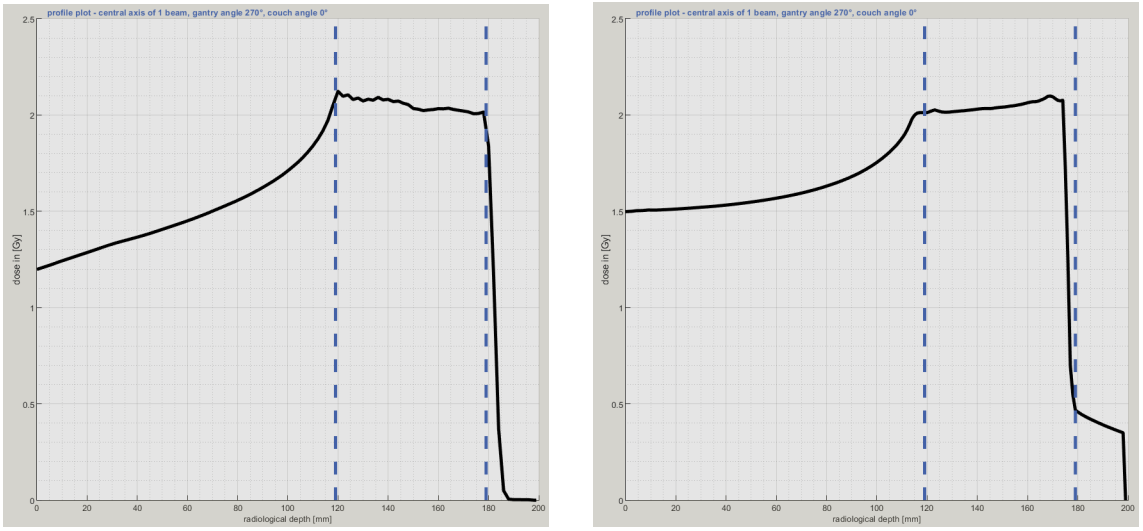

Figure 5.8: SOBP between 120 and 180 mm obtained with *stfNewForwardDoseCalc.m* over a target of dimension [60,40,40] mm using protons (left) and carbon ions (right).

Overall, these codes have proven effective for simulating and optimizing dose distributions in proton and carbon ion therapy within the matRad framework. They allow both forward dose evaluation—analyzing system response with predefined weights—and layer-by-layer optimization consistent with beam physics.
Control over spot coordinates and bixel weights is essential for accurate beam modeling and for controlled SOBP generation, in accordance with CNAO procedures. Integrating these tools with realistic machine files makes matRad not only an educational platform but also a reliable research environment for investigating beam modeling strategies and studying the effects of various physical and geometrical

parameters on the resulting dose distributions.
Not all codes achieve perfect uniformity in the SOBP. Among them, *NewStfNoEnergy.m* stands out as the most accurate, producing a flat and uniform SOBP at the desired dose for both particles. Consequently, this code serves as a reference for comparison with professional TPS results, and all subsequent analyses are based on its performance.

5.5 Observations

This section presents the results obtained using the codes described above. Special attention was paid to dose uniformity, analyzing lateral and longitudinal dose profiles, which ideally should be flat and match the prescribed dose to ensure homogeneous target irradiation and precise SOBP control.
Computation time for each code is also discussed, as it represents a critical practical aspect of dose calculations.

5.5.1 Calculation Time

One of the main challenges encountered during this work was the computational time required for dose calculation and optimization. Calculations using machine files derived from the CNAO database required significantly longer times than those using matRad's default machine files.
Form a theoretical point of view, it was expected that carbon ion calculations would inherently take longer than proton calculations due to the higher physical complexity of carbon beams. Carbon ions, being heavier and having a higher atomic number, are more prone to nuclear fragmentation in the target, generating secondary particles that contribute to the dose. This results in more complex depth–dose distributions requiring detailed physical modeling.
Additionally, carbon ion lateral spread modeling involves wider and more complex dose kernels, increasing the number of voxels contributing to each bixel's dose. The resulting dose influence matrices are larger and denser. Constructing a SOBP for carbon ions also typically requires more energy layers, as individual Bragg peaks are narrower than those of protons. These factors collectively increase computation time and memory usage. Other influences include dose grid resolution and lateral beam width (σ) values specified in the machine files.
Initially, both proton and carbon ion calculations were extremely time-consuming,

especially for irradiation fields corresponding to the simulated experimental setups. A first mitigation involved reducing the dose grid resolution from 1 mm to 2 mm for protons and 1.5 mm for carbon ions, reducing computation time by roughly a factor of five. Nonetheless, some carbon ion calculations still exceeded one hour.
Using MATLAB's parallelization toolbox to distribute computations across multiple cores offered minimal improvement.
A critical enhancement was achieved by inspecting the CNAO-derived machine files. Lateral σ values for carbon ions were an order of magnitude larger than matRad's default, due to a unit mismatch. This modifications were applied also to the proton file. Correcting these values reduced computation time from over 1000 s to approximately 150 s, highlighting the sensitivity of computational performance to correct beam parameterization and underlines the importance of a careful verification of machine file inputs when performing dose calculations.

5.5.2 Dose uniformity

A critical aspect emerging from the analysis of the results concerns the dose distribution uniformity, both laterally and longitudinally. This issue is more evident in the case of carbon ions, but it is also present, although with different characteristics, in calculations performed with protons.

5.5.2.1 Lateral dose profile

Regarding the lateral dose profile, it can be observed that the region in which the dose distribution is effectively uniform and flat is systematically reduced with respect to the region geometrically delimited by the target defined in matRad. In other words, the region of dose uniformity does not coincide with the entire extension of the target volume, but it appears narrower, making it necessary to introduce lateral margins in order to guarantee an acceptable dose distribution in the region to be irradiated, as shown in Fig. 5.9.

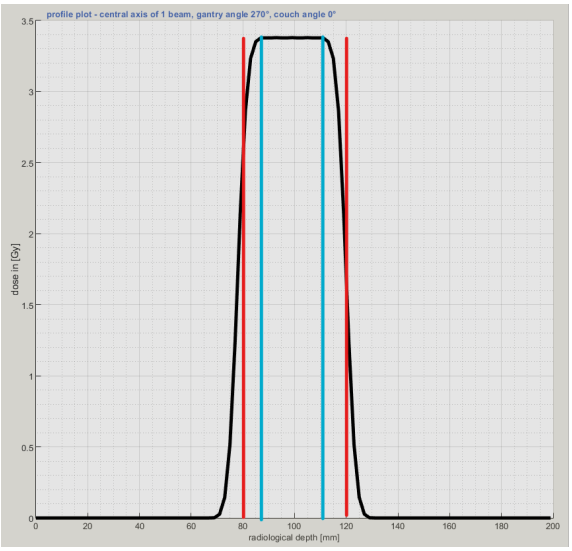

Figure 5.9: Lateral profile of a SOBP obtained with carbon ions over a target of dimensions [60, 40, 40] mm, with lateral margins indicated.

From the analysis of the transversal dose profiles, it was possible to estimate margins of approximately 1 cm on each side. As a consequence, if it is necessary to irradiate a region 20 cm wide with a uniform dose (shown in green in Fig. 5.10), the target used for the optimization should be defined with a total width of 40 cm (shown in red in Fig. 5.10). This approach allows compensation for the effects of lateral penumbra and lateral beam diffusion, ensuring that the central region of the volume receives a sufficiently flat dose distribution.

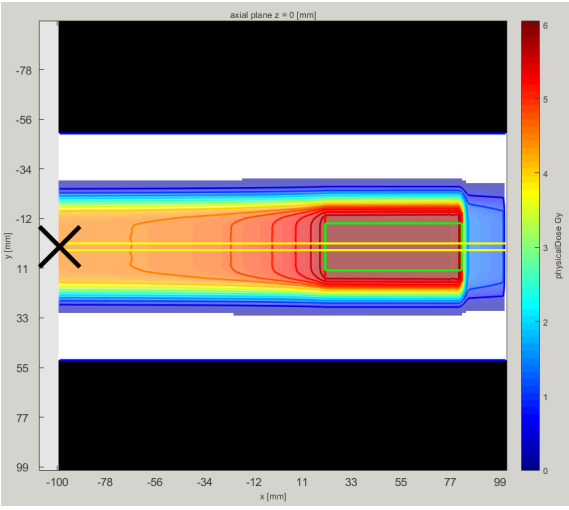

Figure 5.10: SOBP with carbon ions over a target of [60, 40, 40] mm. Green: region of dose uniformity. Red: optimization target region.

5.5.2.2 Longitudinal dose profile

The introduction of lateral margins effectively solves the problem of transversal dose non-uniformity; however, it is not sufficient to guarantee a uniform dose distribution along the longitudinal direction. By analyzing the depth–dose profiles of the SOBP, several anomalous peaks can be observed, both just before the target and inside the SOBP region itself, as shown in Fig. 5.6. This behavior has been observed in calculations performed with protons and, in a more pronounced way, with carbon ions, indicating a critical issue in the construction of the SOBP and in the choice of the energies and the weights associated with the spots.
As already mentioned, a significant improvement in the longitudinal dose distribution has been observed when using *NewStfNoEnergy.m* instead of *NewStf.m*, as shown in Fig. 5.7. In this case, the energies are not manually selected by the user but are determined by matRad according to the information contained in the machine file and in the STF structure. This approach allows for a more appropriate energy discretizations and a better superposition of the individual Bragg peaks, thereby reducing oscillations and undesired peaks inside the SOBP. However, despite the improvement in the dose profile shape, an overestimation of the prescribed dose value is still observed, as displayed in Fig. 5.11.

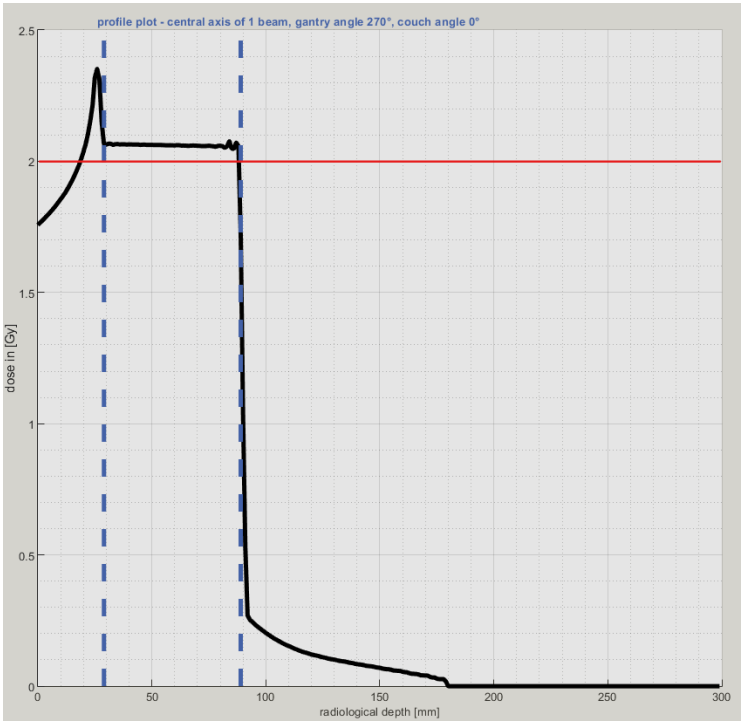

Figure 5.11: SOBP 30–90 mm, carbon ions, without dose constraints.

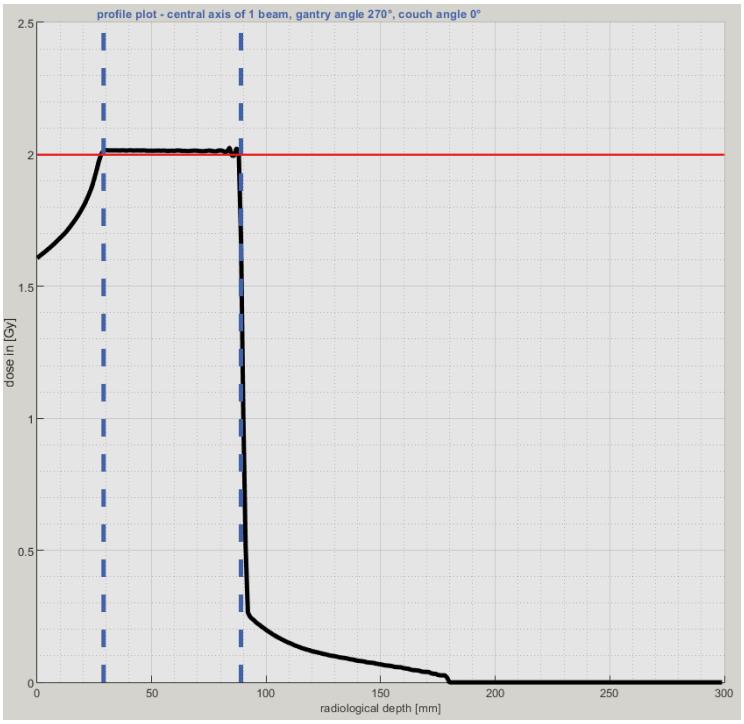

Figure 5.12: SOBP 30–90 mm, carbon ions, with dose constraints.

To obtain a correct dose distribution also in terms of absolute values, it was necessary to explicitly introduce dose constraints into the optimization problem. In particular, three different kinds of constraints were used: *targetOverDose*, which controls the maximum dose to the target; *targetUnderDose*, which controls the minimum dose to the target and is generally set to 90% of the prescribed dose; and *waterOverDose*, which limits the maximum dose to the surrounding water and is set to around 85% of the prescribed dose.

With these constraints applied, the SOBP becomes flatter, and the mean target dose matches the prescription.

A further critical issue emerged in calculations performed with protons, related to the choice of the parameter *longitudinalSpotSpacing*. When this spacing is initially set to 4 mm, as commonly used in clinical practice, matRad selects one energy every 4 mm in depth, effectively skipping part of the available energy spectrum within the considered range. As a consequence, the resulting SOBP appears jagged and non-uniform, as shown in Fig. 5.13. This effect becomes more evident when analyzing the longitudinal dose profile, where the discrete superposition of Bragg peaks produces marked dose oscillations.

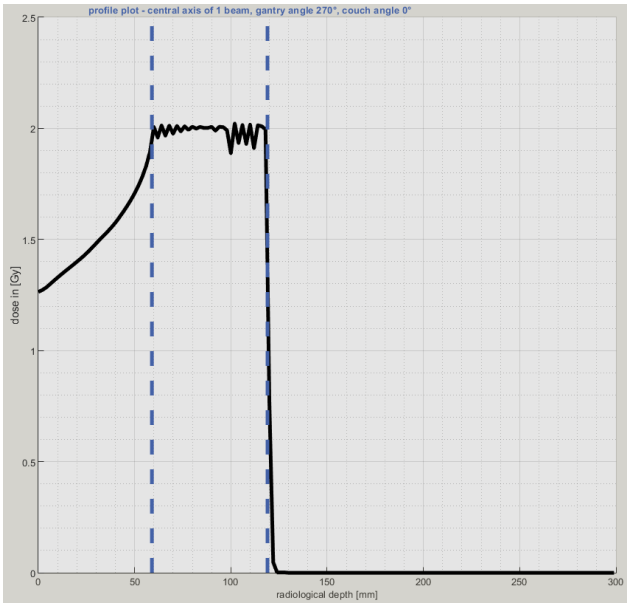

Figure 5.13: SOBP 60–120 mm, protons, longitudinal spot spacing 4 mm.

The adopted solution was to reduce the *longitudinalSpotSpacing* to 2 mm. In this way, matRad makes use of all the energies available in the machine file to cover the depth of the SOBP. As a result, the superposition of the Bragg peaks becomes more continuous and the resulting dose profile is significantly flatter and more uniform, as shown in Fig. 5.14.

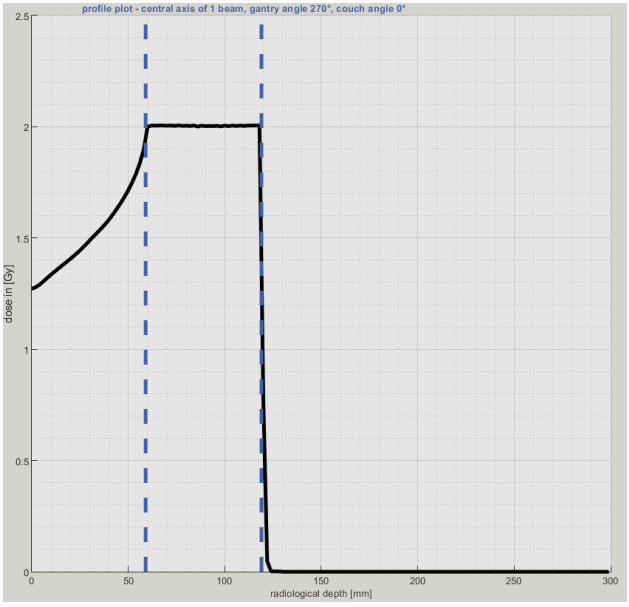

Figure 5.14: SOBP 60–120 mm, protons, longitudinal spot spacing 2 mm.

Overall, these observations highlight how the uniformity of the dose distribution

strongly depends not only on the type of particle used, but mainly on the choices made in the target definition, on the selection and discretizations of the energies, on the introduction of appropriate dose parameters and constraints, and on the geometrical parameters of the beam.

Chapter 6

Results and conclusions

Once the dose uniformity achieved with the code developed in matRad has been verified, it is necessary to assess the compatibility of these results by comparing them with those obtained using professional treatment planning systems (TPS), in order to validate the reliability of the proposed method.

To achieve this result, it is necessary to generate output files that are compatible with the TPS used at CNAO and can be imported as input. A first critical aspect concerns the format of the matRad output files. The software natively returns the dose influence matrix, which represents the dose deposited in each voxel of the simulated geometry. Although this matrix is essential for optimization purposes, it is extremely large and not suitable either for a direct comparison with clinical TPS results or for direct use as input for treatment plan generation in such systems.

For this reason, an alternative output format more consistent with the data structures required by clinical TPS was defined. In particular, a dedicated file named *OutData* was introduced. This file contains, for each spot, the fundamental information needed for plan reconstruction, namely the beam energy, the assigned weight, and the spatial coordinates of the spot.

However, the generation of the *OutData* file alone is not sufficient to enable data import into the TPS, since these systems require input files with a strictly defined syntax and structure. Therefore, an additional code, implemented as the function *writeTxtFile.m*, was developed. This function takes as input the *OutData* file together with auxiliary files containing general information, identifies the type of particle used (protons or carbon ions), and generates a text file fully compliant with the format required by the TPS.

The output file produced by *writeTxtFile.m* contains, for each energy layer, the number of particles assigned to each spot, the corresponding spatial coordinates, and

a set of summary information, including the total number of spots, the total number of delivered particles, the type of radiation used, and the list of energies employed. A relevant aspect concerns the treatment of spots to which a null weight is assigned as a result of the optimization. Since the TPS cannot accept spots with zero particles—because, in clinical practice, a voxel with zero particles corresponds to a specific command for the scanning magnets and follows a dedicated codification sequence—voxels with zero particles in the `.txt` file that do not comply with this codification result in an error when processed by the TPS. For this reason, the function automatically identifies these cases and replaces the null value with a single particle, updating the total number of delivered particles accordingly. This procedure does not introduce significant dosimetric effects, as the number of added particles is negligible compared to the overall total.

Once the input files corresponding to different beam configurations were generated, in particular SOBPs of 60 mm and 30 mm positioned at different depths inside the water box and delivering different dose levels, they were provided to the medical physicists at CNAO. The simulations were then performed using the RayStation TPS. The resulting dose distributions were subsequently exported, analyzed, and compared with those obtained using matRad through post-processing performed in an Excel environment.

From the comparison of the dose profiles produced with protons and carbon ions, some discrepancies immediately emerged between the results obtained with the TPS and those calculated with matRad.

The first discrepancy concerns the presence of a shift along the depth direction in the dose profiles calculated with the TPS with respect to those obtained with matRad. This displacement has been studied in detail and can be mainly attributed to the different spatial resolutions of the two calculation systems. In particular, the TPS uses a grid with a resolution of 2 mm, while in matRad the uncertainty is linked to the dose calculation modality within the single voxel, introducing an additional error of the order of 0.5 mm. This is perfectly consistent with the observed results, since the SOBPs exhibit a shift, evaluated at 80% of the distal fall-off, of about 1.6 mm for carbon ions. On the other hand, for protons, the difference remains below 1 mm, with values ranging between 0.5 and 0.9 mm for both of the considered depths. The behavior of the shift as a function of depth is shown in Fig. 6.1 and Fig. 6.2.

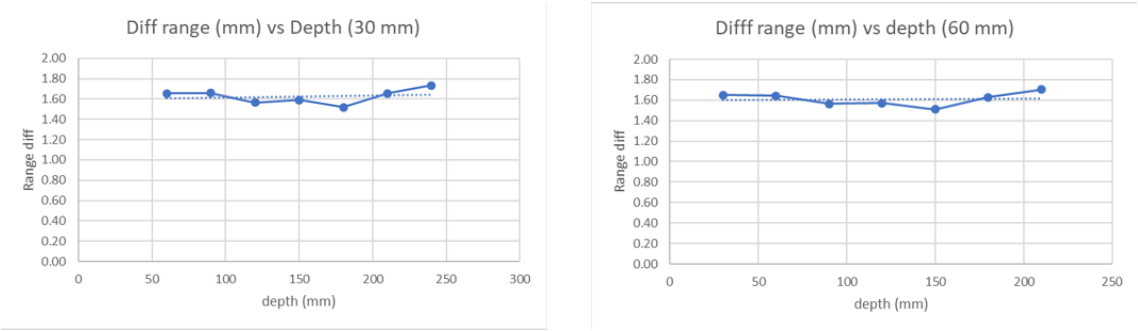

Figure 6.1: Behavior of the longitudinal shift as a function of the different initial depths simulated with carbon ions for SOBP of 30 mm (left) and 60 mm (right).

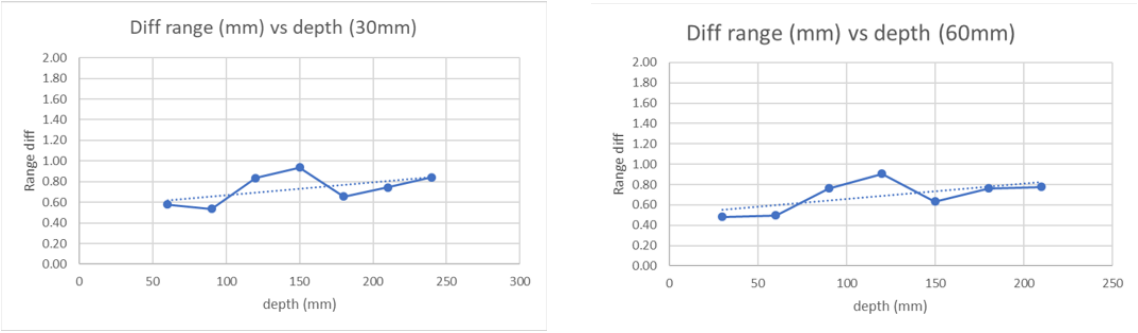

Figure 6.2: Behavior of the longitudinal shift as a function of the different initial depths simulated with protons for SOBP of 30 mm (left) and 60 mm (right).

It is important to note that the measurement of the shift at 80% of the distal fall-off tends to slightly overestimate the actual difference, since the plateaus of the two systems are located at different dose levels. By comparing corresponding points at the same dose height, the average shift for carbon ions is reduced to approximately 1.4 mm, while for protons no significant difference is observed when applying this correction.

This discrepancy can therefore be considered a systematic error, interpretable as a geometrical uncertainty. Such uncertainty is negligible when the interest lies in positioning the target at the center of the SOBP, while it remains manageable, for example through measurements performed at slightly different depths (of the order of a millimeter), if there is a need to place the target close to the distal fall-off of the SOBP.

The second observed difference concerns the dose level at the plateau of the SOBP. In matRad, the plateau consistently reaches the prescribed dose, whereas in the

TPS an underestimation of the dose can be observed. The analysis shows that this underestimation is independent from the prescribed dose level but depends on depth at which the SOBP is located. In particular, it exhibits an approximately linear increase with depth, both for the 30 mm and 60 mm SOBPs and for both particle types, as shown in Fig. 6.3 and Fig. 6.4.

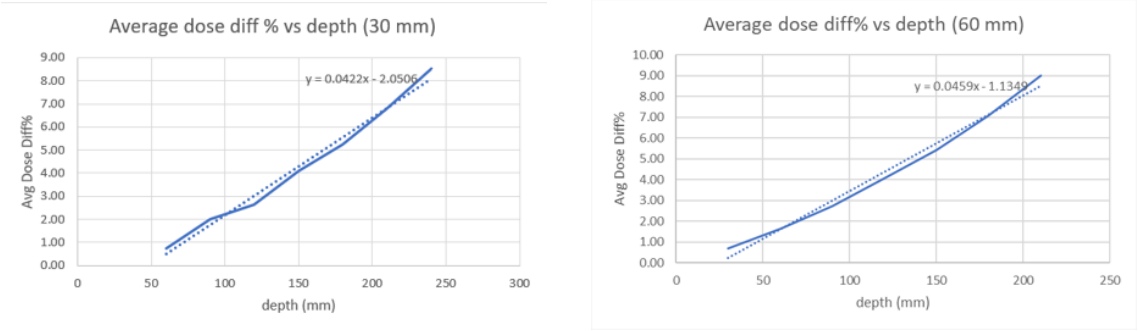

Figure 6.3: Behavior of the average dose difference as a function of the different initial depths simulated with carbon ions for SOBP of 30 mm (left) and 60 mm (right).

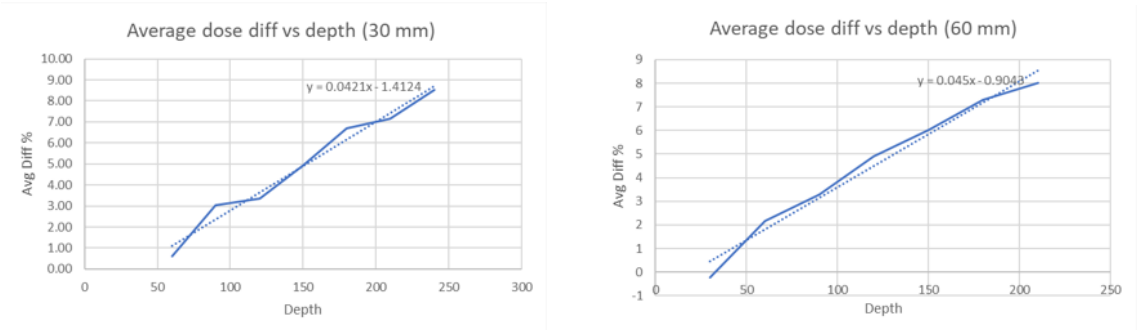

Figure 6.4: Behavior of the average dose difference as a function of the different initial depths simulated with protons for SOBP of 30 mm (left) and 60 mm (right).

Furthermore, especially for carbon ions, the SOBP calculated with the TPS shows a descending profile along the depth, whereas those obtained with matRad are substantially flat.
This discrepancy can be attributed to different factors linked to the distinct calculation strategies and physical modeling approaches adopted by matRad and the TPS. In the first place, a relevant role is played by the lateral beam modeling. In matRad, for protons the lateral distribution is described through a double Gaussian, while for carbon ions a single Gaussian is used. This approach involves rather rough

approximations and is not able to correctly reproduce the complexity of lateral scattering, especially at larger depths in the target.
A further difference arises from the dose calculation engine employed. For carbon ions, both matRad and the TPS adopt an analytical approach based on pencil beam algorithms, with only limited differences in the description of the physical interactions. Conversely, for protons the TPS uses a Monte Carlo approach, which allows for a more realistic representation of the multiple scattering process, whereas matRad still relies on a pencil beam algorithm. In the TPS case, the lateral and out-of-axis contributions are therefore modeled more accurately.
As a consequence, within the SOBP generated by the TPS, a relative lowering of the average dose level—and in some cases the presence of a slope—can be observed with respect to the profiles calculated using matRad.
The combined effects of these modeling differences become progressively more evident with increasing depth in water. In particular, the increasing contribution of scattering and the limitations of the analytical models are reflected in the behavior of the average dose difference percentage as a function of depth in water, as observed in the previous graphs. After a comparison with the medical physicists, it emerged that this behavior is linked to the lateral description of the beam. The TPS uses a triple Gaussian model to describe the lateral broadening of the beam, whereas in matRad a single Gaussian model is used for carbon ions and a double Gaussian model is used for protons. This results in a beam description that is overly approximate and may cause the observed discrepancies in the dose distribution.
Overall, the comparison with a professional treatment planning system represents a fundamental step in the validation of the developed code. Although some discrepancies were observed, especially in terms of SOBP position and absolute dose values, the good qualitative agreement of the dose profiles indicates that the matRad-based approach is physically consistent. Further investigations will be necessary to address the dose differences; in particular, a possible strategy is to study methods to obtain SOBPs in matRad with an ascending profile, similar to those observed in forward calculations, in order to compensate for the differences when reproducing the data with the TPS.

Bibliography

[1] Hyuna Sung and et al. "Global Cancer Statistics 2022: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers". In: *CA: A Cancer Journal for Clinicians* (2024). DOI: 10.3322/caac.21820. URL: https://pubmed.ncbi.nlm.nih.gov/39034807/.

[2] Freddie Bray and et al. "Global, Regional, and National Prevalence of Metastatic Breast Cancer at Diagnosis: A Population-Based Study". In: *JAMA Oncology* (2023). DOI: 10.1001/jamaoncol.2023.4474. URL: https://jamanetwork.com/journals/jamaoncology/fullarticle/2811796.

[3] International Atomic Energy Agency. *Killing More Cancer Cells than Ever Before: A New Era in Radiotherapy*. 2023. URL: https://www.iaea.org/newscenter/news/killing-more-cancer-cells-than-ever-before-a-new-era-in-radiotherapy.

[4] Reshma Jagsi and et al. "Trends in Radiation Therapy among Cancer Survivors in the United States, 2000-2016". In: *Cancer Epidemiology, Biomarkers & Prevention* 26.6 (2017), pp. 963–970. DOI: 10.1158/1055-9965.EPI-16-1025. URL: https://aacrjournals.org/cebp/article/26/6/963/167827/Trends-in-Radiation-Therapy-among-Cancer-Survivors.

[5] Stephanie Chun and et al. "Use of Conventional and Advanced Radiotherapy Techniques in the Palliative Treatment of Bone Metastases in Victoria, Australia". In: *Clinical and Translational Radiation Oncology* 35 (2022), pp. 50–57. DOI: 10.1016/j.ctro.2022.08.006. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC9541909/.

[6] Filippo Alongi and et al. "Patterns of Practice in Palliative Radiotherapy for Bone and Non-Bone Metastases: Results from a Nationwide Survey in Italy". In: *Cancers* 14.19 (2022), p. 4813. DOI: 10.3390/cancers14194813. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC9518763/.

[7] Cornelius A. Tobias et al. “Pituitary irradiation with high-energy proton beams: A preliminary report”. In: *Cancer Research* 18.2 (1958), pp. 121–134. ISSN: 0008-5472.

[8] Marco G. Pullia. “Synchrotrons for Hadrontherapy”. In: *Reviews of Accelerator Science and Technology* 2 (2009), pp. 157–178.

[9] S Agosteo. “Dispense del corso di applicazioni medicali delle radiazioni”. In: (2017).

[10] Fondazione CNAO. *Relazione annuale 2024*. `https://fondazionecnao.it/chi-siamo/relazioni-e-bilanci`. Consultato il 29 settembre 2025. 2024.

[11] Fondazione TERA. *The CNA Accelerator System*. Progetto tecnico. Centro Nazionale di Adroterapia, rapporto interno. Via Cozzi 53, 20125 Milano, Italy: Fondazione TERA, Sezione Milano, c/o Dip. Scienze Materiali, Univ. Milano Bicocca, Sept. 2003.

[12] Varian Medical Systems. *Eclipse Treatment Planning System v18.1 – Feature Sheet*. `https://www.varian.com`. Accessed: 2025-10-28. n.d.

[13] RaySearch Laboratories. *RayStation v2025 – Automated Treatment Planning Using Machine Learning (ECHO)*. `https://www.raysearchlabs.com`. Accessed: 2025-10-28. n.d.

[14] Elekta. *Monaco 5.11 – Up to 4× Faster Calculation Speed*. `https://www.elekta.com`. Accessed: 2025-10-28. n.d.

[15] B. Gottschalk et al. “Multiple Coulomb scattering of 160 MeV protons”. In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 74.4 (June 1993), pp. 467–490.

[16] Niklas Wahl. *About matRad*. `https://github.com/e0404/matRad/wiki/about-matRad`. Pagina modificata il 17 giugno 2021. June 2021. URL: `https://github.com/e0404/matRad/wiki/about-matRad`.

[17] Hans-Peter Wieser et al. “Development of the open-source dose calculation and optimization toolkit *matRad*”. In: *Medical Physics* 44.6 (2017), pp. 2556–2568. DOI: `10.1002/mp.12270`.

[18] E. Pedroni et al. “The dose distribution of proton pencil beams”. In: *Physics in Medicine and Biology* 40.3 (1995), pp. 437–458. DOI: `10.1088/0031-9155/40/3/004`.

[19] L. Hong et al. “Proton dose calculation in heterogeneous media”. In: *Physics in Medicine and Biology* 41.8 (1996), pp. 1305–1326. DOI: 10.1088/0031-9155/41/8/003.

Ringraziamenti

A Simone e Marco, che in questo anno mi hanno seguita con pazienza e mi hanno fatto crescere permettendomi di scoprire l'affascinante ambiente di CNAO e dell'adroterapia. Voglio ringraziare anche la professoressa Francesca Ballarini che mi ha indirizzata verso questo percorso.

A Mamma e Papà, che da ben 24 anni mi sopportate. Non ci sono molte parole per descrivere quanto vi sono grata, più passa il tempo più capisco di essere fortunata ad avere due genitori come voi. La soddisfazione più grande è vedere l'orgoglio nei vostri occhi quando parlate di me. Ma dovete ricordare che se sono arrivata fin qui è solo ed esclusivamente merito vostro, che mi avete supportato sotto ogni punto di vista e mi siete stati vicini nei momenti più difficili. Anche se non sempre mi capivate, avete sempre fatto tutto il possibile per comprendermi.
Spero di poter diventare come voi un giorno e di tramandare ai miei figli quello che voi avete trasmesso a me, ovvero che la felicità si nasconde dietro le cose più semplici, una carezza, un piccolo gesto, e soprattutto nel dare e ricevere amore incondizionato.

All'équipe di Fermo, che da due anni a questa parte mi ha seguita lungo un percorso tortuoso. Ho raggiunto questo traguardo anche grazie a voi: con il vostro aiuto sto piano piano imparando a gestire le montagne russe delle mie emozioni e, col tempo, ad amare me stessa. Decidere di iniziare la terapia con voi è stata una delle scelte migliori che potessi fare, anche se non la più semplice. Mi avete aiutata a fare pace con me stessa e con il cibo, a riappropriarmi delle mie emozioni e ad accettarle come parte di me.
Con voi ho capito che il mio valore non dipende dal mio rendimento né tantomeno dal mio peso. Certo, la strada è ancora lunga, ma sono fiduciosa e determinata ad arrivare fino in fondo.
Grazie perché mi avete aiutata a ricominciare a vivere, anziché sopravvivere.

Ai miei nonni Livio, Carolina e Sandra. Credo che chiunque possa confermare che crescere con i propri nonni accanto sia una delle cose più preziose del mondo, e di certo non sarò io a dissentire. Infatti, grazie alla vostra esperienza, mi avete permesso di guardare il mondo da nuove prospettive e, soprattutto, mi avete fatto il regalo più grande di tutti: il vostro affetto e la vostra presenza.

A tutto il resto della mia famiglia, grazie per tutto l'aiuto e la vicinanza dimostrata in questi anni.

A tutti i miei amici, che mi hanno insegnato che conta la qualità, e non la quantità, del tempo trascorso insieme. Ognuno, a modo proprio, mi ha lasciato qualcosa che porterò con me e di cui farò tesoro negli anni a venire.
Grazie per tutto il supporto e per i momenti passati insieme.

A Marco, nonostante tu sia l'ultimo arrivato nella mia vita, in breve tempo sei diventato una delle persone più importanti. Mi hai teso la mano per farmi uscire dal baratro in cui stavo lentamente sprofondando, e forse non te ne sei nemmeno accorto. Lo hai fatto nel modo più semplice, essendo presente. Mi hai aiutato a riscoprire la felicità che non credevo di meritare, hai rotto la corazza che mi ero costruita negli anni facendo riaffiorare aspetti del mio carattere che avevo da tempo sepolto.
Mi hai stravolto l'esistenza e te ne sono infinitamente grata.

Voglio dedicare qualche parola anche al mio animaletto, o meglio al mio fratellino Beethoven. Sei quello che mi è stato fisicamente più vicino nelle interminabili ore di studio, accucciato tra i miei piedi sotto la scrivania. Istintivamente sei sempre riuscito a capire quando era un momento no e a strapparmi un sorriso.

Per ultimo, ma non per importanza voglio ringraziare e dedicare le parole finali a me stessa. Alla mia tenacia e ostinazione e, a prescindere dal risultato, mi auguro di essere sempre orgogliosa del lavoro fatto in questi anni e di smetterla, almeno per una volta, di considerarmi sempre l'ultima ruota del carro. Quindi Sara, datti una pacca sulla spalla, perché hai raggiunto i tuoi obiettivi nonostante tutto. Cerca di vivere con più serenità ciò che ti aspetta nei prossimi anni e ricorda che anche ciò che non consideri perfetto può esserlo, basta vederlo con gli occhi giusti.

Nei momenti no basta sempre guardarsi indietro e ripensare a tutti gli ostacoli che si sono superati nel corso della vita, dopotutto ognuno ha le proprie sfide e i propri demoni da affrontare. Ricordate che siamo molto più forti di quello che crediamo, basta solo prendere il coraggio necessario per buttarsi nelle sfide di tutti i giorni. Non importa se perdiamo o vinciamo, l'importante è provarci perché meglio imparare da un fallimento che vivere col rimpianto di non aver tentato.

Grazie di cuore a tutti!