POLITECNICO DI TORINO

Faculty of Engineering

Master of Science in Biomedical Engineering

Master Thesis

SCATTERING - INFORMED TOMOGRAPHIC IMAGE RECONSTRUCTION FOR ION IMAGING

Advisors
Prof. Filippo Molinari
Dott.ssa Chiara Gianoli
Prof.ssa Katia Parodi

Candidate
Antonio Micoli

Academic Year 2023-2024

From you, 2000 years ago

Acknowledgements

I am sincerely grateful to my family for all the love and support shown over the years.

To my friends, who always knew what to say and how to say it: thank you for being there when I needed it.

I would also like to thank Prof. Katia Parodi for welcoming me into her department and making this experience possible.

Finally, I owe a special debt to Chiara Gianoli and Ines Butz, for the patience and kindness they had towards me throughout this journey.

Summary

Radiotherapy stands as the leading treatment for cancer. Though the potential of X-rays is proven, the feasibility of using charged particles is gaining growing interest. While photons show an exponentially decreasing dose, ions release ever more energy along their path, culminating at the Bragg peak. Thus, targeting accuracy and preservation of healthy tissues are improved. Nonetheless, ion imaging, which is the native imaging technique for ion beam therapy, is still a long way from being a clinical reality: limited image quality and resolution due to Multiple Coulomb Scattering events across the medium were thoroughly assessed within this project.

Ion imaging relies on two key quantities: the residual energy (stopping power), translating to Water Equivalent Thickness (WET), and (if available) particle deviation from a straight trajectory (scattering power). Upon combining these measurements with an estimation of the ion trajectory, iterative tomographic reconstruction algorithms were adopted to retrieve a map of the ion stopping power ratio relative to water, or Relative Stopping Power (RSP). Two detector setups were explored. List-mode detectors provide spatial and energy data. The Most Likely Path (MLP) algorithm is then applied to reconstruct the individual proton trajectory (scattering line). Since the actual scattering power can be measured only at the entry and exit location, the estimated path is affected by statistical uncertainties in the middle. By introducing a Gaussian scattering model of those uncertainties into the MLP algorithm, a volumetric distribution of all possible trajectories was created (scattering spindle). Integration-mode detectors drop all spatial information, so the pencil beam trajectory is resolved as straight (nominal line). However, as its nominal dimension and direction are known, the scattering model can be adapted as a conical Gaussian distribution (scattering cone). Since WET quantities are carried by the ensemble of protons, a value statistically representative to the whole pencil beam was extracted from WET occurrence histograms. To exploit the whole histogram, a novel approach where a Gaussian cone is realized for each WET component (scattering cones) was explored. The images were reconstructed with and without embedding the Gaussian scattering models into the estimated trajectories. All models were tested in a water- and a tissue-based reconstruction. For list-mode images, the Gaussian spindle provides a better image quality and robustness to noise, with a relative improvement in water up to 10% over the scattering line. Integration-mode images see a relative improvement over 6% for the water cone over the nominal line, while the tissue cone is not able to provide a superior quality. The scattering cones, instead, show a slower convergence due to more abundant reconstruction data, but a lower error when compared to the respective single cone model (over 10% improvement). Though the tissue models are more accurate, the water ones generally offer better results. Since the water scenario usually implies higher WET integrals along the ion trajectory, the respective Gaussian envelop is more pronounced, thus likely adding an intra-reconstruction smoothing effect that favors image stability and convergence. Predominant in integration-mode configuration, the same phenomenon can occur in list-mode images.

Future studies may explore the introduction of intra-reconstruction filtering methods in tissue models to investigate a possible compensation over the water-based scenario. Data-driven approaches may also benefit from these models.

List of Abbreviations

WHO World Health Organization

IARC International Agency for Research on Cancer

DNA Deoxyribonucleic Acid

PBT Proton Beam Therapy

CIBT Carbon-Ion Beam Therapy

BNCT Boron Neutron Capture Therapy

CPT Charged Particle Therapy

RBE Relative Biological Effectiveness

OER Oxygen Enhancement Ratio

SOBP Spread-Out Bragg Peak

PSS Passive Scattering System

ASS Active Scanning System

CT Computed Tomography

IGRT Image-guided Radiotherapy

PET Positron Emission Tomography

MRI Magnetic Resonance Imaging

ART Adaptive Radiotherapy

DIR Deformable Image Registration

RSP Relative Stopping Power

WET Water Equivalent Thickness

iCT Ion Computed Tomography

PSD Position-sensitive detector

RERD Residual energy-range detector

PSI Paul Scherrer Institute

LLU Loma Linda University

UCSC University of California Santa Cruz

NIU Northern Illinois University

pCT Proton Computed Tomography

CSUSB California State university, San Bernardino

SiSDs Silicon strip detectors

FWHM Full-width half-maximum

CNAO National Centre for Oncological Hadron Therapy

MCS Multiple Coulomb Scattering

IR Iterative Reconstruction

DLR Deep Learning Reconstruction

PDF Probability Distribution Function

iRad Ion Radiography

MLP Most Likely Path

OS-SART Ordered Subsets Simultaneous Algebraic Reconstruction Technique

SART Simultaneous Algebraic Reconstruction Technique

GSI Gesellschaft für Schwerionenforschung

HIT Heidelberg Ion Beam Therapy Center

MAE Mean Absolute Error

Contents

List of Figures

List of Tables

I believe that nothing happens by chance.
Deep down, things have their own secret plan,
even though we don't understand it.

[Carlos Ruiz Zafón, The Shadow of the Wind]

Chapter 1

Introduction

1.1 Radiation therapy

Cancer remains one of the main challenges of modern medicine: according to the specialized cancer agency of the World Health Organization (WHO), i.e. the International Agency for Research on Cancer (IARC), at least one in five people worldwide develop malignant tumors during their lifetime. Past decades have witnessed a significant advancement in understanding the underlying mechanisms of cancer and its hallmarks, shaping the path for different types of treatment modalities such as chemotherapy, surgery, immunotherapy, hormonal therapy and, most notably, radiation therapy (Baskar et al. [2012]).

The discovery of X-rays at the hands of German physicist W.C. Roentgen (1895) usually singles out the moment the idea of using radiations as a form of cancer treatment started to take hold. Shortly after, in fact, X-rays were tested for the very first time on breast cancer patients (Hu and Kim [2020]). In 1898, the discovery of radium as a radiation source by Marie and Pierre Curie also helped in advancing the field: indeed, by the start of the next century, an interest for applications of radiations in medicine had been rapidly growing (Gianfaldoni et al. [2017]). Since then, radiation therapy, or *radiotherapy*, has been scientifically recognized as medical specialty for oncologic patients: it is estimated that at least 50% of all cancer patients worldwide undergo radiation therapy, 40% of whom for curative intents (Barnett et al. [2009]).

Radiotherapy relies on ionizing radiations able to alter matter at a microscopic scale by removing electrons from their atoms, i.e. they can ionize atoms. Ionizing radiations passing through tissues are thus capable of induce DNA lesions that, if left unchecked, can lead to different forms of cell death, such as apoptosis (Begg et al. [2011]). Usually, healthy cells can benefit from meticulously controlled DNA repairing processes, but these same mechanisms are instead impaired in cancer cells, resulting in a differential killing in the tumor area (Baskar et al. [2012]). Nevertheless, exposure of healthy cells to ionizing radiations is still a factor to be dealt with: owing to this, the overall dose is usually fractionated over several treatment sessions so to match the healing times of normal cells.

Another factor adding to radiotherapy's undisputed prestige over the years is its versatility. First of all, it allows not only for curative treatments but also for palliative care. Secondly, it may play a role in combination with other strategies: for instance, it can be used before surgery to reduce the cancerous area, or it can be used after surgery so to destroy any remaining cancer cells. Finally, radiotherapy's main feature may very well be the possibility of employing different physical agents for ionization purposes.

In the context of Radiation Oncology, two main types of radiations are at play, i.e. photon and particle radiations.

X-rays and γ-rays are key protagonists when it comes to photon beam radiotherapy. Usually X-rays are produced by accelerating electrons against a high atomic number material such as Tungsten-Molybdenum making up the anode of an X-ray tube. Instead, γ-rays are the result of radioactive decay from substances such as radium and cesium.

On the other hand, three types of particle beam therapies are currently being used for cancer treatment: proton beam therapy (PBT), carbon-ion beam therapy (CIBT) and boron neutron capture therapy (BNCT). Though in 1946 Wilson suggested the possibility of using accelerated protons for cancer therapy, it wasn't until 1958 that the first clinical test was performed, when the Lawrence Radiation Laboratory in Berkeley

treated the pituitary gland of 26 patients with advanced breast cancer. However, the first medical facility to ever use cyclotron-accelerated protons was at the Clatterbridge Oncology Center (UK) in 1989, followed by Loma Linda University (CA, USA), which used a synchrotron instead. In 1994, HIMAC (Chiba, Japan) initiated the first medical treatment for CIBT. Nowadays, the United States of America holds the highest number of treatment facilities worldwide, counting up to forty, followed by Japan with twenty-four facilities, Germany with seven facilities and Russia with five facilities. Most of them are PBT facilities, although CIBT facilities can still be accounted for (Matsumoto et al. [2021]).
It is estimated that, up to 2019, over 250.000 patients had been treated with PBT and CIBT. In fact, charged particles are rapidly proving themselves to be valid candidates for radiation therapy, although the high-cost equipment needed to accelerate them represents right now the greatest challenge to its clinical consolidation. Nevertheless, a critical review between the physical properties and biological effects of photons and ions as treatment agents in human tissues may explain the increasing popularity of ion beam therapy in recent years as a reasonable alternative to photon beam therapy.

1.2 Ion radiobiology

Radiotherapy's main goal is to deliver as much dose as possible to the tumor volume while sparing surrounding healthy tissues. This outcome is greatly affected by the radiation's source, since different radiations interact in different ways with human tissues.
The quality of Charged Particle Therapy (CPT) can be assessed by the *Linear Energy Transfer* (LET) parameter, that is the average energy deposited per unit path length (keV/μ) into the medium by an ionizing particle. After exposure to ionizing radiations, cell survival S can be expressed through the *Linear-Quadratic model* (Durante and Paganetti [2016]):

$$S = e^{-\alpha D - \beta D^2} \tag{1.1}$$

where D is the delivered dose in gray (Gy), while α, the linear term dominating at low doses, and β, the quadratic term dominating at high doses, take into account tissue sensitivity to radiations. Plotting cell survival S (Fig. 1.1) on a logarithmic scale offers a better interpretation on the contribution of these two parameters to cell killing (McMahon [2019]).

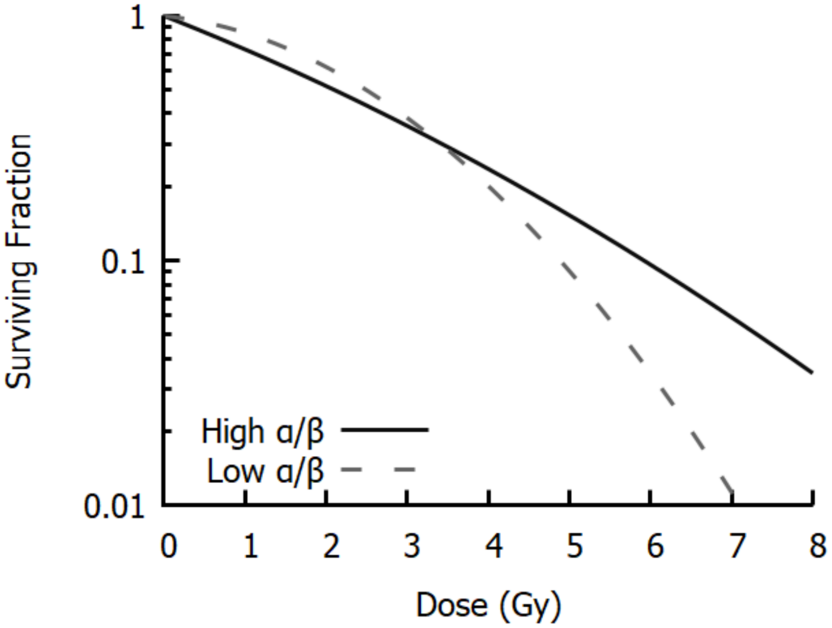

Figure 1.1: Logarithmic plot of Linear-Quadratic model for cell survival. High α/β ratios have a nearly-constant cell killing rate as dose increases, while low α/β ratios present a significant curvature with increasing killing rates as dose increases (McMahon [2019]).

When the α/β ratio is high, the killing rate does not undergo significant variations; instead, low α/β ratios show a pronounced curvature with increasing dose. This behaviour is key for fractionation into smaller doses delivered over time of the total radiation dose (Durante and Paganetti [2016]): eventually, split doses will spare normal tissues (low α/β ratios) while still preserving a substantial killing power towards tumors (high α/β ratios). Densely ionizing radiations generally results in a higher α/β ratio: high-LET radiations, in fact, exhibit an exponential behaviour, which translates to an almost straight line on a logarithm plot; on the other hand, sparsely ionizing radiations, i.e. low-LET radiations, show a plot characterized by an initial curvature turning into an almost straight line at higher doses.
To quantitatively mark different biological effects due to differences in LET properties is the *Relative Biological Effectiveness* (RBE), that is usually defined as the dose ratio between the reference photon and particle resulting in the same biological effect. It was already demonstrated that heavy ions have an high RBE values relative to X-rays due to a steeper survival cell curve resulting from an increased α parameter (Durante and Paganetti [2016]), with carbon ions, for instance, standing at a clinical RBE of 3.0 (Matsumoto et al. [2021]). Furthermore, the benefits of treating tumors with high-LET ions extends to other parameters as well, such as the oxygen enhancement ratio (OER), i.e. the radiation dose required to achieve the same biological effect in an anoxic environment compared to an oxygenated environment: particularly, C-ions have notoriously a low OER, making them optimal candidates against hypoxic tumors.
Nevertheless, the current general trend is towards light particles, that is protons. However, Fig. 1.2 makes it clear that protons' biological outcome is not that far from that of X-rays and γ-rays.

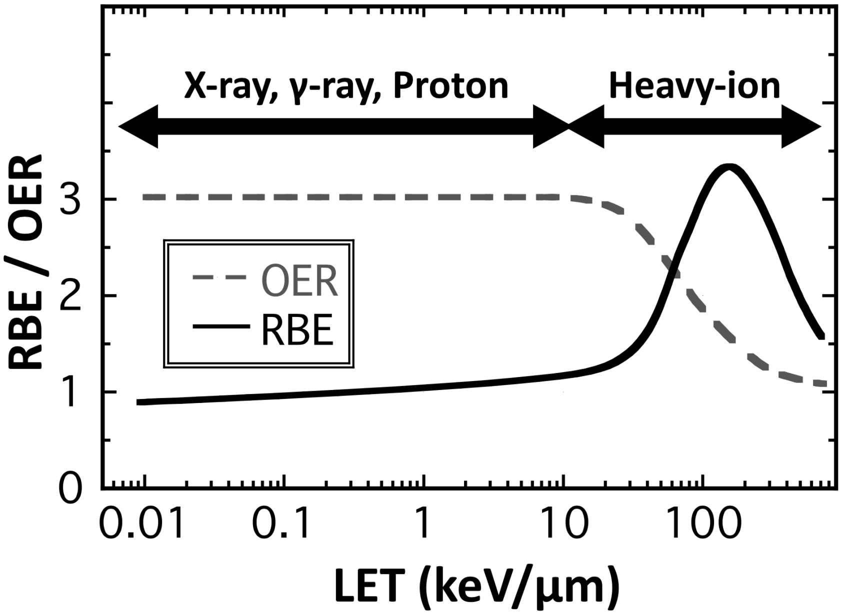

Figure 1.2: Relationship between OER, RBE and LET. X-rays and protons roughly end up with a similar biological outcome, while heavy ions show actual improvement when it comes to their effectiveness in killing cells and in hypoxic enviroments (Matsumoto et al. [2021])

Protons and photons are both low-LET radiations and the international standard for proton RBE is 1.1, making them slightly more effective than photons. Moreover, this RBE value was recommended by the International Commission of Radiations Units and Measurements to be uniformly applied to both cancer cells and healthy tissues. Given that proton radiation does not result in an actual biological advantage (Wang [2015]), the reason that led to its rising popularity in the last years should be traced back to the so called depth-dose distribution profile: it is in fact the distribution of energy deposited per unit mass by charged particles in general that prompted their involvement into the clinical scenario and that provides a compelling justification to prefer protons over photons.

1.3 Ion Beam Therapy: clinical rationale

Dose-delivering properties largely differ when comparing protons and C-ions to photons (Matsumoto et al. [2021]). The profile describing how radiation dose is deposited as it travels through tissue is known as depth-dose distribution profile (Fig. 1.3) and it represents a helpful tool when trying to understand how different ions operate.

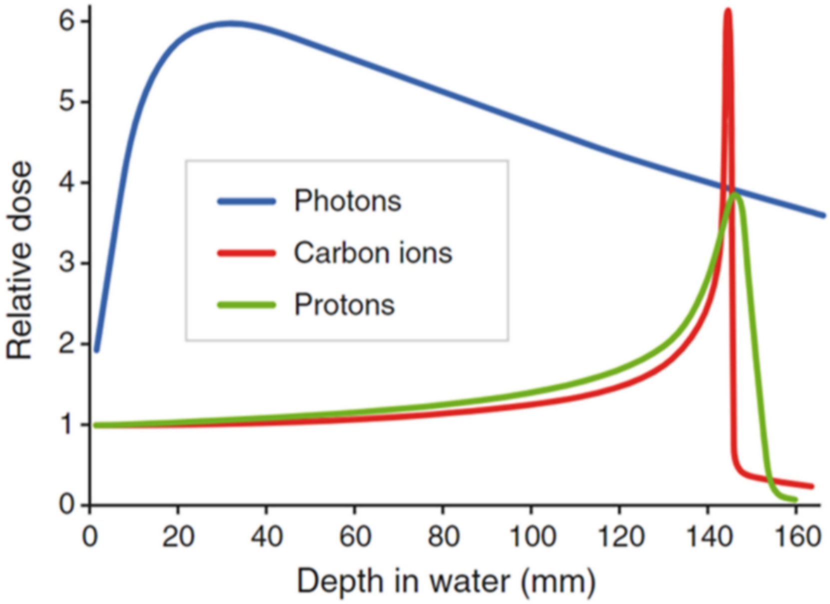

Figure 1.3: Depth-dose distribution profile of protons, carbon ions and X-rays. Photon beams have their highest dose delivery within few centimeters from the patient's body surface. Ion beams deposit their highest dose at larger depths from the body surface. By adjusting their initial energy, ions are able to reach deep-seated tumors with a minimal ionizing effect towards surrounding tissues (Grau et al. [2020]).

Photons are characterized by an exponential dose falloff as they go further into tissues, so in-depth tumors would require higher energy radiations delivering extra dose to tissues in front of the tumor itself. On the other hand, ions show a moderate entrance dose that then starts growing along its trajectory, culminating into what is known as *Bragg peak* (Schneider and Pedroni [1995]). Since charged particles have their highest dose delivery at larger depths and essentially zero dose delivery after the Bragg peak itself, healthy tissues surrounding the cancerous area are thus targeted by a minimal radiation dose. By adjusting the ion's incoming energy from the accelerator it is possible to modulate the depth of their Bragg peak position, increasing the targeting accuracy. Although very similar, differences in the absorption profile of carbon ions and protons are mainly due to C-ions higher mass and charge: particularly, the Bragg peak is sharper, allowing for a more precise energy deposition at the tumor site and a narrower penumbra when the dose decreases to background levels. Also, carbon ions show a fragmentation tail definitely more extended than in protons: particles heavier than protons are affected by several nuclear interactions leading to a more gradual decrease in dose after the Bragg Peak due to the contribution of secondary particles produced by these interactions.

As a clinical standard, the pristine Bragg peak should be extended to fit the entire tumor area (Fig. 1.4): this so-called spread-out Bragg peak (SOBP) is formed by combining single peaks of different energies and intensities (Matsumoto et al. [2021]). This can be done either through a passive or an active modulation of the primary beam (Fig. 1.5). Particularly, passive scattering systems (PSS) rely on a rotating scatterer able to enlarge the beam vertically and horizontally, while a collimator and a bolus are added so to adapt the beam to the tumor's shape. In active scanning systems (ASS) the beam is instead expanded by using electromagnets. While ASS no longer need to rely on a collimator and bolus and though scanning provides a superior dose distribution in comparison to PSS, one major problem with this approach is hidden in the difficulty of managing moving organs caused, for instance, by the patient's breathing during scanning. In fact, the position-accurate pencil beams used in ASS are much more sensible to the geometrical variations of the target's contours while PSS is designed to cover a broader region by spreading the beam on the tumor and its surroundings (Durante and Paganetti [2016]).

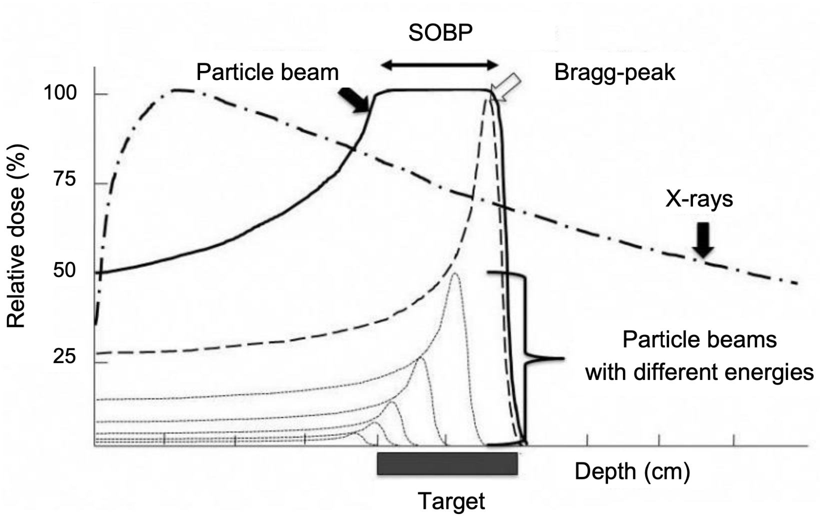

Figure 1.4: Spread-out Bragg peak (SOBP). The depth-dose distribution profile for a particle beam's SOBP is formed by overlapping of several pristine Bragg peaks (Matsumoto et al. [2021]).

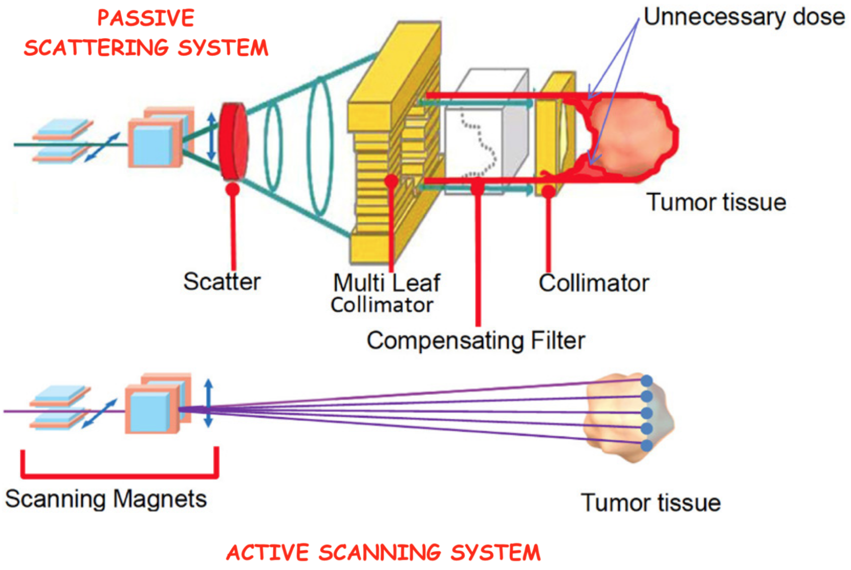

Figure 1.5: Passive and active beam delivery systems. Pencil beam scanning provides greater flexibility and precision in dose delivery, enabling highly conformal treatments even for complex tumor shape. Passive modulation is easier, although less adaptable to irregularly shaped tumors, and allows for a simpler management of moving targets such as lung tumors (Durante and Paganetti [2016]).

In short, despite the limited sample of patients, clinical results do support the rationale for ion beam radiotherapy based on an improved depth-dose distribution for charged particles in general and enhanced radiobiological effects of heavy ions in particular. Still, the high costs associated with particle accelerators remain a significant challenge. The ongoing debate regarding their cost-benefit ratio can only be resolved through the construction of new facilities to carry out additional clinical trials and the implementation of advanced technologies (Durante and S.Loeffler [2010]).

1.4 The role of imaging in radiation oncology

A crucial step in planning radiation treatment is to ensure that the anatomical borders of the cancer area are well defined, so to avoid the possible exclusion of tumors from the targeted region or the accidental inclusion of normal tissues into it (Puspasari and Sulistyo [2022]). The introduction of Computed Tomography (CT) by Cormack and Hounsfield in 1972 allowed for an actual visualization of the tumor geometry, leading to greater precision in dose distribution and patient positioning (Pereira et al. [2014]).
As shown in Fig. 1.6, the role of imaging can be declined into offline imaging outside the delivery room and online imaging, i.e. image-guided radiotherapy (IGRT) inside the treatment room (Dawson and Ménard [2010]).

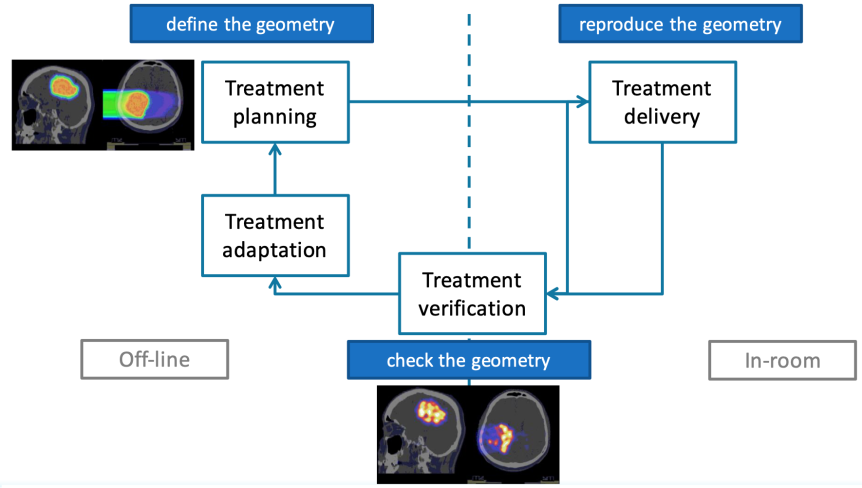

Figure 1.6: Imaging workflow in radiation oncology. CT scans acquired in treatment position before delivery remain essential in contouring tumor target volumes and identifying the geometry of surrounding *organs at risk*. Furthermore, the problem of moving organs has been addressed by image-guided radiotherapy by relying on real-time imaging and motion management to precisely target tumors, adjusting for organ movement during dose delivery.

Usually, treatment begins with a radiation treatment planning based on X-ray CT images acquired while the patient is immobilized. Furthermore, complementary imaging modalities can be adopted as well: in fact, multi-modality imaging is usually exploited to get a patient model covering the anatomical information through X-ray CT primary images but also functional and anatomical information through PET and MRI secondary images.
Treatment delivery is the in-room procedure required to match the patient model to the actual anatomy at play and imaging at this stage also represents a critical need nowadays. Previously, only bony structures could be visualized on the treatment couch of the linear accelerator during each radiotherapy fraction.
Modern developments have now led to equipment of advanced on-board imaging systems on linear accelerators: most notably, cone-beam CT (CBCT) allows to monitor the position of the tumor and surrounding organs over the course of the treatment, enabling immediate adjustments if any movement is detected. Moreover, MRI-linear accelerators are able to provide further reliability in imaging assessment, both where organ motion is involved and where tumor and healthy tissues share similar densities (Chandra et al. [2021]).
Treatment verification is also highly dependent on image acquisition. To deal with dosimetric uncertainties in photon beam radiotherapy, an Electronic Portable Imaging Device (EPID) installed onto the linear accelerator can be exploited to ensure that the dose is delivered as planned: they are able to capture images of the radiation field during treatment, allowing for confirmation of the treatment's outcome. Ion beam therapy can instead rely on indirect verification because of secondary radiations emitted due to nuclear interactions of charged particles with surrounding tissues (Kraan [2015]). On one hand, when an ion beam interacts with the patient's tissue, it can induce nuclear reactions in the atoms of that tissue, leading to the formation of positron-emitting isotopes (β^+ emitters): exploiting PET imaging for detection of 511 keV gamma pairs produced in the annihilation of positrons, clinicians can verify that the ion beam has reached

the intended depth and has accurately targeted the tumor by comparing PET images with the planned dose and contoured volume. On the other hand, another class of secondary emissions are prompt gamma photons produced by fast de-excitation of excited nuclei following the interaction between charged particles and organic tissue components. This second strategy allows for real-time treatment verification since the gamma photons are emitted almost immediately during the irradiation process, contrary to annihilation photons which are delayed due to decaying times of β^+ emitters.

Finally, since new information can emerge during treatment, it is of the essence to account for day-to-day patient settings and alignment as well as uncertainties about the extent of the disease (Puspasari and Sulistyo [2022]): adaptive radiotherapy (ART) can often benefit from advanced tools for deformable image registration (DIR) to map changes in the patient anatomy over time, facilitating the alignment of images from different time points and aiding in the accurate adjustment of treatment plans.

1.5 Calibration of the CT data

Since Computed Tomographic images provide the basic information for treatment planning and verification (Schneider and Pedroni [1995]), the imaging modality should be as representative as possible of the underlying interactions between radiation and traversed matter. Nevertheless, this is not entirely the case when it comes to Ion Beam Therapy.

To treat a cancer patient with ions, two important quantities are to be calculated. The first one is the *ions' range* needed for the Bragg Peak to be deposited at the desired depth: in other words, oncologists need to precisely determine the tumor's position inside the targeted volume, so to set that incident energy necessary for ions to stop at the tumor's location. Secondly, the dose delivered during treatment procedures should also be carefully planned prior to the positioning of the patient in the treatment room (de Assis et al. [2005]). These steps could be performed relying on an imaging approach that uses the ion beam itself to visualize both tumor and healthy tissue, derive treatment planning information and study the alignment of patient and ion beam in the treatment room. Particularly, the potential of ion radiography and ion computed tomography for delineation purposes had already been investigated (Schneider and Pedroni [1995]): during the late 1960s it was demonstrated that, by using parallel-sided objects with a thickness close to the range of the incident protons, it was possible to obtain proton radiographs showing unusually high contrast in comparison to X-radiographs taken under the same conditions, but also a poor spatial resolution (Fig. 1.7).

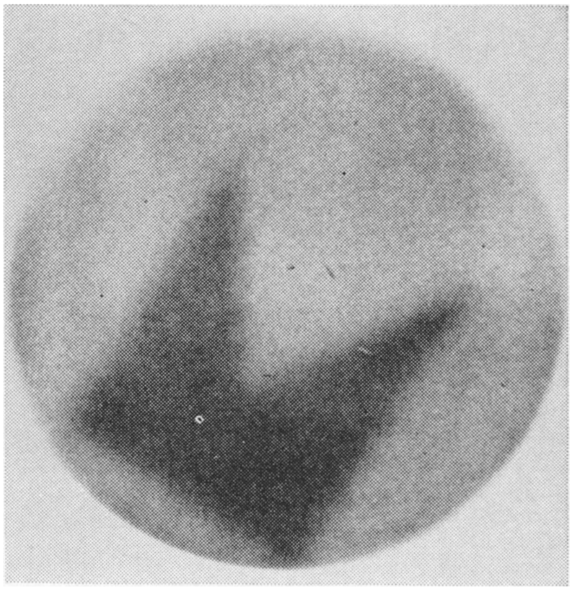

Figure 1.7: Proton radiography. A proton radiograph of an aluminum absorber, which is 7 cm in diameter and 18 g/cm² thick, includes an extra 0.035 g/cm² of aluminum foil shaped like a pennant. This foil is inserted at a depth corresponding to 9 g/cm². Adding 0.2 percent to the overall thickness results in a significantly darker area on the film (Koehler [1968]).

However, the interest in this imaging modality soon disappeared due to the increasing success of X-ray imaging during the latest decades, leading to its predominance as a standard routine even in small hospitals (de Assis et al. [2005]). Thus, it is nowadays well-established to derive the patient's model from X-ray scans for both photon and ion therapy as first step in a treatment plan: the most critical flaw in this approach is that photons and ions propagate very differently in the patient's tissues, making an X-ray CT ultimately inappropriate for range calculations.

When a monochromatic photon radiation enters some homogeneous material, its initial intensity I_0 will be attenuated to a value I due to energy absorption phenomena across the thickness L of the material itself. This process is described by Lambert Beer's equation:

$$I = I_0 e^{-\mu L} \tag{1.2}$$

where μ is the linear attenuation coefficient, which varies according to tissue type, density and energy of the X-ray beam. In CT imaging, the projection data, or radiography, refers to the intensity of the transmitted X-ray after passing through the body: these measurements are recorded by detectors surrounding the patient and, when acquired for multiple rotation angles, they are collectively labeled as Radon Transform. The tomographic image of the attenuation coefficients μ is enabled by Lambert Beer's law that models the projection as a line integral of the physical variables. In fact, for a given projection angle θ and position t, the projection data is given by:

$$P_\theta = ln \frac{I_0}{I_\theta(t)} = \int_{L(t)} \mu(x, y) dl \tag{1.3}$$

Therefore, the resulting distribution of the photoabsorption coefficient μ(x,y) can be reconstructed by inverting the Radon Transform. Finally, since human tissues are mostly made up of water, the radiodensity of the reconstructed image is usually expressed as "CT units", also known as Hounsfield units (de Assis et al. [2005]):

$$HU(x, y) = 1000 \frac{\mu(x, y) - \mu_{water}}{\mu_{water}} \tag{1.4}$$

While X-ray CT relies on maps of photons' linear attenuation, ion CT is typically operated in terms of energy loss or *residual range* of the ions traversing the patient. The mean energy loss per unit path length, also known as *stopping power*, can be calculated through the well-known Bethe-Bloch formula (Schneider and Pedroni [1995]):

$$S = \frac{dE}{dx} = K \frac{\rho N_g}{\beta^2} \left\{ ln \frac{2 m_e c^2 \beta^2}{I(1 - \beta^2)} - \beta^2 \right\} \tag{1.5}$$

where K is a constant, βc is the velocity of the projectile, I is the mean ionization energy of the absorbing material, m_e is the mass of the electron and ρN_g is the electron density of the traversed matter. In the context of ion CT, the ion's residual energy is measured by a detector and then converted to Water Equivalent Thickness (WET), i.e. the projected data or the *ion radiography* (iRad). Water Equivalent Thickness represents the thickness of water leading to the same amount of energy loss caused by the actual material being penetrated and can be mathematically modeled as:

$$WET = -\int_{E_{in}}^{E_{out}} \frac{1}{S(I_w, E)} dE \tag{1.6}$$

where E_{in} is the ion's incoming energy, E_{out} is its outgoing energy, I_w is the mean excitation energy of water and $S(I_w, E)$ is the stopping power of water for charged particles of energy E (Hurley et al. [2012]). Finally, the reconstructed quantity measures the ability of the imaged object to slow ions down relative to a reference material, typically water: this is usually referred to as Relative Stopping Power (RSP), that is the charged particle stopping power relative to that of water (Krah et al. [2022]).

Since software involved in the in dose distribution planning for ion therapy rely on range calculation, it becomes clear that the photon attenuation coefficients from the X-ray CT has somehow to be converted to RSP so to match the physical properties of the therapeutic radiation. This is usually done through means of semi-empirical calibration curves based on phantoms of tissue equivalent materials. However, due to different chemical composition, mass density and mean ionization energy of real tissues, uncertainties are deeply embedded into this conversion mechanism: the relation between HU and RSP ceases to be one-to-one and multiple values of RSP can be associated to the same HU value (Schneider and Pedroni [1995]). Specifically, the magnitude of these inaccuracies was estimated to be ranging from 1.1% in soft tissue to 1.8% in cortical bone, up to 5% in the lungs, as shown in Fig. 1.8. In general, the average error introduced by the HU-RSP calibration can be estimated to be around 3% (Gianoli et al. [2020]).

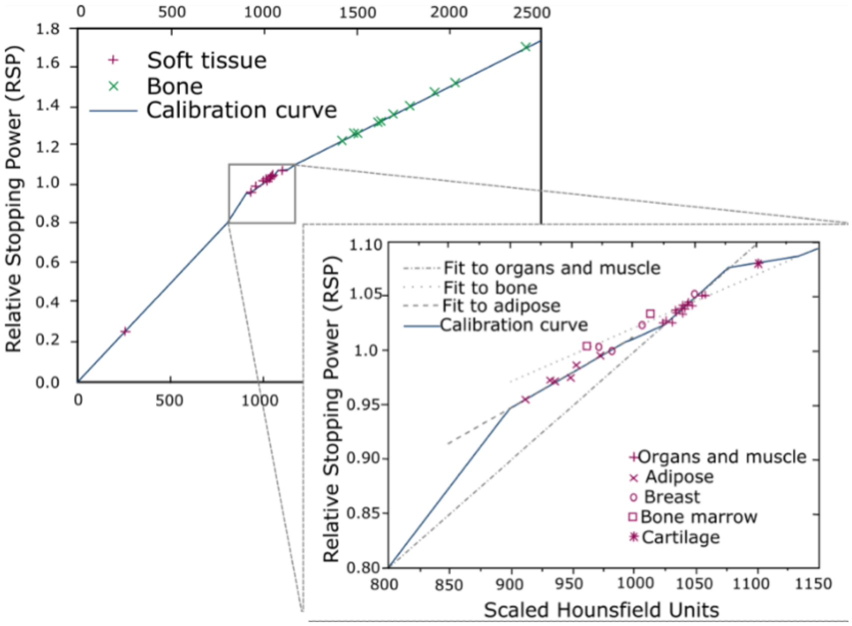

Figure 1.8: Calibration curve for conversion from Hounsfield Units to Relative Stopping Power. The calibration curve is defined by piecewise linear fitting of the theoretical HU, parameterized based on experimental HU of tissue equivalent materials with known elemental composition, and theoretical RSP, calculated according to the Bethe-Bloch model.

Because of these limitations, it stands to reason that a single imaging modality able to bypass the need for a translation from a physical domain to the other is quite appealing. Transmission ion imaging is, in fact, the native imaging paradigm in ion beam therapy because it is already representative of charged particle behavior.

1.6 Ion Imaging

The goal of ion CT is to reconstruct a 3D RSP distribution inside an imaged object so to define a more accurate dose deposition with treatment beams.

Ion imaging is implemented by measuring the residual energy of the ions that is then converted to WET. By introducing an estimation of the ion trajectory into reconstruction algorithms, it is eventually possible to obtain a map of RSP for the object of interest (Oria et al. [2018]). The mathematical ground for iCT reconstruction lies in the relation between RSP and WET, where WET can be modeled as the line integral of the object's RSP along a certain concept of ion trajectory L that depends on the detection system (Meyer et al. [2021]):

$$WET = \int_L RSP(\vec{r})dl \tag{1.7}$$

The solution to the inverse problem of equation (1.7) allows to directly retrieve the spatial distribution of RSP through dedicated reconstruction algorithms, thus effectively eluding the inaccuracies intrinsically integrated in the calibration process from HU to RSP.

The same state-of-the-art reconstruction techniques usually employed in medical imaging (especially X-ray CT) can also be applied to solve the inverse problem in iCT image reconstruction. Analytical approaches, i.e. the Filtered Back Projection (FPB) would provide for a fast and relatively easy solution to the problem: the measured data are projected back to the imaging space while applying a filter (usually a ramp function) so to equalize the contributions of all frequencies in the Fourier Transform domain (Desai and Kulkarni [2010]). Though computationally efficient, FBP-reconstructed CT images suffer greatly from noise amplification and streak artifacts introduced by the filtering step (Schofield et al. [2020]).
While more demanding in terms of memory consumption and computation times, iterative reconstruction (IR) algorithms are less influenced by noise and, with ever increasing computing advancements, they are slowly replacing FBP as a reconstruction method. IR assumes the slice or section of the object to be imaged as an array of unknowns to be determined by defining a set of algebraic equations pointing to the measured data. All algorithms within this family work to iteratively improve an estimate of the object by minimizing the discrepancy between the measured projections and those simulated from the current image. Since ion imaging often needs to address large, sparse, sometimes inconsistent datasets, IR may offer a better handle than analytical approaches. Furthermore, the FBP is not as suited as iterative algorithms for modeling the complex physics underlying ion-matter interactions and the reconstruction of the particle's path (discussed further in section 1.7), resulting in a lower spatial resolution (Penfold and Censor [2015]).
Finally, recent years have also witnessed the engaging of data-driven, deep-learning-based tomographic image reconstruction, i.e. *deep learning reconstruction* (DLR). Relying on Ground Truth data to train the network, the objective in deep tomographic reconstruction is to find the parameters of the mapping function that infers the Ground Truth itself based on supervised predictions. Despite it being an emerging field, DLR aims to improve image quality and eventually reduce radiation dose, laying the grounds for new reconstruction approaches able to outperform current techniques and redefine the future landscape of image-guided radiotherapy (Koetzier et al. [2023]).

1.7 Statistical limitations in ion imaging

Taking into account the different nature of the ion-matter interactions compared to photons has deeper implications than the single need for calibration: poor spatial resolution and presence of artifacts are rooted problems when it comes to the quality of the reconstructed ion image. Particularly, these limitations are due to two contributing processes: Multiple Coulomb Scattering (MCS) and energy-loss straggling (Collins-Fekete et al. [2020]).

1.7.1 Energy-loss straggling

Variations in the energy loss are mainly due to the ion's interactions with the atomic electrons of the traversed medium (Poludniowski et al. [2015]). Whereas the stopping power expresses a mean value, the energy transfer itself happens over many individual inelastic collisions between the ions and the sample's elements with the amount of energy lost varying due to the stochastic nature of the process (Lohmann et al. [2023]). Since this effect is cumulative along the ion's entire path, the exact energy loss experienced by the single ion suffers from statistical variations (Chibani [2002]) leading to a spread in the energy distribution that can be characterized by its variance called *energy-loss straggling*. To better understand this process, it is useful to refer to the probability distribution function (PDF) describing the distribution of energy loss Δ of an ion traversing an absorber of thickness x (Newhauser and Zhang [2015]). Focusing specifically on fast and light ions such protons, the moments of the energy PDF can be schematically compacted as:

$$M_n = \int_0^{\Delta_1^{max}} \Delta^n f(\Delta) d\Delta \tag{1.8}$$

where f(Δ)dΔ is the probability of energy loss in a infinitesimal range $[\Delta,\ \Delta + d\Delta]$ and Δ_1^{max} is the highest energy loss possible in a single encounter. The order-zero moment is the total collision cross section, while the first-order moment represents the mean value of the energy distribution that is, as a matter of fact, the ion's stopping power. The second moment is thus the variance of the straggling distribution as expressed in equation (1.9):

$$M_2(\Delta) = \sigma_\Delta^2 = \overline{v} \int_0^{\Delta_1^{max}} \Delta^2 f(\Delta) d\Delta \tag{1.9}$$

where $\overline{v}$ is the average number of primary collisions that a proton undergoes during its traversal. Different theories have been proposed to assess the straggling of the energy distribution based on material's constitution experienced by the ion. The possible effects of the absorber's nature are summarized in Fig. 1.9, where the PDFs are shaped according to the respective model.

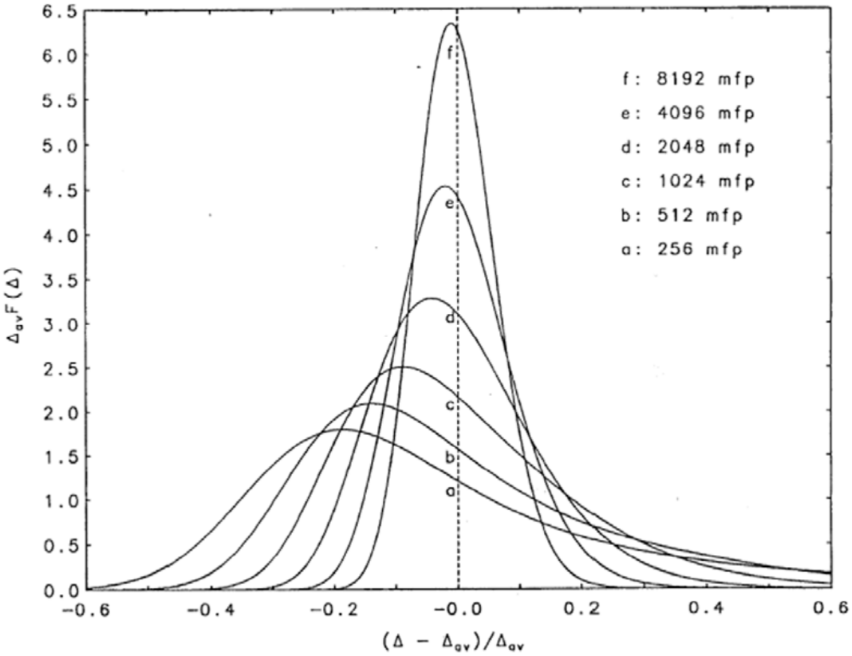

Figure 1.9: Energy loss straggling distributions in water absorber. Energy loss PDFs vary according to the thickness of the traversed material, here expressed in units of mean free path (mfp). The energy loss for a single event is represented as a fraction of the average energy loss over the entire thickness of the absorber. For thin absorbers, the distribution is asymmetric with a long tail and can be modeled either by Landau's or Vavilov's theory. As the material gets thicker, the distribution follows a Gaussian trend accurately expressed by Bohr's theory (Newhauser and Zhang [2015]).

Bohr's theory (Bohr [1915]) applies to thick materials and postulates that the straggling behavior of energy follows a Gaussian PDF described as:

$$f(\Delta)d\Delta = \frac{1}{\sigma_{\Delta}\sqrt{2\pi}} e^{\frac{-(\Delta-\overline{\Delta})^2}{2\sigma_{\Delta}^2}} \tag{1.10}$$

where the variance amounts to

$$\sigma_{\Delta}^2 = 2\pi r_e^2 m_e c^2 z^2 \frac{NZ}{\beta^2} \Delta_1^{max} \rho x \tag{1.11}$$

where ρ is the mass density and x is the absorber's thickness. The intrinsic assumptions backing this theory up are that the material is thick enough so many individual collisions take place, that the particle's velocity does not decrease that much along its path and that the electrons are unbound in the absorber (Newhauser and Zhang [2015]).

Landau's theory (Landau [1944]) adapts to thin absorbers assuming that in this kind of medium there are fewer collisions but higher energy transfers. This leads to an asymmetric distribution of the energy PDF known as Landau distribution:

$$f(\Delta, \rho x)d\Delta = \frac{1}{\xi} \phi_{\mathrm{L}}(\lambda_{\mathrm{L}}) \tag{1.12}$$

where ξ is an approximation of the mean energy loss and $\phi_L(\lambda_L)$ essentially represents the deviation from it. Finally, Vavilov (Vavilov [1957]) extended Landau's theory to absorbers of intermediate thickness relying on the exact expression for Δ_1.
In clinical-like scenarios, a good first-order approximation neglects range fluctuations relying only on the average energy loss provided by Bethe-Bloch's formula (Newhauser and Zhang [2015]). In fact, models such Bohr's or Landau's are based on specific simplifications and assumptions and may fail in capturing the complexity of the actual process: for instance, the diverse range of densities and compositions in human tissues not always can be accounted for in theoretical models. Nevertheless, since the combined effect of all these small variations in the energy loss plays an important role in shaping the Bragg peak curve, this phenomenon can heavily limit the overall energy resolution for ion beam analysis (Lohmann et al. [2023]) and can also induce inaccuracies in determining the exact position of the interactions.
Lastly, it is interesting to note that these particular conditions are mainly reflected in light ions (protons); energy straggling becomes instead less meaningful when it comes to heavy ions (Jäkel [2006]). Due to their larger mass, heavy ions move more slowly when traversing a medium resulting in fewer high-energy impacts with surrounding electrons. Because of this, they tend to lose energy more gradually than light ions and are thus less sensible to energy loss fluctuations.

1.7.2 Multiple Coulomb Scattering

The other issue lies instead in the elastic and semi-elastic interactions of the ions with the Coulomb field of the nuclei of the surrounding matter: this effect is known as *Multiple Coulomb Scattering* (Fig. 1.10) and specifically refers those deflections that a charged particle undergoes when traversing a certain material (Schneider and Pedroni [1995]).

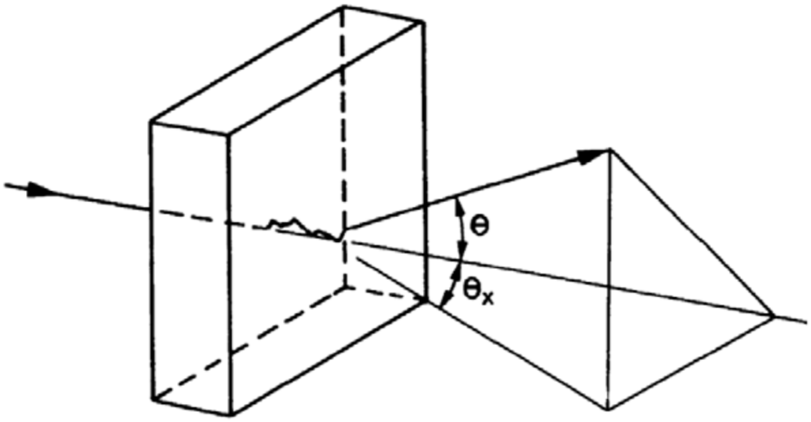

Figure 1.10: Schematic representation of Multiple Coulomb Scattering. The ion trajectory is deflected from its nominal path due to repulsive forces from positively-charged nuclei of the surrounding medium. Though the energy lost due to this process is not significant, the deviations in the trajectory may add up to severe implications in terms of therapeutic and diagnostic procedures (Newhauser and Zhang [2015]).

The cumulative effect of MCS across the entire depth can lead to a significant deviation from the initial trajectory. Because of this, the projectile's original path cannot be entirely recovered and reconstructed, once again degrading the spatial resolution of the ion image (Burker et al. [2020]).
The scattering power T can be employed as a parameter to somehow quantify the amount of scattering an ion beam experiences when traversing a certain object:

$$T = \frac{d < \theta^2 >}{dx} \tag{1.13}$$

where $< \theta^2 >$ is the mean squared scattering angle and x is the material's thickness. When talking about MCS, the number of scattering events for an ion across a given depth (considering a minimum relevant length scale of 1 mm in proton radiotherapy) is clearly high, usually more than 20 (Newhauser and Zhang [2015]). A thorough modeling of the overall scattering phenomenon is difficult to achieve, but one of the most acclaimed is Molière's theory (Molière [1948]). The probability density function of the scattering angles themselves is given as a series expansion by:

$$P(\theta)d(\Omega) = \eta d\eta\left(2e^{-\eta^2} + \frac{F_1(\eta)}{B} + \frac{F_2(\eta)}{B^2} + ...\right) \tag{1.14}$$

where $d\Omega$ is the differential solid angle into which the projectiles are scattered, $\eta = \frac{\theta}{\theta_1\sqrt{B}}$ is a dimensionless parameter relying on the scattering angle θ and a characteristic angle θ_1 which depends on the properties of the material and the incident particles; moreover, the functions F_1, F_2, ..., are specific corrective terms dependent on η, while B is a parameter related to the material's thickness and properties.
While extremely comprehensive, Molière's theory is also quite difficult and computational implementations may be too time consuming: for this reason, simpler approaches have been proposed. Widely recognized is the Gaussian scattering model proposed by Fermi-Eyges (Eyges [1948]) which computed a joint distribution between projected angle and spatial displacement, and the correlation between them. According to this approach, the probability of finding a particle across a certain depth z in intervals $[x, x+dx]$ and $[\theta_x, \theta_x+d\theta_x]$ is expressed in eq. 1.15 as [1]:

$$P(x,\theta_x)dxd\theta_x = \frac{1}{2\pi\sqrt{A_0A_2 - A_1^2}}e^{-\frac{1}{2}\frac{A_0x^2-2A_1x\theta_x+A_2\theta_x^2}{A_0A_2-A_1^2}}dxd\theta_x \tag{1.15}$$

where the parameters A_n represent the first three moments of the distribution:

$$A_n = \int_0^z (z-z')^n T(z')dz' \tag{1.16}$$

with *T(z)* being the scattering power.
Afterwards, Highland [1975] proposed a revisited approach to Moliere's theory in order to determine the scattering angle θ in a slab of thickness l assuming it was thin enough not to significantly slow the ion down:

$$\theta = 14.1MeV\frac{Z}{pv}\sqrt{\frac{l}{X_0}}\left(1 + \frac{1}{9}log_{10}\frac{l}{X_0}\right) \tag{1.17}$$

where X_0 is the radiation length, p is the initial momentum of the projectile, while v is its initial velocity and Z the atomic charge number.

Later, Gottschalk et al. [1993] adapted Highland's work to thick absorbers. When introducing also the corrective terms proposed by Lynch and Dahl [1991], the first three moments of the projected angular-spatial distribution, assuming an entry depth at z_1 and exit depth at z_2, can be generalized as:

$$A_n(z_1,z_2) = E_0^2\left[1 + 0.0038ln\left(\frac{z_2-z_1}{X_0}\right)\right]^2\int_{z_1}^{z_2}\left(\frac{Z}{pv}\right)^2\frac{(z_1-z')^n}{X_0}dz' \tag{1.18}$$

with $E_0 = 14.1MeV$.
Despite offering a solution to somehow quantify the elastic scattering events, all these approaches hide their most daunting weakness in the small-angle approximation. The Central Limit Theorem poses that the contribution of different *small* displacements tends to a Gaussian distribution. However, rare large-angle scattering events may occur when an ion passes too close to the nucleus itself or has a direct collision with it. Whereas small-angle scatterings collectively follow a Gaussian curve, the large-angle ones lead to a heavier falloff in the tail that does not fit a Gaussian distribution (Embriaco et al. [2017]). Through his higher-order terms, Molière was effectively able to unify the Gaussian and the non-Gaussian region of the scattering angle

---

[1] This formula and the following are adapted from the work of Embriaco et al. [2017] and refers to a particle moving in a 3D homogeneous material with nominal trajectory along the z-axis, direction described by the angles θ_x and θ_y and kinetic energy E.

PDF in a compact and precise way. While the small-angle approximation is true for thin or low-atomic-number materials, ion imaging may happen to deal with matter that could be somewhat dense or have high atomic number, resulting in underestimations of the particle's path and, consequently, lower accuracy in the reconstructed image.
Though it is factual that MCS is a limiting factor, heavy ions (such as carbon ions or oxygen ions) are less susceptible to changes in their trajectory than protons. Since protons have a lower mass than heavier ions, they will also have lower inertia and thus they will more easily be deflected by Coulomb forces from the atomic nuclei. While protons are currently the main focus for charged particle therapy and imaging, heavier ions may very well represent a suitable alternative when minimizing scattering phenomena and blurring effects is crucial for therapeutic purposes.

1.8 Detector configurations for ion imaging

If ion beams have enough energy to traverse the patient, an ion radiography can be obtained by directing those ions towards suitable sensors. When it comes to ion imaging, two different types of detectors can be adopted, shaping the data acquisition and the reconstruction process.

1.8.1 List-mode detector configuration

List-mode ion imaging has gained growing interest over the years as a methodology to improve range accuracy in treatment planning: in fact, this family of detectors allows for one-by-one ion tracking, leading to a remarkably precise estimation of the ion's trajectory which is quite beneficial to the spatial resolution of the reconstructed image (Sølie et al. [2020]). Ion-tracking systems usually consist of a number (typically from one to four) of position-sensitive detectors (PSD) and a residual energy-range detector (RERD). As shown in Fig 1.11, PSDs are positioned upstream and downstream of the object to be imaged, so to provide information about the projectile's scattering power through position and direction measurements at the entry point and the exit point; the RERD is located downstream of the tracking system so to retrieve data about the ions' energy at the end of their trajectory (i.e. the stopping power). The energy detector can rely on a single absorption and detection block as well as multiple absorption layers interleaved by detection layers measuring multiple energy losses.

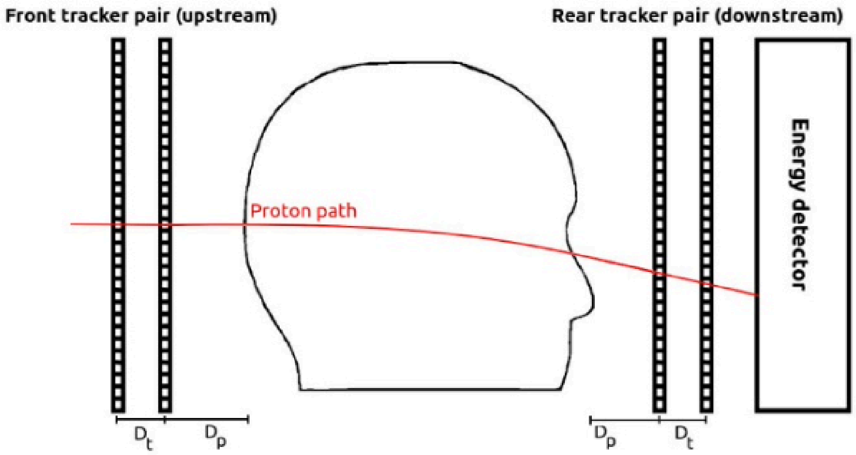

Figure 1.11: Ion-tracking system. A schematic of the ideal ion-traking system for ion radiography/ion CT. D_t is the distance between the position-sensitive trackers in each tracker pair and D_p is the distance between the inner tracking plane and the object to be imaged. (Sølie et al. [2020])

A significant amount of studies have been published over the last decades marking the emergence of list-mode technology. Poludniowski et al. [2015] reported a summary of recent developments in this field produced by different research groups.

In 1999, an experiment conducted at the PSI Center for Nuclear Engineering and Sciences culminated in one of the first proton-tracking device. The system consisted in two PSDs (one before and one after the patient). The tracking units were scintillating fibre hodoscopes (Sci-Fis) consisting of two orthogonal planes of 2×2 mm^2 plastic fibres made of plastic scintillators and coupled each to a channel of a photomultiplier tube. As for the RERD, it was a range telescope consisting of 64 closely packed and optically isolated scintillator tiles of 3-mm thickness. The light from each tile was collected by a wavelength-shifting fibre coupled to a photomultiplier channel. The fibres were designed to gather scintillator emissions and efficiently transfer light quanta to the photon sensor at a wavelength matched to the spectral sensitivity. Although suitable for proton CT (pCT), no indication of the system being used for this purpose is recorded and it was used only to produce a proton radiography of a live canine head.

In 2007, a new initiative was carried out by an Italian group for the development of a PRoton IMAging (PRIMA) device. Four silicon strip detectors (SiSDs) were used as PSDs, while a crystal calorimeter was adapted to RERD. Particularly, the RERD was based on four $YAG : Ce$ crystals arranged in 2×2 array and coupled to photodiodes. The reconstructed CT images promised great potential for spatial resolution, with a 0.9 mm full-width half-maximum (FWHM).

In 2008, the Tera Foundation (Novara, Italy) was funded by the Italian National Centre for Oncological Hadron Therapy (CNAO) (Pavia, Italy) to develop a series of devices concerning the Advanced Quality Assurance in Hadron Therapy project. The trackers were realized by a three-foil gas electron multipliers with a read out pitch of 400 μm. Since the tracking system included only two PSDs downstream of the patient, it was only possible to obtain information on the exit direction and position of the protons. The RERD was a stack of 48 plastic scintillators with an area of 30×30 cm^2, each connected to a silicon photomultiplier.

In 2010, a collaboration involving Loma Linda University (LLU), University of California Santa Cruz (UCSC) and Northen Illinois University (NIU) revealed the first results of a new prototype. Four PSDs were adopted in total, two before the patient and two after the patient, allowing for the tracking of the particles' ingoing and outgoing position and direction. Each PSD consisted of two silicon strip detectors with a sensitive area of 9×9 cm^2. As for the RERD, this was a calorimeter where 18 $CsI : Tl$ crystals were arranged in a 3×6 matrix; the light was collected by a photodiode paired to each crystal. Though the scanning times took several hours due to the dead-time of the calorimeter, the acquired CT images had a promising accuracy.

In 2011, LLU, UCSC and California State University, San Bernardino (CSUSB) started working on a second generation tracking system. Four PSDs identical to the first generation ones were used. The main difference resided in the energy detector: this consisted of a stack of five fast plastic scintillators read out by photomultiplier tubes, proving a more accurate measure than the first-generation calorimeter. Early results showed a good RSP accuracy and high image quality.

Finally, a more recent implementation led to the development of 4D tracking detectors, allowing for a simultaneous acquisition of the ion's position and time. In fact, nowadays most tracking systems are still struggling with fast acquisition rates, not at all comparable with standard CT procedures (<1 min). A rather intriguing solution is the introduction of iCT scanners using 4D-tracking for both ion's trajectory estimation and residual energy determination through time-of-flight (TOF) measurements. This fast tracking device would certainly help in improving the data rate performances of the imaging system and the spatial resolution of the reconstructed image, but it would also ideally allow for the integration of the TOF through the patient into the imaging process, such as by using the energy loss within the patient along with the corresponding TOF to filter out nuclear events (Ulrich-Pur et al. [2022]).

1.8.2 Integration-mode detector configuration

Integration-mode detector configuration allows to retrieve stopping power data through an energy detector consisting either in a single absorption and detection block or in multiple absorption layers interleaved by detection layers measuring multiple energy losses (i.e. the Bragg Peak signal). Nevertheless, ion-integrating systems lack the capability of resolving the single ions because the tracking component is not embedded in this family of detectors. Since the ion's trajectory can not be modeled, ion-integrating systems operate based on the pencil beam's mean trajectory concept (Fig. 1.12). Thus, the integrated signal is exploited to obtain the RSP volumetric information of the object traversed by the pencil beam (Meyer et al. [2017]).

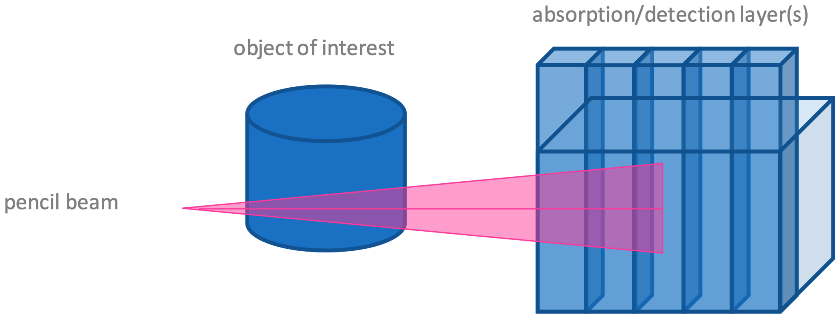

Figure 1.12: Ion-integrating system.The integration-mode detector configuration, intended for pencil beams, is composed by only the absorption detector without tracking system.

Different detector technologies have been proposed in the field of ion-integrating radiography (Poludniowski et al. [2015]).
In the early 2000s, a gadolinium oxysulfide scintillator screen was coupled to a charge-coupled device and the signal then converted to WET. Using the Felkamp algorithm for tomographic reconstruction, the image showed recognizable phantom slices, although consistent edge artifacts were still present at interfaces between materials owing to MCS. Muraishi et al. [2009] reproduced the same system concept exploiting heavier ions (^4H and ^{12}C), observing a superior image quality.
In 2012, a carbon-ion radiography was produced relying on a gadolinium oxysulfide scintillator coupled to an amorphous silicon matrix array. A filtered back projection reconstruction led to a RSP accuracy of 1% and a spatial resolution dominated by pixel size (0.8 mm).
In 2013, a prototype 2D diode-array was tested at Massachusetts General Hospital (Boston, MA): the detector had a 12 cm field size and contained 249 semi-conductor diodes in an octagonal array with a 7 mm diagonaltransvect pitch. Iterative algorithms provided identifiable CT phantom images, though the spatial resolution was low due to severe scattering events and the RSP accuracy was unacceptable for treatment planning purposes.
In 2013, the GSI Helmholtz Center for Heavy Ion Research developed an efficient prototype for carbon-ion based transmission imaging in collaboration with the Heidelberg Ion Beam Therapy Center (HIT) (Rinaldi et al. [2013]). The integrating-system detector consisted in 61 ionization chambers (IC) interleaved with absorbers of homogeneous thickness (3 mm polymethyl methacrylate) working as range degrader. The chambers had a 300 × 300 mm^2 area allowing for a minimum scanning field size of 250 × 250 mm^2. Each IC section bears a capacitance of 130 pF. Aluminized kapton foils of 20 μm were exploited as collecting electrodes supplied with a voltage of 1000 V. The system was able to handle energies from 89 to 430 MeV/u, leading to ranges between 20 to 300 mm in water. The read-out electronics consisted in two I3200 thirty-two channel digital electrometers, synchronized and buffered by an A500 module as real-time controller and able convert the electric current signals into readable measurements through a gated integration mechanism. The set-up was tested generating a 30 × 30 mm^2 field of monoenergetic ^{12}C ions at 270.55 MeV/u and beam spot size of 4.1 mm; the raster scanning system was set to move in 1 mm steps along the x and y directions, while the field was applied without any object placed in front of the IC stack. The proposed detector showed a high enough signal quality even with the smallest tested number of beam particles, which amounted to 100, corresponding to a 0.4 mGy dose for each delivered raster point. This result is encouraging for integration-mode transmission ion imaging as a low dose imaging technique, although the ratio of the number of beam particles measured by the ionization chambers relative to the total number of beam particles delivered to the patient should be optimized by minimizing the dead times of the read-out electronics.
Though it clearly appears that list-mode detectors are capable of achieving far greater accuracy and resolution, it is also worth noticing the scarcity of attempts towards the development of simpler detector models (Telsemeyer et al. [2012]). In fact, ion-tracking devices generate a large amount of data, since they record individual particle events, resulting in long acquisition times as well as storage and processing issues. Integration-mode detectors could provide for a significant reduction in terms of costs and data collection and analysis, enabling a provisional step in the clinical standardization of ion imaging.

Chapter 2

Materials and methods

2.1 Monte Carlo simulation framework

X-ray CT images of head-and-neck patients being treated with intensity modulated photon therapy at the Department of Radiation Oncology of the university hospital of the Ludwig-Maximilians-Universität München (Munich, Germany) were collected. The size of the images was 512 × 512 pixels in the imaging plane, then cropped to 256×256 pixels for storage purposes, and with a slab of slices varying depending on the specific patient, also cropped so to enclose the anatomy of interest. The voxel size was $1.074\times1.074\times3\ mm^3$. Relying on the GPU-based FRED Monte Carlo platform (Schiavi et al. [2017]), all iCT and iRad simulations aimed to mimic pencil beam scanning, where the individual beam is implemented as if deflected according to a certain angle θ, based on the empirically validated phase space files of the Heidelberg Ion Beam Therapy Center (Tessonier et al. [2016]). This configuration allows to keep track on charge, mass, energy, coordinates and direction cosines for each particle crossing the Beam Application and Monitoring System (BAMS) located downstream of the scanning magnets. Following this model (Fig. 2.1), the radiation transport begins when the beamline leaves the monitoring system positioned 112.6 cm upstream of the isocenter. After that, the beam is focused on a simulated raster point grid with 1.074 mm and 3 mm spacing in the horizontal and vertical direction respectively, in accordance with the resolution of the X-ray CT image. The positions to be irradiated are selected according to an irradiation plan defined prior to the simulation.

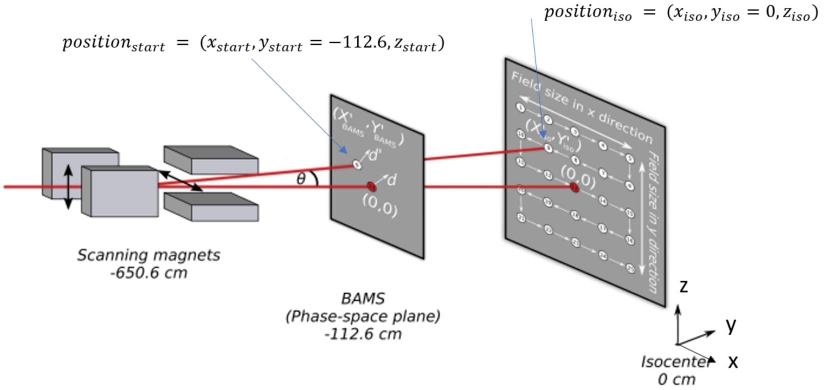

Figure 2.1: Schematic representation of the FRED simulation environment. Upon modulation of the magnets, the beam is focused on a raster grid with a certain displacement from the isocenter on the x-z plane according to the scanning angle. The y coordinate marks the nominal direction of the ion beam (figure adapted from Meyer [2019]).

According to this scenario, the X-ray CT DICOM files were imported in FRED so to simulate a tomographic acquisition with a total of 180 projections starting from 0° and spaced by 1°. For each projection angle, the number of recorded pencil beams was adapted in accordance to the interactions of the ions with the traversed anatomy, represented by the X-ray CT voxelized geometry checked by the Monte Carlo code at each step: particularly, trajectories passing through only air were excluded. Each pencil beam contained 400 primary protons. In all cases, the same nominal energy E222 of the Heidelberg Ion Beam Therapy Center (around 200 MeV) was used, corresponding to a proton range of roughly 26 cm in water.

Adding to this implementation, two scoring planes were also simulated to record position, direction and energy for every individual ion. They were located before and after the patient at 13.7472 cm from the isocenter: considering the physical dimension of the image, this distance was calibrated so to have the two planes right at the edges of the imaging space. Relying on this approach, it was possible to model an ideal proton-tracking system; on the other hand, only energy scoring was taken into account when modeling proton-integrating systems.

Finally, the X-ray CT data framed in the Monte Carlo environment allowed for the extraction of a clinical-like (monotonic and bijective) HU-RSP baseline calibration curve based on 27 linear segments, according to the standard set by Schneider et al. [2000] and extended by Parodi et al. [2007]. The start and end points of these segments are randomly deviated by applying to the respective RSP value either a multiplicative error up to ±5% for HU smaller or equal to 0 or an additive error up to ±0.05% for HU greater than 0. This allows for the determination of slice-specific calibration curves, thus defining slice-specific Ground Truth images. The same principles were adopted to retrieve the correlation between HU and both density and radiation length.

2.2 Tomographic iCT reconstruction

In order to reconstruct the 3D RSP map, the continuous model in equation (1.7) was fully discretized. Under this paradigm, the problem can be easily expressed in matrix notation as:

$$Ax = b \tag{2.1}$$

where A is the $I \times J$ system matrix containing elements a_{ij} describing the intersection volume of the i-th ion trajectory with the j-th voxel, x a is J-dimensional vector containing an estimation of the imaged RSP and b is a I-dimensional vector containing the measured WET (Penfold and Censor [2015]). Following this, the discretized model linking RSP and WET is given by:

$$WET_i = \sum_j a_{ij} RSP_j \tag{2.2}$$

Though this is the general framework used to obtain the RSP images, the definition and implementation of the system matrix and the ion radiography are meticulously declined in order to represent the nature of the detecting system and what they actually bring to tomographic reconstruction. In fact, the reconstruction algorithm must take into account the specific needs and issues of the actual sensor configuration so to adapt to the physics and geometry of the problem and provide for the highest possible quality in the ion image. Furthermore, this work particularly aims to address those resolution limitations owing to elastic collisions of the protons. The MCS is embedded in the trajectory estimation according to different models either in list-mode or integration-mode: while this does not result in a different role of the system matrix in the reconstruction workflow, it is important to keep in mind its true physical implications throughout this entire work.

2.3 Proton-tracking systems

List-mode detectors are able to track each ion individually, so the system matrix will contain the single ion trajectories belonging to a certain pencil beam. This also means that each ion is associated to a specific WET value.

The estimation of the ion trajectory is a crucial step, especially for the introduction of the scattering model. The approach followed in this study was the *Most Likely Path* (MLP) algorithm as proposed by Schulte et al. [2008] for uniform water and then extended by Collins-Fekete et al. [2017a] in non-uniform materials.

2.3.1 Most likely path in homogeneous medium

The MLP formalism combines the generalized MCS model from Eyges [1948] with Bayesian statistics in order to compute the most probable path followed by the proton across a uniform absorbing material. Considering the physical reference system associated to the pCT detectors, it is possible to define a u-axis orthogonal to the detector planes along which the proton travels according to its nominal direction, and a t and v axis parallel to the detector planes instead. At any given depth u_i across the medium, the proton's position is defined by the lateral t_i and vertical v_i coordinate, while its direction is given by the θ_i and ϕ_i angles relative to the u-axis. In a homogeneous material, scattering in the two directions perpendicular to the nominal one can be considered as independent statistical problems, so it is possible to derive the MLP formula in just one plane, bearing the same outcome in the other. Assuming to decline the formalism in the u-t plane (Fig. 2.2), the proton's position t_1 and direction θ_1 at a certain depth u_1 is expressed by the 2D parameter vector:

$$\mathbf{y1} = \begin{pmatrix} t_1 \\ \theta_1 \end{pmatrix} \tag{2.3}$$

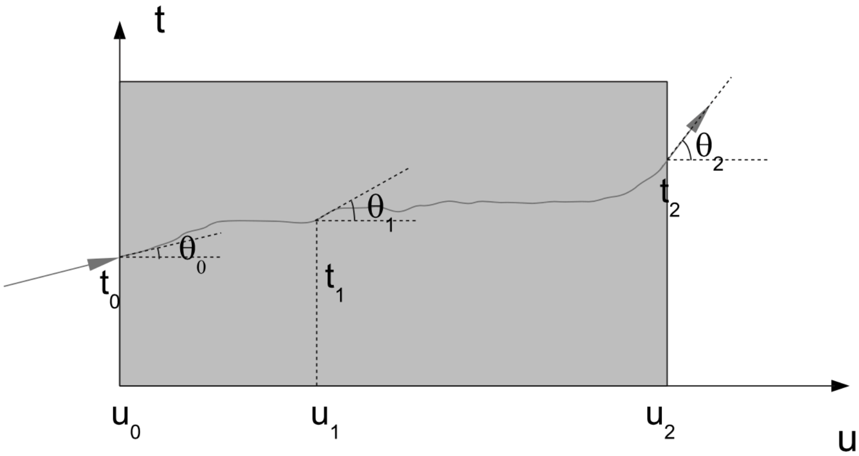

Figure 2.2: Model of a proton undergoing scattering events in the *u-t* detector plane. The MLP formalism employs Bayesian processes to estimate the proton's deviation by keeping track of its position and direction at any given depth inside the traversed object (Schulte et al. [2008]).

The algorithm aims to determine the most probable state vector at a specific depth within a Bayesian framework and it is computed from four input parameters: the proton's position and direction at the entry point (t_0 and θ_0) and at the exit point (t_2 and θ_2) of the image reconstruction volume. Bayesian terminology defines a prior likelihood $\mathcal{L}(y_1|y_0)$ of finding the vector y_1 at a depth u_1 given knowledge on its initial state y_0 at u_0, and a posterior likelihood $\mathcal{L}(y_1|y_2)$ of finding the vector y_1 in u_1 given y_2 in u_2. Prior and posterior likelihood are linked by the likelihood $\mathcal{L}(y_2|y_1)$ of observing the particle with exit state y_2 given y_1 at depth u_1 by the following relation:

$$\mathcal{L}(y_1|y_2) = \mathcal{L}(y_2|y_1)\mathcal{L}(y_1|y_0) \tag{2.4}$$

According to the generalized Fermi-Eyges theory of MCS (see section 1.7.2), the likelihood functions can be expressed as a bivariate Gaussian. Assuming a proton entering the reconstruction volume at u_0 with zero displacement and angle, the prior likelihood is represented by:

$$\mathcal{L}\left(\boldsymbol{y_1}|\mathbf{y_0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}\right) = exp\left(-\frac{1}{2} y_1^T \Sigma_1^{-1} y_1\right) \tag{2.5}$$

where Σ_1 is known as scattering matrix and contains the variances and covariances of t_1 and θ_1 acquired between u_0 and u_1:

$$\mathbf{\Sigma_1} = \begin{pmatrix} \sigma^2_{t_1} & \sigma^2_{t_1\theta_1} \\ \sigma^2_{t_1\theta_1} & \sigma^2_{\theta_1} \end{pmatrix} \tag{2.6}$$

The elements of the scattering matrix Σ_1 can be computed as the first three Eyges moments following the approach of Gottschalk et al. [1993] adjusted through the corrective terms introduced by Lynch and Dahl [1991] (see equation 1.18):

$$\sigma^2_{t_1}(u_0, u_1) = E_0^2\left(1 + 0.0038ln\frac{u_1 - u_0}{X_0}\right)^2 \int_{u_0}^{u_1} \frac{(u_1 - u)^2}{p^2(u)v^2(u)}\frac{du}{X_0} \tag{2.7}$$

$$\sigma^2_{\theta_1}(u_0, u_1) = E_0^2\left(1 + 0.0038ln\frac{u_1 - u_0}{X_0}\right)^2 \int_{u_0}^{u_1} \frac{1}{p^2(u)v^2(u)}\frac{du}{X_0} \tag{2.8}$$

$$\sigma^2_{t_1\theta_1}(u_0, u_1) = E_0^2\left(1 + 0.0038ln\frac{u_1 - u_0}{X_0}\right)^2 \int_{u_0}^{u_1} \frac{u_1 - u}{p^2(u)v^2(u)}\frac{du}{X_0} \tag{2.9}$$

where $v^2(u)$ is the squared velocity of the proton relative to the speed of light c, $p^2(u)$ is its momentum at depth u, $E_0 = \frac{13.6\ MeV}{c}$ is an empirical constant representing the proton's rest energy and $X_0 = 36.1\ cm$ is water's radiation length.
To compute the likelihood of finding a particle at a depth u_1 with state vector y_1, it is necessary to implement a *shear mapping* (or *transvection*) of the measurement points towards the reconstruction point at u_1 (Fig. 2.3).

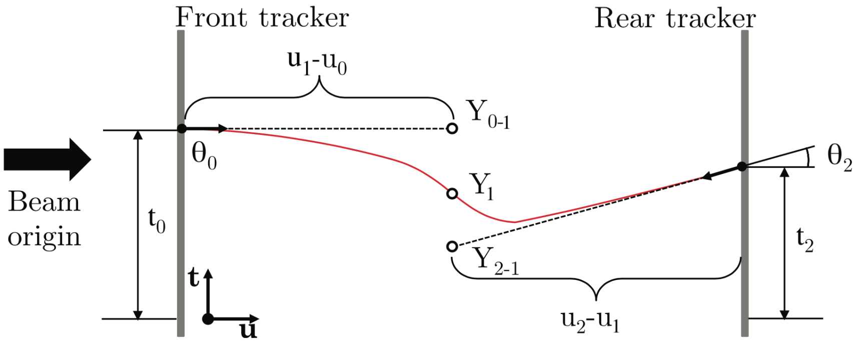

Figure 2.3: Transvection of the local reference coordinate system of the proton. The observed entry and exit point are propagated according to location and orientation of the proton path at the reconstruction depth u_1 (figure adapted from Collins-Fekete et al. [2017b]).

As it is reasonable to assume relatively small entry angles for ion imaging purposes, the small-angle approximation ($sin\theta \approx \theta$) is therefore adopted. Operating under these conditions, the introduction of the transvection matrix

$$\mathbf{R_0} = \begin{pmatrix} 1 & u_1 - u_0 \\ 0 & 1 \end{pmatrix} \tag{2.10}$$

allows to define the new state vector as:

$$y_1' = y_1 - R_0 y_0 \tag{2.11}$$

By following the same steps, it is possible to draw equivalent conclusions for the exit state vector y_2 at depth u_2 given y_1 at depth u_1. The likelihoods in the second half of equation 2.4 can be therefore re-interpreted as:

$$\mathcal{L}(y_1|y_0) = exp\left(-\frac{1}{2}(y_1^T - y_0^T R_0^T)\Sigma_1^{-1}(y_1 - R_0 y_0)\right) \tag{2.12}$$

$$\mathcal{L}(y_2|y_1) = exp\left(-\frac{1}{2}(y_2^T - y_1^T R_1^T)\Sigma_2^{-1}(y_2 - R_1 y_1)\right) \tag{2.13}$$

The product of these two likelihoods as expressed in equation 2.4 allows to determine the posterior likelihood as:

$$\mathcal{L}(y_1|y_2) = exp\left(-\frac{1}{2}((y_1^T - y_0^T R_0^T)\Sigma_1^{-1}(y_1 - R_0 y_0) + (y_2^T - y_1^T R_1^T)\Sigma_2^{-1}(y_2 - R_1 y_1))\right) \equiv exp(-\chi^2) \tag{2.14}$$

At this point, the most probable parameters describing the ion's path are found by maximizing the likelihood, i.e. by minimizing the χ^2 factor. The derivative of the posterior likelihood with respect to y_1 is:

$$\nabla\chi^2 = (\Sigma_1^{-1} + R_1^T\Sigma_2^{-1}R_1)y_1 - \Sigma_1^{-1}R_0 y_0 - R_1^T\Sigma_2^{-1}y_2 \tag{2.15}$$

Setting this to zero and solving for y_1 results in the following compact expression for the MLP:

$$y_{MLP} = (\Sigma_1^{-1} + R_1^T\Sigma_2^{-1}R_1)^{-1}(\Sigma_1^{-1}R_0 y_0 + R_1^T\Sigma_2^{-1}y_2) \tag{2.16}$$

2.3.2 Most likely path in heterogeneous medium

The omission of heterogeneities when applying the MLP formalism in a water medium may introduce potential discrepancies in the reconstructed trajectory. Though it is true that human tissues are mostly made up of water, the different stopping powers encountered by the proton when moving across heterogeneous materials should not be underestimated in terms of accuracy. In this scenario, the terms depending on the material's properties in the scattering matrix formula change at different depths along the ion's path, so the energy evaluation has to be adapted to take into account these specific conditions. The estimation of the elements expressed in the scattering matrix from equations 2.7, 2.8 and 2.9 requires not only the computation of the product between the proton's momentum $p(u)$ and velocity $v(u)$ as a function of the depth in the medium, but also that of the local radiation length $X_0(u)$ and the density $\rho(u)$ for each step of the MLP algorithm.

Briefly, the momentum-velocity product can be expressed in terms of a parameterized function $\xi(u)$ mapping the particle path at a specific depth u to a 3D vector position by estimating the MLP on each orthogonal plane relative to the nominal direction:

$$pv(\xi(u)) = \frac{E(\xi(u)) + 2m_pc^2}{E(\xi(u)) + m_pc^2}E(\xi(u)) \tag{2.17}$$

where m_pc^2 represents the proton's rest mass. The local energy $E(\xi(u)) : \mathbb{R}^3 \mapsto \mathbb{R}$ can be inferred based on prior knowledge of the 3D RSP map:

$$E(\xi(u)) = E_{in} - \int_{\xi(u_0)}^{\xi(u)} RSP(\xi(u'))S(I_w, E(\xi(u')))d\xi(u') \tag{2.18}$$

where E_{in} is the proton's initial kinetic energy, I_w is water's mean excitation energy and $S(I_w, E(\xi(u')))$ is the stopping power from Bethe-Bloch's formula for an ion in a water absorber. The term u' represents the integration variable along the ion's path from the entry point u_0 to the current depth u: therefore, the integral sum allows to determine the energy loss from the start of the path to the current position. Following the standard MLP formalism, the ion's trajectory across the medium is divided into smaller segments where the properties are now different according to the material's nature. The total energy loss along the entire path is thus cumulatively retrieved by progressively adding the current energy loss at each small segment. The variances and covariances in the scattering matrix are then easily derived as Eyges moments following the same principle.

In this work, the models proposed by Schulte et al. [2008] and Collins-Fekete et al. [2017a] were used to reconstruct trajectories in water and tissues respectively. The information needed as prior knowledge for the two scenarios was extracted from MC-simulations as well as the available X-ray CT data. The detailed process is explained in the following sections.

2.3.3 List-mode filtering

Before the actual trajectory estimation process, data are filtered so to remove all irrelevant or inconsistent events.

First of all, all particles that were not protons were removed, along with some anomalous secondary particles produced during the simulations. Protons that were not recorded on both detectors were excluded to ensure that only complete tracking history was considered. Moreover, protons with energy loss below 1 MeV or WET below 1 mm were filtered out, since they were likely noise or not relevant for accurate trajectory reconstruction. Finally, a 3σ cut was implemented (Schulte et al. [2008]). The need to include this step arises from Fermi-Eyges Gaussian model of MCS: to compensate for this approximation, large-angle scattering events are removed. In this work, if less than 10% of the protons were left, the respective pencil beam was removed entirely before the cut. Otherwise, the joint distribution of the relative position (i.e. difference between proton exit and entry position) and angle (i.e. difference between proton exit and entry angle) were considered: if the proton was outside either one of the joint distributions, then it was discarded.

2.3.4 Path estimation models: the scattering line

The MLP formalism describes the concept of ion trajectory by statistically combining the multiple Coulomb scattering model and the Bayesian inference: in order to do this, it is necessary to measure the ion's position, direction and energy at specific locations throughout its run inside the patient. Because of this, list-mode detectors are especially suited for such a task.

To reconstruct the trajectories under a water-based assumption, masks of the patient's anatomy, along with the baseline calibration curve from MC simulations, were utilized as follows. A water map was simulated by setting unit RSP inside the anatomy itself, while the outside of the anatomy was accounted as air by setting each pixel to the RSP corresponding to a HU value of -999. According to the same strategy, a density map and a radiation length map were achieved by considering the respective quantity for a HU value of 0 inside the mask and a HU value of -999 outside. An example of this for one patient's slice is reported in Fig. 2.4.

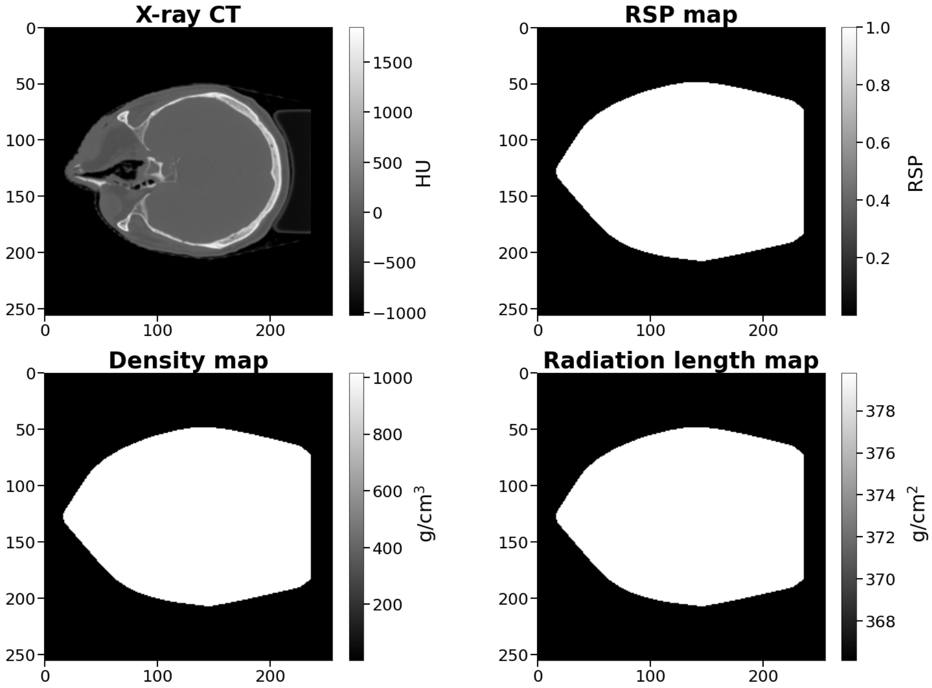

Figure 2.4: Distribution maps of RSP, density and radiation length in water-based scenario. When no prior knowledge about the traversed object is available, the water-base assumption is used. Exploiting the correspondence from HU to RSP, density and radiation length, it was possible to obtain binary masks of these physical quantities accounting for water throughout the entire anatomy and for air outside of it.

The cross-section marking the entry slice of the proton in the imaging volume for all three maps is then rotated by an angle opposite to the current projection angle and transposed; along with the information about the entry and exit position and direction of the proton along the three coordinates and its initial energy, these data are used to compute the MLP. The ion's nominal direction is set to follow the y-axis moving along the depth of the object to be imaged, so the MLP was actually computed in order to find the coordinates on the orthogonal axes x and z. Particularly, the depth-dependent energy vector was calculated applying equation 2.18, where the RSP value corresponding to current depth was extracted from the column of the rotated RSP image marking the entry of the proton in the imaging space, while the mean excitation energy of water for the stopping power computation was set to 75 eV. After this, the scattering matrices were retrieved as explained by Schulte et al. [2008] and Collins-Fekete et al. [2017a], where the integral sum was estimated through Simpson's rule and accounting for local density and radiation length across the entry column of the respective rotated image. Variances and covariances were assessed moving along the nominal direction considering a spacing of one from the entry point to the exit point along the y-axis in pixel coordinates, progressively adding the integral contribution at every location. Once the transvection matrices were also implemented, the state vector along the x and z coordinates was easily derived from equation 2.16. Next, the physical coordinates of the MLP were linearly interpolated so to fit the trajectory into the image frame: specifically, a unit value was weighted on a $3\times3\times3$ grid centered around the physical point itself. The output volume has the same dimension of the X-ray CT in the imaging plane (256×256 $pixels$) with a slab of 7 around the slice of interest. Trajectories that did not fit the imaging volume in the orthogonal directions were excluded by returning an all-zero matrix. Each trajectory thus reconstructed is then transposed and rotated back according to the projection angle and finally stored as a column of a sparse system matrix.

The reconstruction of the ion path in non-homogeneous materials was approached with a similar strategy. The prior knowledge data to use for the trajectory's estimation had in fact to be adapted to the properties and heterogeneities of the traversed tissues. In particular, the X-ray CT image was interpolated according to the baseline calibration curve: in this way, a 3D RSP, density and radiation length map (Fig. 2.5) were derived to supply to the MLP algorithm.

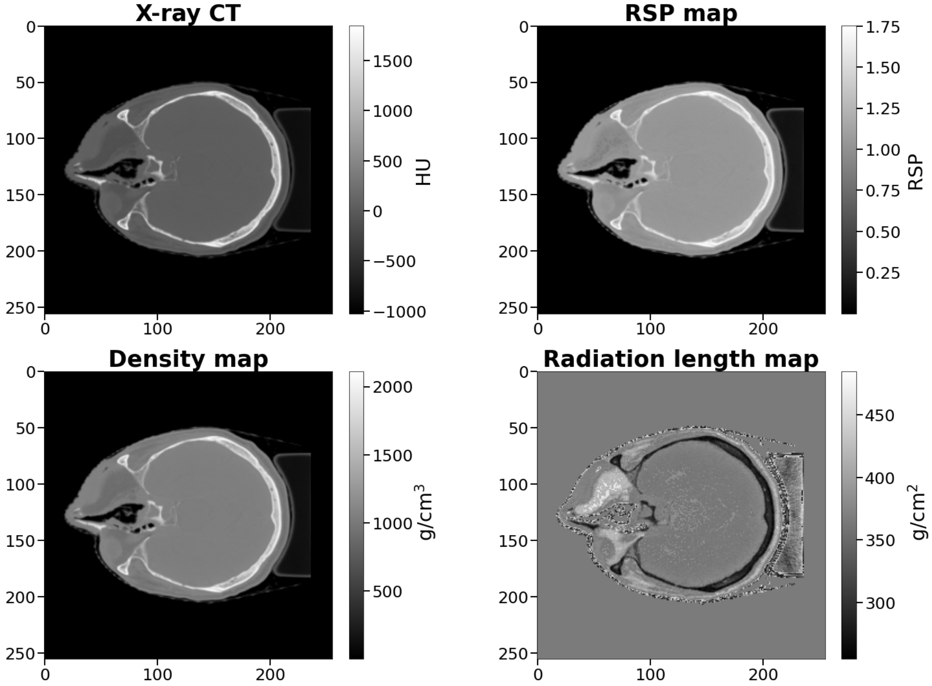

Figure 2.5: Distribution maps of RSP, density and radiation length in tissue-based scenario. By means of calibration curves from MC simulations, interpolation across the entire HU spectrum of the X-ray CT provided the spatial distribution of RSP, density and radiation length for the patient. A cross-section from a certain level of the patient slab is showed in this figure.

The following computational steps remained identical to the water-based scenario.

The strategy of rotating the maps of the physical quantities needed for the MLP allowed to reconstruct the trajectory without accounting directly for the projection angle, as if the proton itself was moving in the same direction as the normal to the detector plane while the patient was rotating instead of the gantry. This was done to decrease computational times and also errors during the interpolation process. Furthermore, the image maps are transposed after the rotation because, as per designing choice, the model was made to work along the columns of the imaging space so to stay consistent with the projectile's nominal direction. The interpolated trajectory is transposed back exactly to compensate for this mechanism. Finally, the choice of retrieving the physical information from the entry slice and column was carried out as a simplifying assumption: the model in fact tries to determine how much a proton would scatter if it would travel on a straight path through the column of that image slice.

The MLP formalism represents a powerful tool for imaging purposes. As a matter of fact, modeling the trajectory as a straight line may underlie inaccuracies in the representation of the physical problem as well as poor spatial resolution in the image itself. Factoring in the scattering events, the MLP allows for a more accurate interpretation of the particle's physical behavior, providing for trajectories that are more curved. The trajectory showed in Fig. 2.6 is a prime example of a more realistic reconstruction of a proton's path across a medium. Since a single trajectory is reconstructed by computing the proton's most likely location one after the other across the entire object, this model will be referred to as *scattering line* in this work.

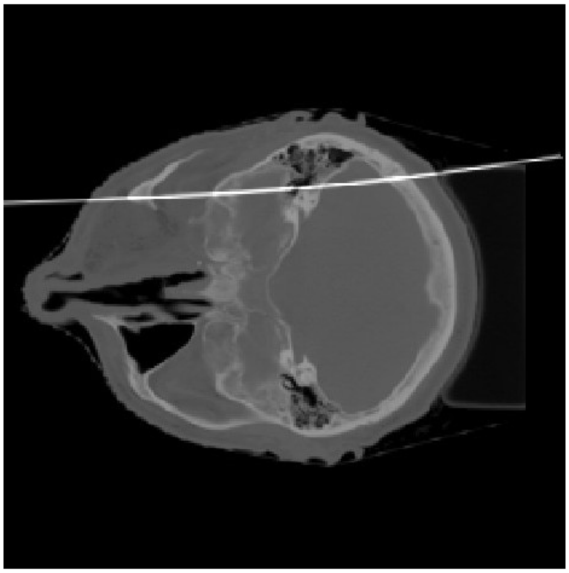

Figure 2.6: Scattering line model. Most probable trajectory of a proton across the imaging space. Though the scattering is minor through air and soft tissue, upon collision with bony structures, the particle is deflected significantly. Here, the proton image was darkened to highlight the particle's path.

2.3.5 Path estimation models: the scattering spindle

Though the MLP represents an accurate method, this is still a probabilistic approach, implying that the most likely trajectory is not necessarily the actual one. Thus, a further step can be taken to account for the variability in the reconstructed path. One of the perks of using a Gaussian approximation for MCS is that the uncertainties in the MLP-estimated trajectories can be modeled as a volumetric distribution of all possible trajectories at any given depth (Fig. 2.7).

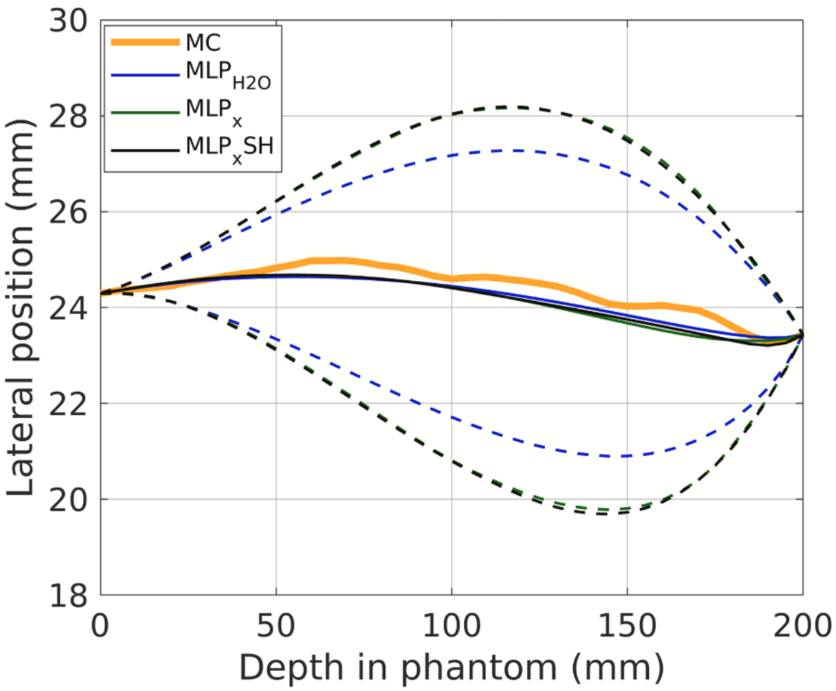

Figure 2.7: Gaussian error envelop around most likely trajectory. The lateral displacement of the trajectory at any given depth across the traversed volume can be modeled as an error envelop in the nominal direction to create a Gaussian volumetric distribution (dashed line) of all possible trajectories around the most probable one (solid line). The line in bold represents the actual Monte Carlo proton tracks (figure from Brooke and Penfold [2020]).

Schulte et al. [2008] proposed an addendum to their formalism. Considering once again position t_1 and angle θ_1 at a certain depth u_1, a curvature matrix can be defined as such:

$$\alpha_{t_1\theta_1} \equiv \frac{1}{2}\frac{\partial^2\chi^2}{\partial t_1 \partial\theta_1} \tag{2.19}$$

where

$$\frac{\partial^2\chi^2}{\partial t_1 \partial\theta_1} = \Sigma_1^{-1} + R_1^T\Sigma_2^{-1}R_1 \tag{2.20}$$

From this, it is possible derive an error matrix as the inverse of the curvature matrix:

$$\epsilon_{t_1\theta_1}(u_1) = 2(\Sigma_1^{-1} + R_1^T\Sigma_2^{-1}R_1)^{-1} \tag{2.21}$$

The element in the first row and column of the error matrix returns the variance in the lateral displacement at the depth u_1.
By embedding this mechanism in the MLP computation, it was possible to derive a vector containing the variances in the lateral displacement for each point of the ion trajectory, once again sampling the proton's path pixel by pixel in its nominal direction across the entire depth of the imaging reconstruction volume. The square root of these elements returned the standard deviations of the error envelop as a function of the current depth: the declination of the MCS model in a list-mode application results in a volumetric spindle of cross-sectional Gaussian distributions along the trajectory.
In order to implement the Gaussian volumetric distribution, the physical points of the trajectory were calculated according the MLP formalism followed for the scattering line model (see section 2.3.4), meanwhile retrieving the array containing the standard deviations. The linear interpolation instead needed to be adjusted accordingly. The scattering line grid was extended along the first and third dimension proportionally

to 3 times the standard deviation at the considered depth, while the size along the second dimension remained untouched. This was done to accommodate a Gaussian filtering: for each point of the trajectory, this size-varying kernel was in fact smoothed out by a 3D Gaussian filter defined by the respective standard deviations (scaled to pixel size) along the first and third dimension, while the one in the second dimension was practically[2] set to zero. This allowed to impose matter-of-factly a 2D filtering directed across the plane perpendicular to the ion's nominal direction (i.e. the y-axis) where the MLP was effectively acquired. The kernel's dimensions were tailored according to the empirical rule of statistics, i.e. to include 99.7% of the data. The kernel was then interpolated to image coordinates and the resulting volume allocated in the system matrix.

It is worth pointing out that the filtering plane would not generally match the ion's transversal one since applying the Gaussian filter normally to the nominal direction does not entirely account for the trajectory's curved shape. Nevertheless, the smoothing process is carried out voxel-wise along the trajectory: owing to the 3σ cut and the small-angle approximation, the proton's deviation from a straight path at that specific depth was considered negligible.

The lateral displacement experienced by the proton, and therefore the shape and size of the Gaussian envelop, are dependent on the material composition of the medium traversed by the particle. The water-based scenario only provides a binary environment so that the particle is either moving in air or in water; in case of inhomogeneous medium, the proton interacts with various kinds of heterogeneities dictating a behavior that is reflected in the probability envelop. Fig. 2.8 shows an example of how the anatomy's properties may result in different types of Gaussian distributions around the trajectory. Due to the resemblance to a spindle-shape, this model will henceforth be termed as *scattering spindle*.

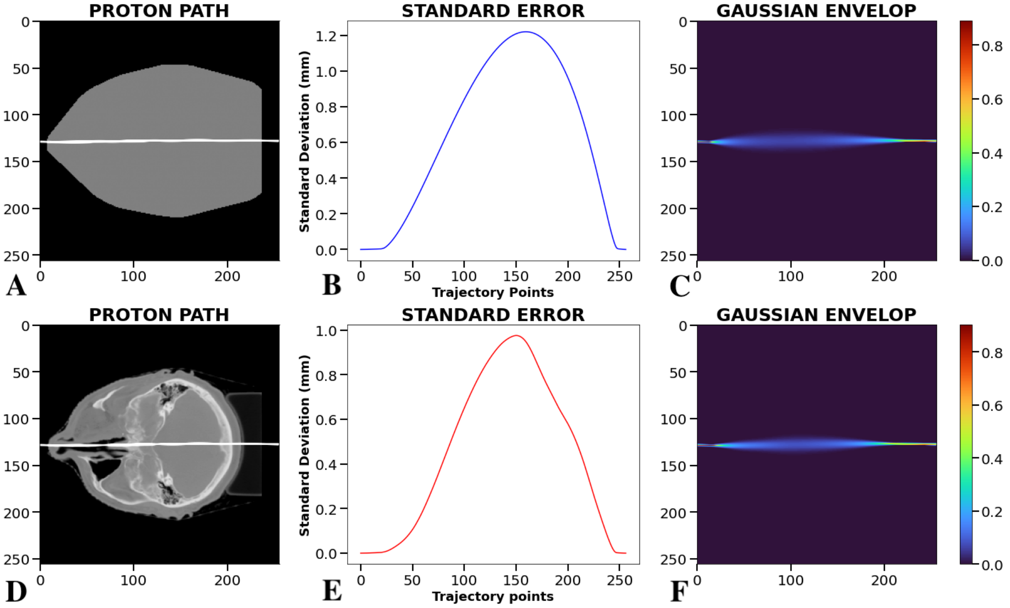

Figure 2.8: Scattering spindle model. This picture illustrates how the probability envelop is affected by the proton's history across the traversed medium, specifically water (top row) and tissue (bottom row). Here, the trajectory refers to a 179° projection angle, so the proton is moving from the right side of the image to the left.
Note: for representative purposes, the spindle trajectories shown in picture **C** and **F** were realized increasing by a factor of 5 the respective standard deviations in picture **B** and **E**).

In this case, the water spindle closely follows the theoretical envelop, though a slight shift of the Gaussian profile can be noticed due to the uneven amount of air encountered by the proton when entering and leaving the anatomy. This is also a possible indication that, when assuming no prior knowledge, the highest error in

---

[2] The coding language used in this work (Python 3.11, https://www.python.org) did not allow to set the standard deviation value as input to the Gaussian filter directly to zero, so it was approximated to the smallest possible number.

the MLP approximation, on average, unfolds downstream from the center of the object (Schulte et al. [2008]). The standard deviations are instead more sensitive to changes in the anatomy composition, as show in Fig. 2.8-E: in particular, a hump-like feature in the Gaussian profile follows the discontinuity in the passage of the proton from soft tissue to bone tissue. Furthermore, though this may not always be the case, the spindle in Fig. 2.8-C is bigger than the one in Fig. 2.8-F, as pointed out by the peak of the Gaussian in the two cases. In this instance, the overall RSP contribution determines a WET experienced by the proton that is higher in uniform water than in tissues, where the particle is going through air not only outside the patient's geometry but also inside cavities in the anatomy.

The introduction of the spindle model into the MLP formalism may help in a more thorough estimation of the proton's scattering events inside the object to be imaged and may represent an important step forward in improving ion image reconstruction and spatial resolution.

2.3.6 Ordered Subsets Simultaneous Algebraic Reconstruction Technique

Once the trajectory models were assessed, the images were reconstructed. The 3D RSP map was derived by approaching the solution of equation 2.1 with a reconstruction algorithm from the family of projection methods: by repeatedly employing projections onto convex sets, they are well suited when it comes to feasibility-seeking or optimization problems (Penfold and Censor [2015]). The iterative algorithm adopted in this project to retrieve list-mode images is known as Ordered Subsets Simultaneous Algebraic Reconstruction Technique (OS-SART). The OS-SART algorithm is a valuable resource for list-mode tomographic ion imaging since the processing of large datasets, such as the system matrices, is accelerated. In fact, instead of accessing all the projections at the same time, the image is updated based on subsets of ordered projections.

In this work, for each projection angle, a system matrix per pencil beam was created: the matrix was thus defined by as many rows as the number of voxels in the imaging reconstruction volume (256 × 256 × 7) and as many columns as the number of protons whose trajectories actually intersected the reconstruction volume. Next, 36 subsets were assembled by adding sequentially pencil beam by pencil beam picking the matrices with a step equal to the number of subsets itself; this was done for all projection angles to ensure that each subset contained information about the entire tomographic acquisition. To avoid repetitions in the same subsets or overlapping data in different subsets, pencil beams already used in previous assemblies were excluded from the remaining ones. By following the same logic, the WET values associated with each proton in the subset were compacted into an array.

An iteration of the algorithm is divided into sub-iterations carried out on a single subset of projections: a single OS-SART iteration can be considered complete once all subsets have been used to update the image. Particularly, the updating formula for the reconstruction of the image is expressed in the following equation:

$$\overrightarrow{f^{(s+1)}} = \overrightarrow{f^{s}} - \frac{\sum_{i \in S(s)} \left(a_{ij} \cdot \left(\frac{\overrightarrow{f^{s}} \cdot \overrightarrow{a_{ij}}^{T} - \vec{g}}{\sum_{j} a_{ij}} \right) \right)}{\sum_{i} a_{ij}} \tag{2.22}$$

where f is the image vector, S is the subset of projections at the current iteration s and $\vec{g}$ is the measured data, i.e. the ion radiography.

The reconstruction process starts with an initial guess of the image vector $\overrightarrow{f^{0}}$. For each subset $S(s)$, the forward projection of the current image estimate is computed by a matrix multiplication between the transposed system matrix A^{T} and the current image estimate $\overrightarrow{f^{(s)}}$; the measured data $\vec{g}$ is then subtracted to the forward projection to find the discrepancy, as Euclidean distance, between the two. This discrepancy is then back-projected to the image space through another multiplication with the system matrix. The normalizing terms are introduced to account for the contribution of every element in the system matrix. Throughout each step, the algorithm aims to find that image vector that minimizes the error between the forward projection and the ion radiography: if the convergence is reached, the reconstructed image is assumed validated. A summary workflow of the OS-SART iterative reconstruction process is illustrated in Fig 2.9. The image is updated over the course of numerous iterations over the chosen subsets: this introduces information redundancy during the reconstruction process that is helpful in mitigating the inaccuracies in the forward projection model.

Finally, a drawback of the OS-SART algorithm is that the increased convergence speed and reduced

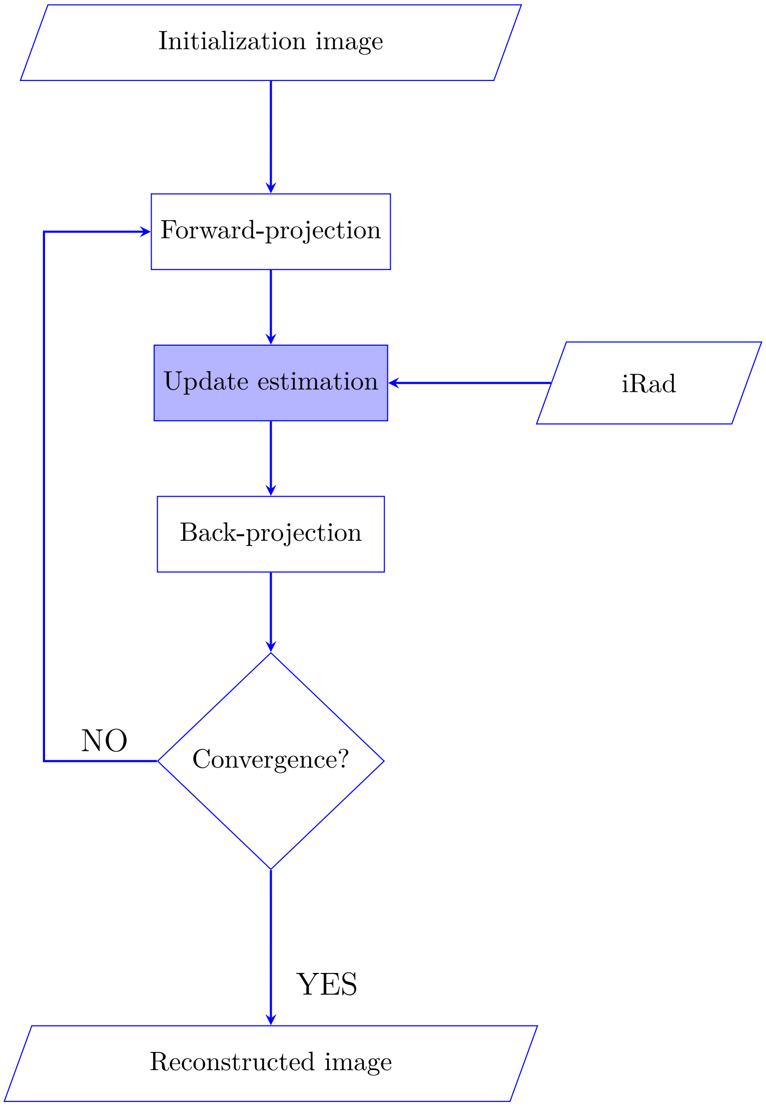

Figure 2.9: Flowchart of the OS-SART algorithm. The converge is reached through a series of iterations during which the Euclidean error between the forward projection of the current image estimate and the ion radiography is back-projected to imaging space.

memory requirements come at the cost of a much higher noise amplification across the reconstructed image. This phenomenon is usually known as *noise break-up* and can be observed as the image darkens and starts to be perturbed by the so-called salt-and-pepper noise as the number of iterations progresses. The iterative update of the forward projection is most likely impaired because of inconsistencies in the measured data: in fact, scattering processes due to ion-nuclei interaction, statistical variations in the proton's energy (see section 1.7) and inaccuracies in the estimated trajectories may lead to a gradual halt of the monotonic increase in image quality. Because of this, the optimal solution minimizing the objective function of the tomographic image reconstruction algorithm is not necessarily the solution that best resembles the Ground Truth image (Penfold et al. [2010]). Different approaches have proposed to compensate for noise amplification: for instance, a straightforward strategy could be to simply stop the algorithm before the solution diverges. A more sophisticated alternative relies on superiorization techniques introducing controlled perturbations in the solution at each iteration to try and find another solution that is superior to the current one according to a merit function.

2.4 Proton-integrating systems

Integration-mode detectors allow to simplify the data acquisition process by removing the tracking systems and synchronizing the total absorption detector with the pencil beam delivery system. To match this clinical scenario, position and direction measurements from MC simulations were ignored, while keeping information about the residual energy of the pencil beam. As a natural consequence, the reconstruction of the trajectories could not rely on the MLP formalism, lacking its input parameters.

The pre-processing of the data was also adapted to this family of detectors: particularly, since the joint angle or position distributions were no longer available, the filtering was limited to insufficient energy requirements or anomalies in the MC simulations themselves.

2.4.1 Range measurements in integration-mode configuration

For integration-mode detector configuration, the WET information is carried by the ensemble of protons in the pencil beam itself. Usually, when the energy detector consists in a single absorption layer, the mean energy is converted to a single WET value associated to that specific pencil beam. Nevertheless, information about WET inhomogeneities can be introduced by discretizing the total absorption detector into a stack of absorbers and ionization chambers (i.e. channels). The mixed energy is thus statistically resolved and converted to a histogram of WET occurrences expressing the weight of each WET component. Theoretically, the complete WET histogram is retrieved by means of linear decomposition of the signal (Meyer et al. [2017]) by solving the following system of linear equations:

$$\overrightarrow{BP} = LUT \cdot \overrightarrow{WET} \tag{2.23}$$

where $\overrightarrow{BP}$ is the discretized Bragg peak signal, $\overrightarrow{WET}$ is the unknowns vector of WET occurrences and LUT is the look-up-table of individual Bragg peak signals for each WET component.

In the context of this work, it was possible to rely on range measurements coming directly from the MC platform. The WET data carried by the individual protons were distributed into a histogram binned at pixel-sized steps from 0 mm to 258.74 mm for a total coverage of 241 channels (Fig. 2.10). The rightmost edge value of the WET histogram was chosen so to approximately accommodate the range of protons in water. WET histograms were generated for each pencil beam across all projection angles.

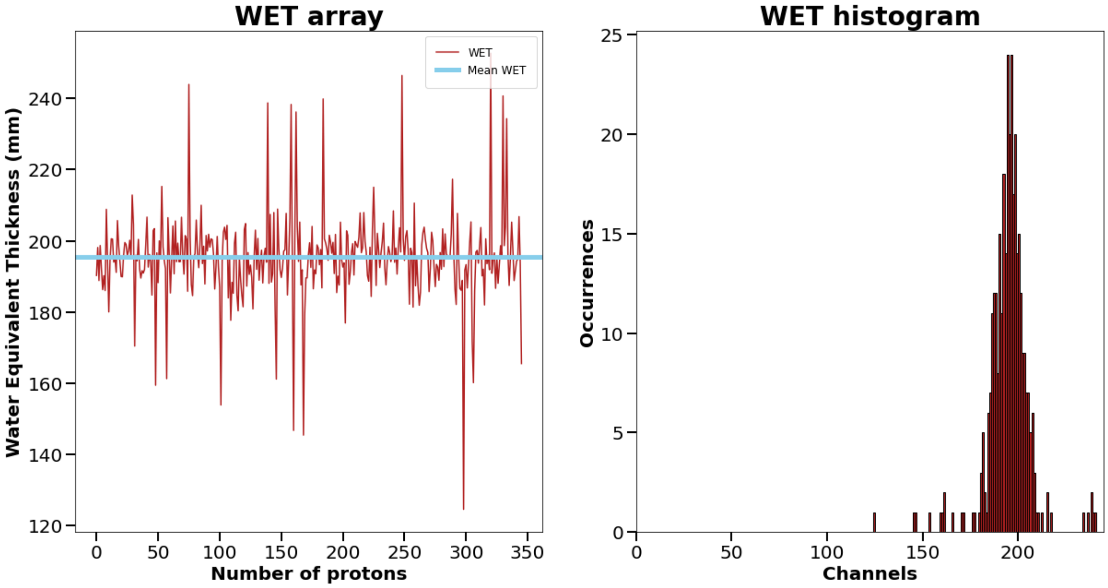

Figure 2.10: WET histogram. This figure shows an example of a WET array binned into the respective histogram of WET occurrences for a pencil beam. As indicated by its mean value, the majority of WET components are distributed around 200 mm, so the highest occurrence in the histogram is roughly at channel 200.

The composition of the WET histogram is of the utmost importance for integration-mode tomographic ion imaging: in fact, since the reconstruction is based on the general concept of beam trajectory, one specific WET value has to be chosen to depict the entire set of protons in the pencil beam. The histogram helps in an accurate estimation of the WET heterogeneities encountered by the protons, so to pick the WET value that is deemed more statistically representative.
The reconstruction specifics adopted in this work for integration-mode imaging will be discussed in later sections.

2.4.2 Path estimation models: the nominal line

Proton-integrating systems are not as convenient as list-mode detectors when it comes to the reconstruction of the trajectory. Not having the possibility of following the individual ion's path across the imaged volume, the integration lines were assumed as straight. Relying on the FRED working set-up (see Fig. 2.1), the deflected beam was parameterized as a straight line between the point at the exit of the BAMS and the irradiated raster point. Assuming the imaging volume to be centered around the isocenter itself, the line was sampled in an interval ranging from -13.7472 cm to 13.7472 cm along the beam's nominal direction with a step of 0.1074 cm, that is the voxel size along the y-axis: this assured that the number of sampled points amounted to number of pixels in the imaging space. This procedure returned an array containing the physical coordinates of the pencil beam in the normal $x-z$ plane, then interpolated and rotated as explained in section 2.3.4 for the scattering line model. The resulting volume was thus vectorized and stored in the system matrix.
Since the reconstructed trajectories were simply approximated as straight lines along the y-axis, this model was referred to as *nominal line* during this study. An example is provided in the figure below.

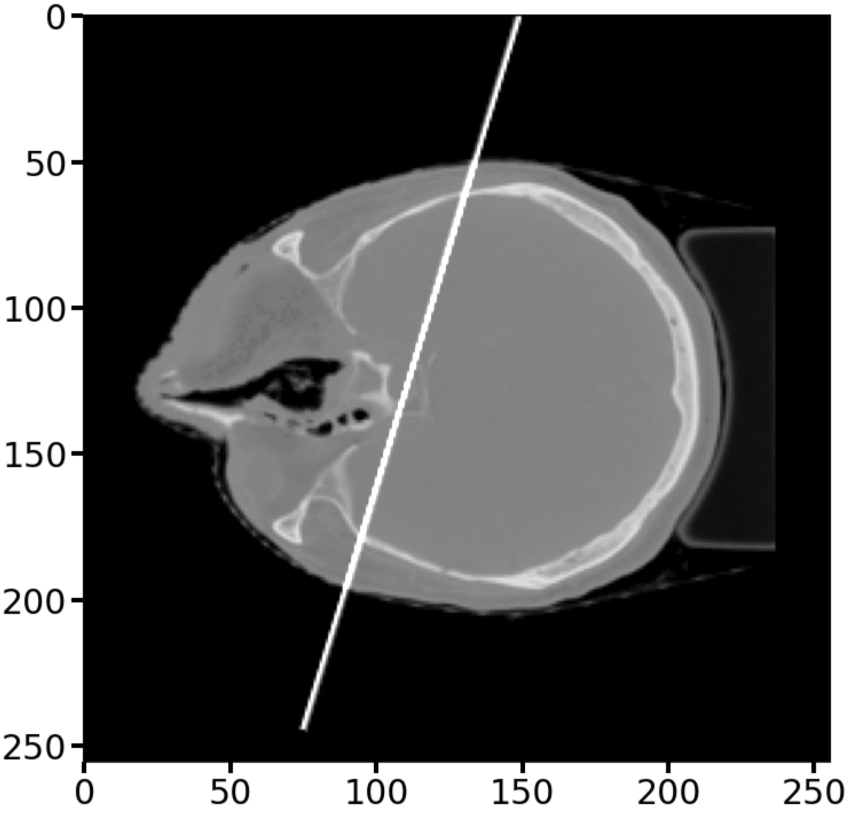

Figure 2.11: Nominal line model. Since no information is available on entry and exit position and direction, the concept of ion beam trajectory is modeled as a straight line for integrating systems, thus lacking knowledge on the scattering phenomena addressed in the list-mode counterpart.

Due to the nature of the delivery system, the inclination of the pencil beam is modulated not only by the projection angle but also according to the deflection angle imposed by the scanning magnets. Because of this, the straight trajectories have a tendency to symmetrically diverge from the nominal line at the center of the image, resulting in a fan-spread set of pencil beams for that projection angle (Fig. 2.12). The scanning angle is instead neglected for trajectory reconstruction in list-mode configuration: in fact, not only the MLP algorithm allows to meticulously trace the individual proton's path, but the entry and exit point across the imaged object are explicitly known.
Finally, since scattering events are neglected, the straight trajectory does not require any data about the traversed medium and its properties: the nominal line model is thus valid in water as well as tissues.

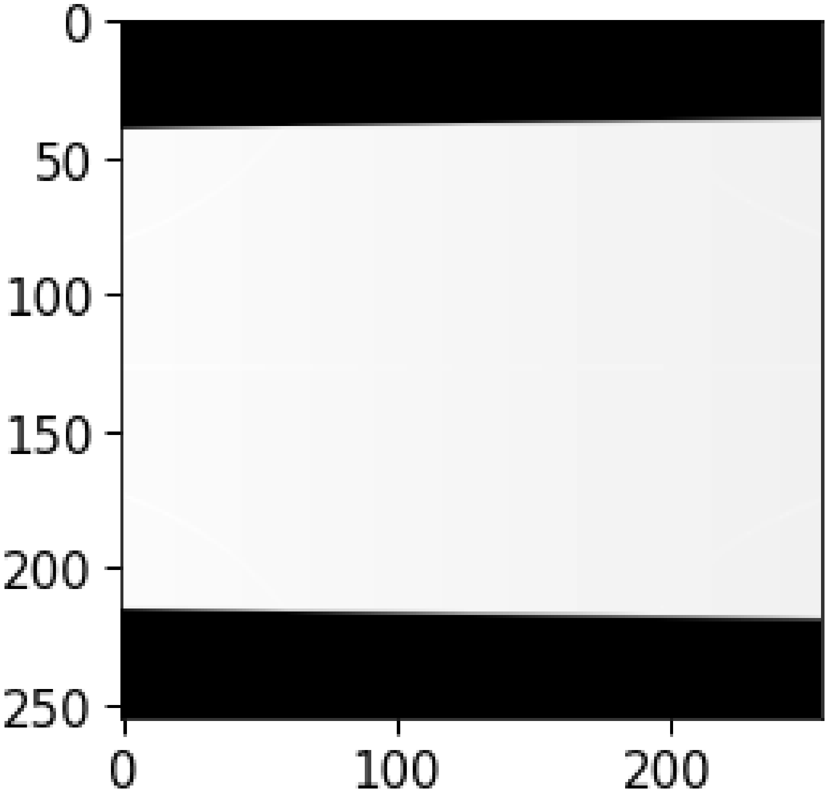

Figure 2.12: Pencil beam geometry for integration mode. Since the deflected pencil beam is modeled as a straight line between the BAMS and the Raster Grid, the angle imposed by the scanning system creates a divergence from the center of image in the beam's nominal direction. The overall effect results in a spread-out fan of pencil beams for the projection angle.

2.4.3 Path estimation models: the scattering cone

As stated in the previous section, the scattering line concept as most probable ion trajectory turns into a straight line along the pencil beam for integration-mode data. The inaccuracies underlying this model are evident since it just describes the most likely range (or distance) a proton will travel in a straightforward path. The nominal line model simplifies to its core ion-matter interactions: though it is true that heavy ions may travel in a basically straight trajectory, the behavior of protons is not fully accounted for within this pattern.
The use of scattering models consistent to integrating systems has been recently introduced both in ion tomography and radiography: more precisely, given that the nominal pencil beam direction and dimension are known, the beam trajectory can statistically be described by the Gaussian Multiple Coulomb Scattering

model proposed by Eyges [1948]. To imbue the trajectory with information about the scattering power, data about the beam's surroundings need to be provided: once again, the RSP information, as well as density and radiation length, is derived from the patient's anatomy whether homogeneous or heterogeneous.

In order to wrap the Gaussian distribution around the nominal pencil beam, a framework similar to that of the scattering spindle model was adopted (see section 2.3.5). The variances describing the positional uncertainties along the nominal trajectory were retrieved as the second Eyges moment (see 2.7) and the their square root used to define the Gaussian filter in the normal plane. Once again, the physical parameters needed for energy and scattering computation were extracted along the entry columns of the RSP, density and radiation length map, assuming the entry point to be the first point in the parameterized line. An example of the scattering model for integration mode is illustrated in Fig. 2.13. Since it corresponds to a flared conical Gaussian distribution around the trajectory, this model is called *the scattering cone*.

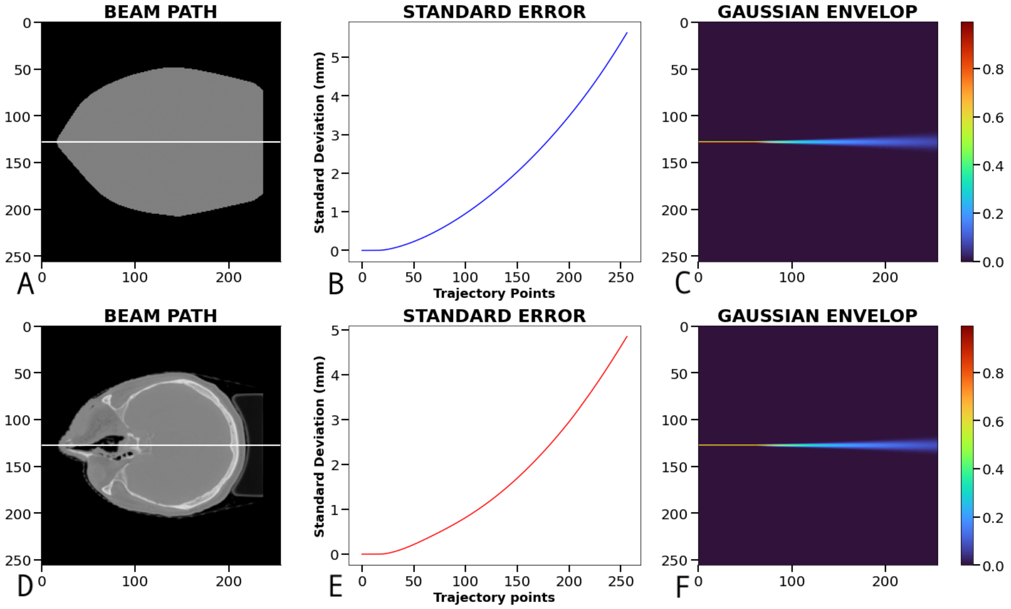

Figure 2.13: Scattering cone model. This figure shows how the original Gaussian approximation of the generalized Fermi-Eyges MCS theory behaves along the nominal trajectory, both in uniform water (top row) and in heterogeneous tissues (bottom row). Shape and size of the Gaussian cone witness the pencil beam's history across the traversed volume.

As the scattering power depends on the geometrical and physical properties experienced by the beam across the medium, the scattering cone reflects the nature of these interactions. For instance, Fig. 2.13-C shows a cone that is larger than the one in Fig. 2.13-F: in fact, while the RSP values are uniform along the beam's path in water, the heterogeneities in tissues may result in a lower integral sum, i.e. the cone in tissues is associated to a lower WET value than the one in water. This is of course imprinted in the standard deviation profile leading to the construction of the cone itself (Fig. 2.13-B and Fig. 2.13-E).

The positional variances are computed the same way both in list-mode and integration-mode configuration. Nevertheless, the declination of the scattering power in the MLP Bayesian framework allows to confine the statistical envelop of all possible trajectories between the measured entry and exit locations, providing for a distribution that necessarily starts from one single point and converges to one single end point. Not being able to rely on spatial information for integration-mode detectors, the scattering model accounts only for the cumulative progression of RSP along the pencil beam's trajectory resulting a diverging profile of the variances and, therefore, the cone.

In order to model the Gaussian uncertainties across the entire length of the beam, the Gaussian filter was applied voxel-wise to the trajectory: under the same assumptions made for the scattering spindle model, the deflection angle imposed to the pencil beam by the scanning magnets was considered locally negligible, allowing for a more straightforward filtering in the plane orthogonal to the nominal direction.

2.4.4 Simultaneous Algebraic Reconstruction Technique

Whereas in list-mode configuration, each pencil beam contains a certain number of protons, the integration-mode system matrix is limited to the single integration line. Projection algorithms for iterative reconstruction aim to minimize the error between the estimated projections and the actual measurements. If the projection data are insufficient, the system of equations describing the problem may fail in representing the image itself: particularly, noise amplification problems may arise directly from the beginning of the iterations owing to the fact that high frequencies, that are more sensitive to noise, are accentuated being the first spectral components to converge. To avoid fluctuations in the image reconstruction process, the system matrices from all the 180 projection angles were assembled into one subset used during all the iterations. Since the overall projection data is no longer split into different subsets, the reconstruction algorithm adopted in this work for integration-mode configuration should fall under the proper name of *Simultaneous Algebraic Reconstruction Technique* (SART).

For the comparative update of the iterative algorithm, it was necessary to reduce the WET histogram to a single WET component. Defining p_{ik} the WET component and w_{ik} the WET occurrence for the i-th pencil beam and the k-th channel, two theoretical approaches are usually implemented: the first one consists in a weighted average of the WET components as:

$$WET_i = mean\{p_{ik}w_{ik}\} \tag{2.24}$$

The second alternative was instead adopted in this study and is based on picking the WET component with maximum occurrence across all channels, that is the *mode WET* or WET_{mode}:

$$WET_i = max\{p_{ik}\} \tag{2.25}$$

However, the mode WET can become ambiguous when two or more channels share the same occurrence. In such an instance, the neighbour of these channels was explored: the one with the highest total occurrence defined the WET_{mode}.

Once the system matrix and the respective WET_{mode} values were defined, the updating procedure followed the same workflow outlined in section 2.3.6 and depicted in Fig. 2.9.

2.4.5 Path estimation models: the scattering cones

The statistical description of the integral scattering power as a Gaussian distribution around the nominal trajectory should guarantee a better estimation of the pencil beam's path across the object to be imaged. Even so, the image reconstruction is based on a single WET value associated to each integral line: though the mode offers a good indication of the central tendency of the WET distribution, this oversimplification may still lead to consistent information losses regarding the cross section's actual anatomy and properties. To make the most of integration-mode data, a novel approach is proposed where a scattering cone for each range component of the pencil beam is adopted.

Under the water-based assumption, the original patient's mask was exploited (see 2.3.4). The integral sum of the RSP values along the beam's nominal trajectory was calculated to retrieve the $WET_{nominal}$ value associated to that ensemble of protons. Each component from every channel of the WET histogram was then divided by the $WET_{nominal}$ value in order to obtain a factor f needed to scale the unit RSP value set inside the anatomy and the RSP value accounting for the air around it. By means of linear interpolation on the MC calibration curves, the CT numbers corresponding to the scaled RSP were computed. By the same logic, the values of density and radiation length associated to the scaled HU quantities were derived. These parameters were thus used to create binary RSP, density and radiation length map providing the required data for the computation of the variances and the Gaussian cone.

Exploiting prior knowledge, it was possible to extend this mechanism to the entire HU spectrum. The $WET_{nominal}$ value was derived by the RSP integral sum on the RSP image obtained from the original calibration curves: thus, the scaling factor f was retrieved for each WET component at every channel. The subsequent interpolation between HU and scaled RSP resulted in a scaled RSP map which was in turn interpolated to produce an X-ray map with HU corresponding to the scaled RSP. The new X-ray image was finally interpolated to density and radiation length.

All these steps (see algorithm 1) lead to the creation of a 4D matrix where the last dimension contains the

channels: a single pencil beam for a specific projection angle is thus declined into 241 cones, one for each channel. Due to this nature, this model is thereby termed as *the scattering cones*.
In both water and tissues, the scaled 3D physical maps ensure that the range traveled by the protons actually corresponds to the WET value extracted from the WET histogram at that specific channel. However, it has to be noted that the adoption of the scaling factor is equivalent to redefine the composition and properties of the traversed medium to those of some other material rather than uniform water or the original tissues. Because of this, characteristics such as the traveled range have to be considered within a certain margin from the initial delineations. In order to avoid anomalies in the computed trajectories, two conditions were imposed on both the residual energy and the variances of the Gaussian scattering model. Whenever the range did not show the typical monotonically decreasing trend across the traversed depth, an all-zero column was stored in the system matrix. Similarly, when the variances exhibited a behavior diverging from the monotonically increasing curve across the volume, then the resulting trajectory was set to zero. These atypical traits may arise due to the nature of the algorithm itself. In fact, as the number of channels progresses, the f factor also increases: the respective scaled RSP, density and radiation length define a somewhat unknown material that may result in uncommon physical outcomes (such as too-high scattering angles or unforeseen energy straggling) as the protons move across the patient's anatomy.

Algorithm 1 Scattering cone for each channel

```
for $angle = 0 : N_{projections}$ do
    Computation of $\mathrm{WET}_{nominal}$
    for channels = 1 : $N_{channels}$ do
        if $WET_{nominal} = 0$ then
            trajectory = all-zero array
            Skip to next channel
        end if
        $f = WET_{ch}/WET_{nominal}$
        Scale RSP
        Interpolation from HU to scaled RSP on X-ray CT → scaled iCT
        Interpolation from RSP to HU on scaled iCT → scaled X-ray CT
        Interpolation from HU to density on scaled X-ray CT → scaled density map
        Interpolation from HU to radiation length on scaled X-ray CT → scaled radiation length map
        Compute energy and variances
        if not energy monotonically decreasing then
            trajectory = all-zero array
            Skip to next channel
        end if
        if not variances monotonically increasing then
            trajectory = all-zero array
            Skip to next channel
        end if
        Interpolate Gaussian cone into image frame
    end for
end for
Final dimension ($256 * 256 * 7 * N_{channels}$)
```

The scattering cones allow to enlarge the statistical distribution around the nominal trajectory. The cone on its own represents a probabilistic distribution where the likelihood of the data points decreases moving farther and farther from the nominal trajectory. The progressive trend in the number of channels signals also that the cumulative WET values for that trajectory increases: physically, this implies that the ensemble of ions is going through more material or, at least, denser material. This results in a greater potential for scattering or deviation from the nominal trajectory, that is statistically reflected in the construction of increasingly larger cones (Fig. 2.14) expressing the growing uncertainty or variability in path estimation due to more material or regions with higher stopping power along the beam's path. By ranging across the entire WET histogram, the cone built relying on a scaling factor $f \approx 1$ represents the actual WET value encountered by the nominal beam: the addition of the fourth dimension (i.e. the channels) allows to provide information about the overall range for a more accurate model and prediction of the radiation path which is especially important in medical applications where understanding the variability in tissue density is crucial for precise dose delivery.

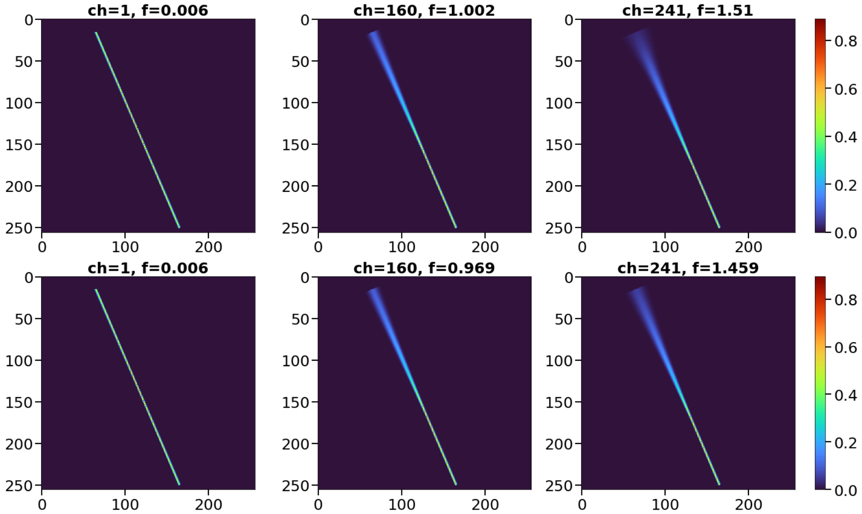

Figure 2.14: Scattering cones model. As the number of channels progresses, the scaling factor f increases leading to the construction of ever larger cones, both in water (top row) and tissues (bottom row). This represents the spread in the uncertainties associated to more consistent ion-matter interactions.

2.4.6 Scattering cones: dedicated tomographic image reconstruction algorithm

In the conventional integration-mode scenario, the image reconstruction relies on the identification of the maximum of the Bragg peak signal as the WET component with maximum occurrence. When adding the fourth dimension to the system matrix, the contribution of all channels should be introduced in the reconstruction algorithm, along with the impact of each cone on the scattering model dictated by the number of occurrences on each channel (Fig. 2.15)

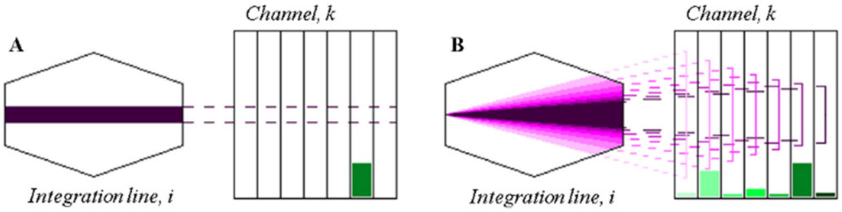

Figure 2.15: Concept of reconstruction algorithms for integration-mode detector configuration. When the integration lines are modeled in a straight path, the conventional tomographic reconstruction (A) extracts the value with the highest occurrence from the WET histogram. When a 3D Gaussian cone is implemented along the trajectory for each channel, the reconstruction algorithm is able to exploit the entire WET spectrum by weighting each cone contribution for the number of occurrences associated to that channel (B) (figure from Oria et al. [2018]).

Following the approach of Oria et al. [2018], a dedicated algorithm for tomographic reconstruction including the entire WET histogram is proposed. Defining k = {1, ..., K} the number of channels (i.e. the number of WET components in the histogram) and i = {1,..., I} the number of integration lines, the projection data is expressed as:

$$p_i = \sum_k W_{ik} p_{ik} \tag{2.26}$$

where p_{ik} represents the WET components for each integration line. These terms are weighted for a probabilistic interpretation of the WET occurrences: the parameter W_{ik} represents in fact the WET occurrence on the k-th channel for the i-th trajectory divided for the total amount of occurrences in the histogram:

$$W_{ik} = \frac{w_{ik}}{\sum_k w_{ik}} \tag{2.27}$$

where w_{ik} is the WET occurrence for each integration line.
The system matrix was assembled in a similar manner, statistically weighting the cones on each channel of the i-th integration line. The weighted channels were then summed together resulting in a single column of the system matrix for that pencil beam:

$$m_i = \sum_k W_{ik} a_{ik} \tag{2.28}$$

where a_{ik} describes the scattering power assigned to the integration line at each channel. The matrices thus created, and the respective WET values, were assembled into 180 subsets for each projection angle. These subsets were in turn combined into a single projection set exploited for the SART algorithm. A summary of the reconstruction technique is reported in algorithm 2.

Algorithm 2 Dedicated algorithm for integration-mode detector configuration

```
for $angle = 0 : N_{projections}$ do
    for $i = 1 : N_{beams}$ do
        $WET_i = (WET_{occurrences}/TOT_{occurrences}) * WET_{components}$
        for $k = 1 : N_{channels}$ do
            $m_i = m_i + (WET_{ik}/TOT_{occurrences}) * a_{ik}$
        end for
        Store matrix for i-th integration line
    end for
    Store subset for a-th projection angle
end for
Apply SART algorithm
```

2.5 Data analysis

Two slices were reconstructed for three patients according to all path estimation models presented in this work. Particularly, list-mode reconstruction was limited to 30 iterations for computational reasons, whereas 1000 iterations were imposed for integration-mode reconstruction so to extensively analyze the convergence process. The Mean Absolute Error (MAE) was adopted as quantification metric to estimate the difference between the Ground Truth and the reconstructed image at each iteration of the reconstruction algorithm. Based on this, the percentage increase of the tissue models over the water ones was computed for both detector configuration. Furthermore, the same criterion was applied to list-mode configuration to compare the performances of the scattering spindle over the scattering line; for integration-mode configuration, this was done to measure the performances of the scattering cone against the nominal line but also for the scattering cones over the single cone model. Finally, every path estimation model was compared to the scattering line in water for list-mode configuration and to the nominal line for integration-mode since they represent the standard paradigm for trajectory estimation for the respective detector configurations.

Chapter 3

Results

3.1 Patient A

The cross-section images considered for Patient A are shown in Fig. 3.1.

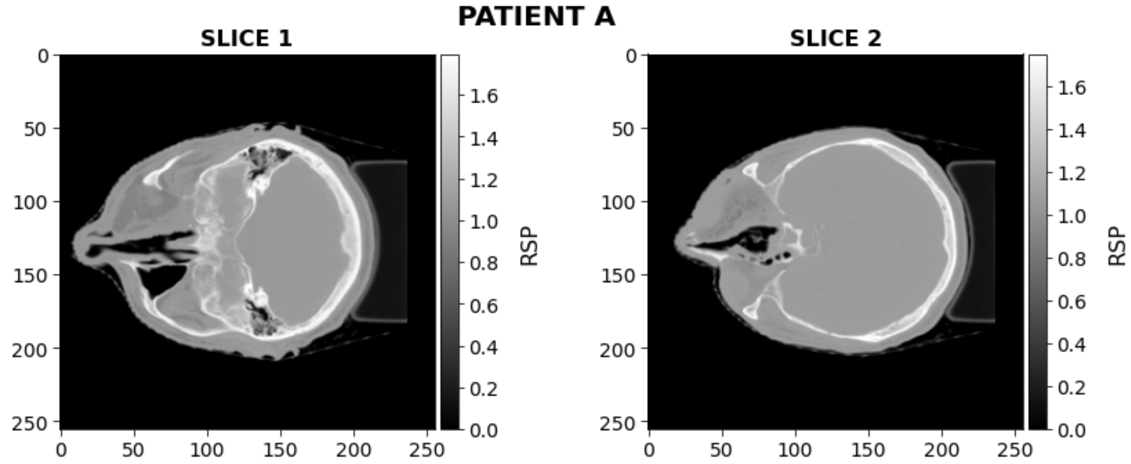

Figure 3.1: Ground Truth A. For patient A, the Ground Truth slices depicted in this figure were reconstructed.

3.1.1 Slice 1

- **List-mode**

List-mode images for slice 1 at different reconstruction steps are shown in Fig. 3.2.
The reconstruction process behaves similarly for the image reconstructed both in water and tissues. In particular, at the second iteration (Fig. 3.2a), the image is still affected by uncertainties and blurring effects, although it is already possible to make out the anatomical landmarks and the general shape of the slice. Moreover, Fig. 3.2b marks the stability of the image already at the 10-th iteration. Starting from iteration 20 (Fig. 3.2c), noise phenomena are more consistent and more evidently so at the last iteration (Fig. 3.2c) when streak artifacts and grainy features are more prominent.

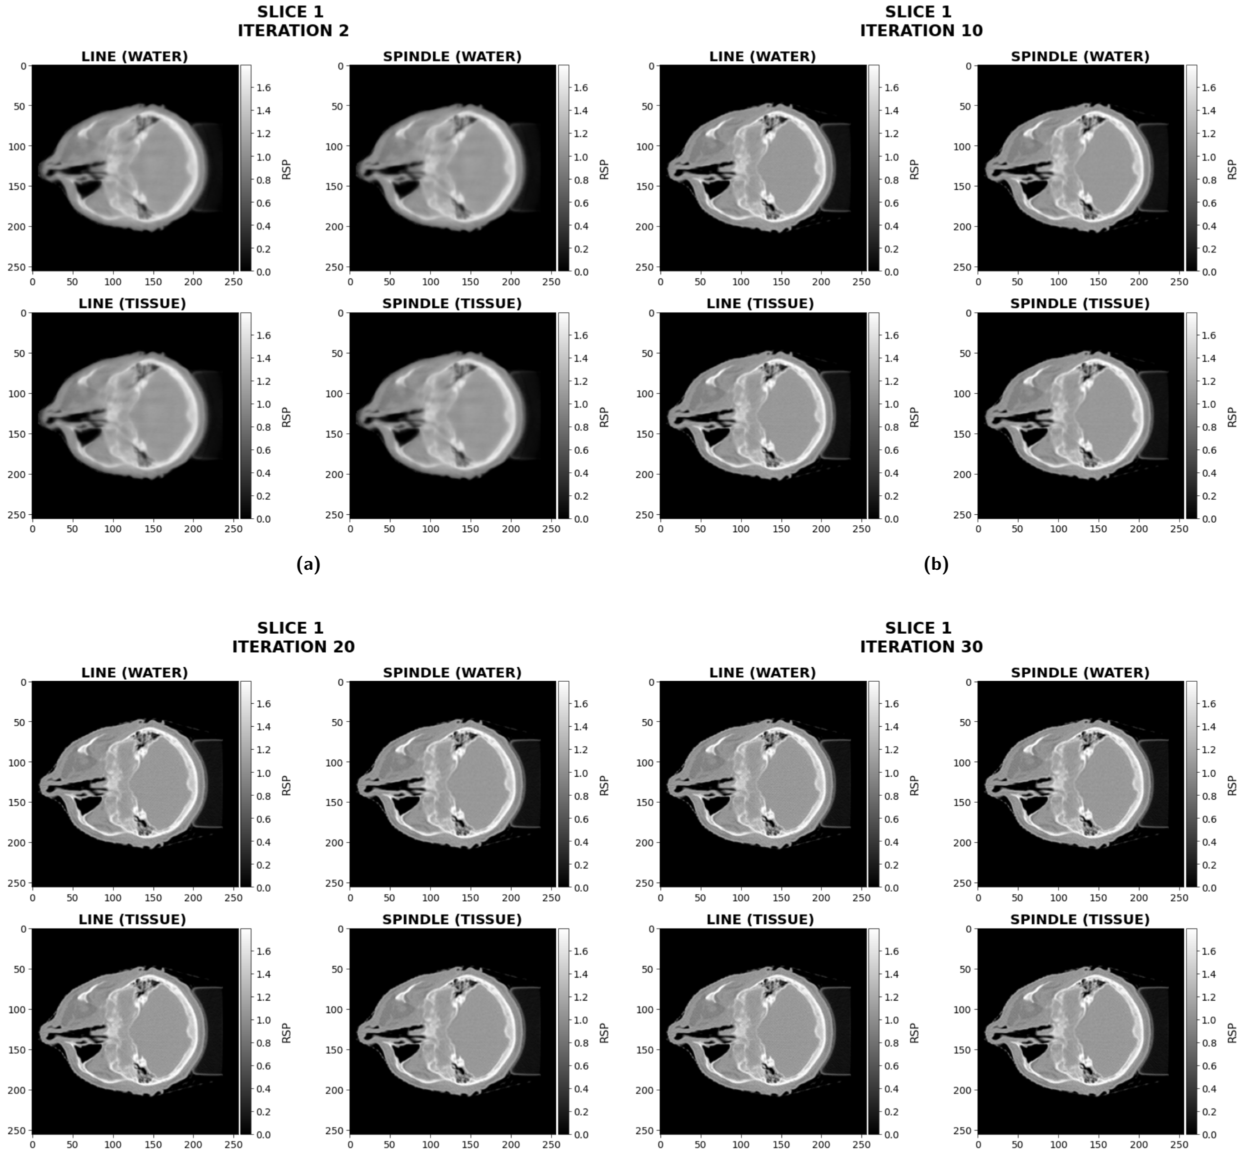

(a) (b)

(c) (d)

Figure 3.2: List Mode - Progression of the reconstruction algorithm for slice 1 of patient A. The blurring effects noticeable during the first iterations (a) tend to disappear as the image stabilizes already at iteration 10 (b) and 20 (c). Noise perturbation is minimal, but more pronounced towards the last iterations (d). The image shows acceptable quality both without and with prior knowledge, i.e. water-based and tissue-based scenario respectively.

The fitness of the specific models is better understood by looking at the quantification metric (Fig. 3.3). Though following a similar trend, the images reconstructed relying on prior knowledge about the scattering model provide overall a better quality than the water-based scenario. On an average 3% error introduced by the HU-RSP calibration in clinical scenarios (Gianoli et al. [2020]), the line model in tissues offers a relative improvement over the water line up to 1.07% at iteration 11 and 0.65% at iteration 30; an even greater performance in this regard can be observed for the spindle model, with an improvement in tissues over water of 2.27% at iteration 8 and 0.89% at iteration 30. As for the scattering spindle in tissues, a maximum percentage increase of 6.06% by iteration 30 over the scattering line in water can be noted.
Figure 3.3 also allows for a comparison between the two proposed models for proton-tracking devices. When the trajectories are reconstructed in water, the scattering spindle remains closer to the Ground Truth than the scattering line as the number of iterations increases, up to 5.22% by the 30-th iteration. Similarly for

the tissue-based scenario, iteration 30 sees a 5.45% relative improvement of the spindle over the line.

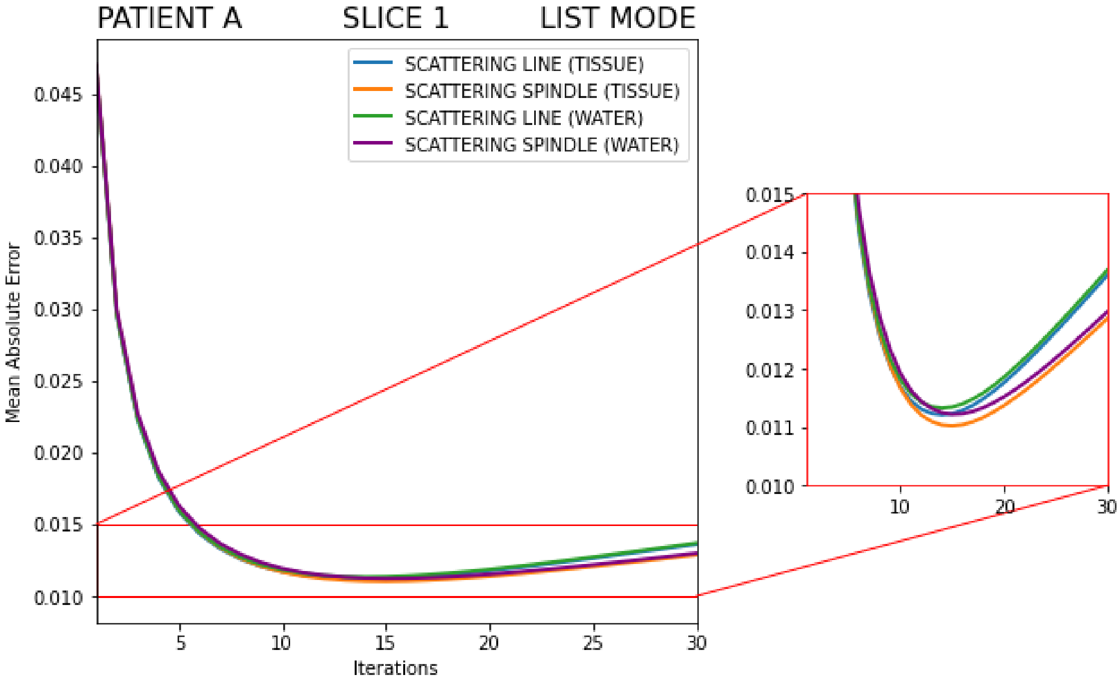

Figure 3.3: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient A. Though all models are able to provide for a generous quality in the reconstructed image, the scattering spindle models appear to be more accurate and stable than the scattering line ones as the number of iterations progresses. Reconstructing the images relying on prior knowledge about the scattering behavior from the X-ray CT files also helps in improving image quality.

The RSP relative improvements are summarized in Table 3.1. For both line and spindle, it can be noticed that the models in tissues are ever so slightly more exposed to noise amplification since the error relative to water tends to increase after reaching an optimal point. Nevertheless, comparing the two path estimation models demonstrates that the spindle in tissues possesses better performances than the one in water, providing not only a greater improvement relative to water but also a faster development over the line model between iteration 3 and 10, whereas this happens after iteration 10 for the water-based scenario.

Iterations	C1	C2	C3	C4	C5
2	0.19%	0.68%	-1.11%	-0.61%	-0.42%
10	1.05%	2.17%	-0.92%	0.22%	1.27%
20	0.80%	1.29%	2.85%	3.31%	4.09%
30	0.65%	0.89%	5.22%	5.45%	6.06%

Table 3.1: List Mode - Relative improvement for all case scenarios for slice 1 of patient A. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

- **Integration-mode**

Sampled stages of the SART-reconstructed images for slice 1 of patient A are reported in Fig. 3.4.

By iteration 50 (Fig. 3.4a), it is clear that the WET$_{mode}$ models are helpful in reproducing the general shape of the patient's geometry, but details and separation of the anatomical structures suffer from poor

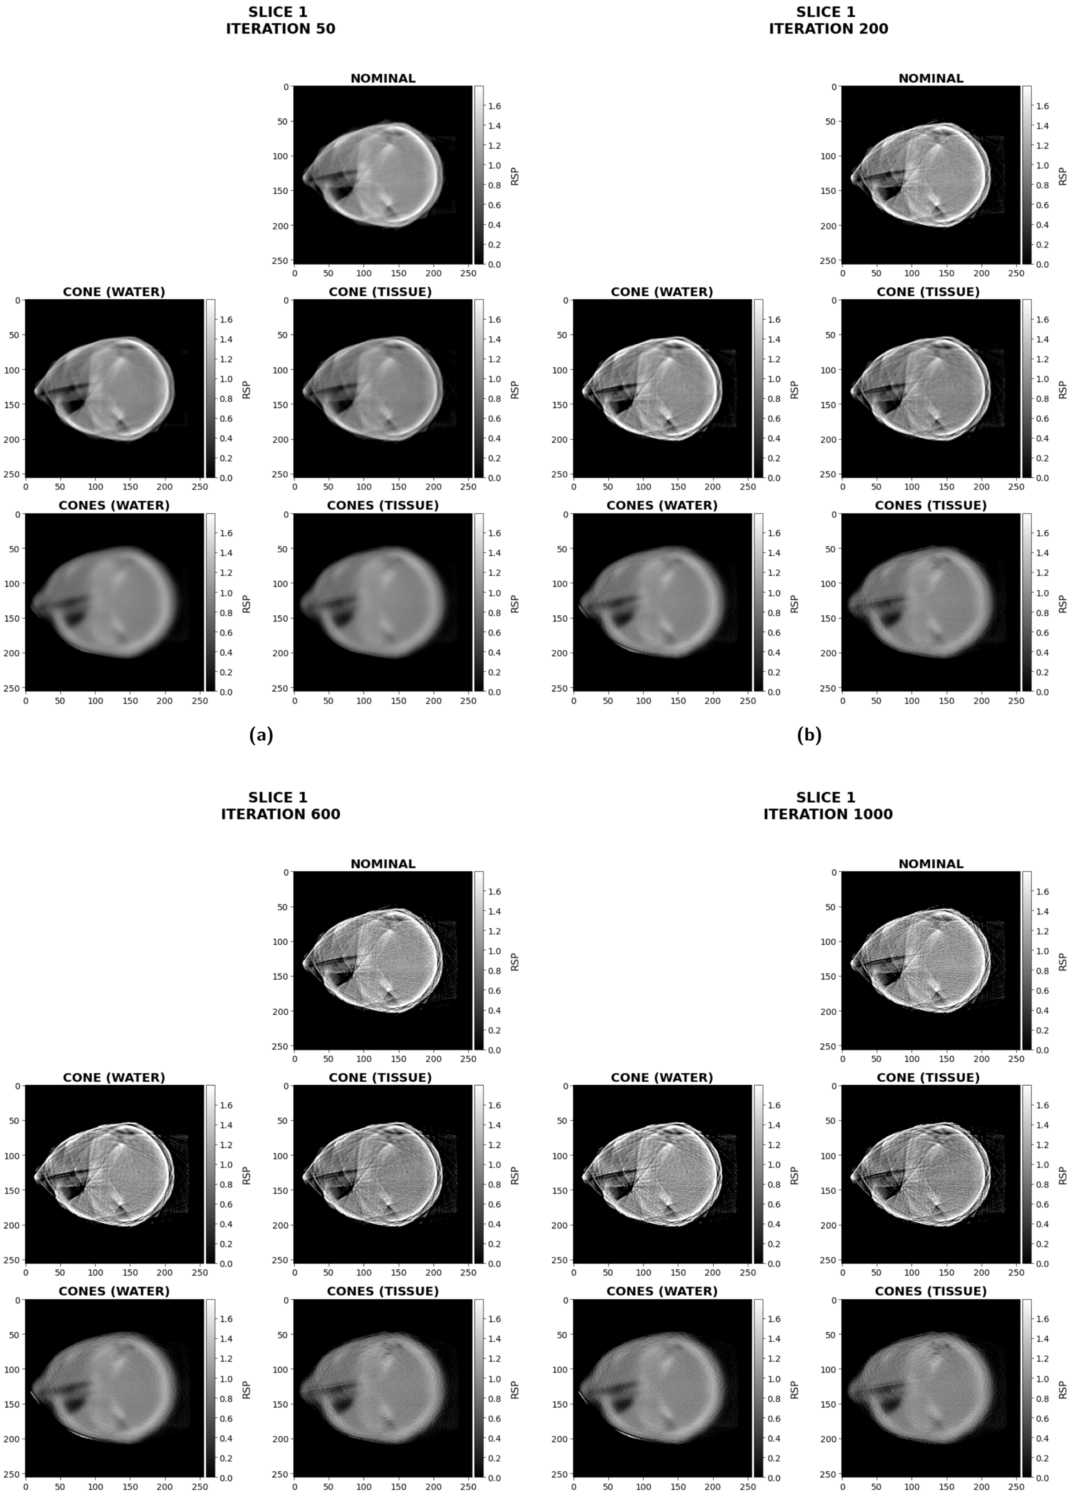

(a) (b) (c) (d)

Figure 3.4: Integration Mode - Progression of the reconstruction algorithm for slice 1 of patient A. The scattering cones images offer a higher spatial resolution and preservation of anatomical details when compared to the WET$_{mode}$ models. However, as evident streak line artifacts can be observed for the nominal line and the scattering cone both in water and tissues, edges artifacts and graininess affect the scattering cones images as the number of iterations increases.

spatial resolution, appearing extremely segmented. Furthermore, the borders of the head district are not properly identifiable as little protrusions can be observed against the overall rounded shape of the skull. Even if noticeably blurred, the scattering cones images are more precise in rendering the natural delimitations and transitions in the considered anatomy, although edge artifacts can already be seen for the water-based scenario. As the tomographic reconstruction progresses, noise break-up begins to affect the images: at iteration 200 (Fig. 3.4b), streak line artifacts compromise image quality, especially for the WET$_{mode}$ models and, to a smaller extent, also for the scattering cones in tissues. By iteration 1000 (Fig. 3.4d), these effects become heavier for the nominal line and the scattering cone, while the scattering cones seem to be more sensitive to graininess and edge artifacts, though still preserving an acceptable rendering of the cross-section geometry.

Fig. 3.5 allows to explore the behavior of these models in a quantitative manner comparing the reconstructed images to the respective Ground Truth by MAE.

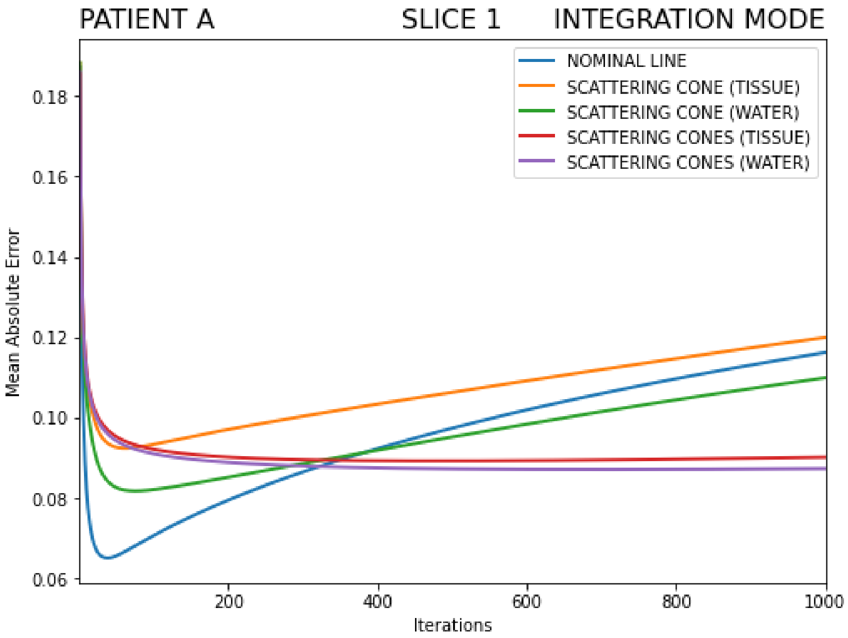

Figure 3.5: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient A. The faster convergence of the nominal line is compensated by a higher robustness to noise of the scattering cone in water, while the scattering cone in tissues show a higher error than both of these models. The scattering cones are instead able to outperform the WET$_{mode}$ models with a slower but better convergence. Also in this case, the water-based scenario prevails over the tissue one.

Being the simplest model, the nominal line is characterized by the fastest convergence, reaching a minimum error of 0.065 at iteration 38. However, this also results in a more concrete sensitivity to noise: in fact, the nominal line is outperformed by the water cone at iteration 381, leading to a relative improvement of 5.42% by the last iteration. The tissue cone is instead unable to compete with the other WET$_{mode}$ models, with a difference of 3.20% in favor of the nominal line and 9.11% in favor of the water cone by the last iteration. The scattering cones show specific traits to be attributed to the inclusion of the entire WET spectrum during tomographic reconstruction. In fact, they converge slower than the other models, but they also present lower error as the number of iterations progresses, i.e. higher stability to noise. In the water-based scenario, starting with a 14.15% disadvantage at iteration 45, the scattering cones are then able to outperform the single cone with a relative improvement of 20.58% at iteration 1000. Similarly, the tissue cones see a difference of 3.66% in favor of the respective single cone at iteration 33 only to be followed by a trend reversal at iteration 84, leading to a relative improvement of 24.82% by the last iteration. Furthermore, the scattering cones are able to provide better performances than the nominal line already at iteration 322 for the water-based scenario and at iteration 348 for the tissue-based scenario, ending up with a relative improvement of 24.88% and 22.42% respectively by the last iteration. However, when comparing the two reconstruction scenarios, the scattering cones in water represent the dominant model, with a 3.28%

difference over the tissue cones at iteration 1000.
Table 3.2 reports the percentage difference among all proposed models.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-12.13%	-1.10%	-25.82%	-41.09%	-14.08%	-2.85%	-43.53%	-45.11%
200	-13.99%	-1.49%	-7.13%	-22.12%	-4.34%	7.10%	-11.79%	-13.45%
600	-10.95%	-2.51%	3.46%	-7.12%	11.46%	18.20%	14.52%	12.38%
1000	-9.11%	-3.28%	5.42%	-3.20%	20.58%	24.82%	24.88%	22.42%

Table 3.2: Integration Mode - Relative improvement for all case scenarios for slice 1 of patient A. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

As a last remark, though the scattering cone in tissues shows a higher error than the nominal line throughout the entire reconstruction process, the relative difference seems also to be decreasing as the number of iterations increases, hinting that, however marginally, the tissue cone may still behave better against noise.

3.1.2 Slice 2

- **List-mode**

Fig. 3.6 displays the list-mode reconstructed images for slice 2 of patient A at sampled stages of the reconstruction process. Once again, list-mode configuration provides a generous image quality for all studied models. As the blurring effect of the early iterations (Fig. 3.6a) dissipates, the reconstructed images stabilize towards the convergence point. Already at iteration 10 (Fig. 3.6b), noise artifacts in the shape of light streak lines can be noticed across the scattering line images, while the spindle model seems to be more consistent. Nevertheless, by iteration 20 (Fig. 3.6c), streak lines begin to afflict more significantly also the images reconstructed relying on the Gaussian probabilistic envelop. Even with noise break-up attributes, the anatomical structures are still notably defined and spatial resolution in all four cases is well preserved (Fig. 3.6d).

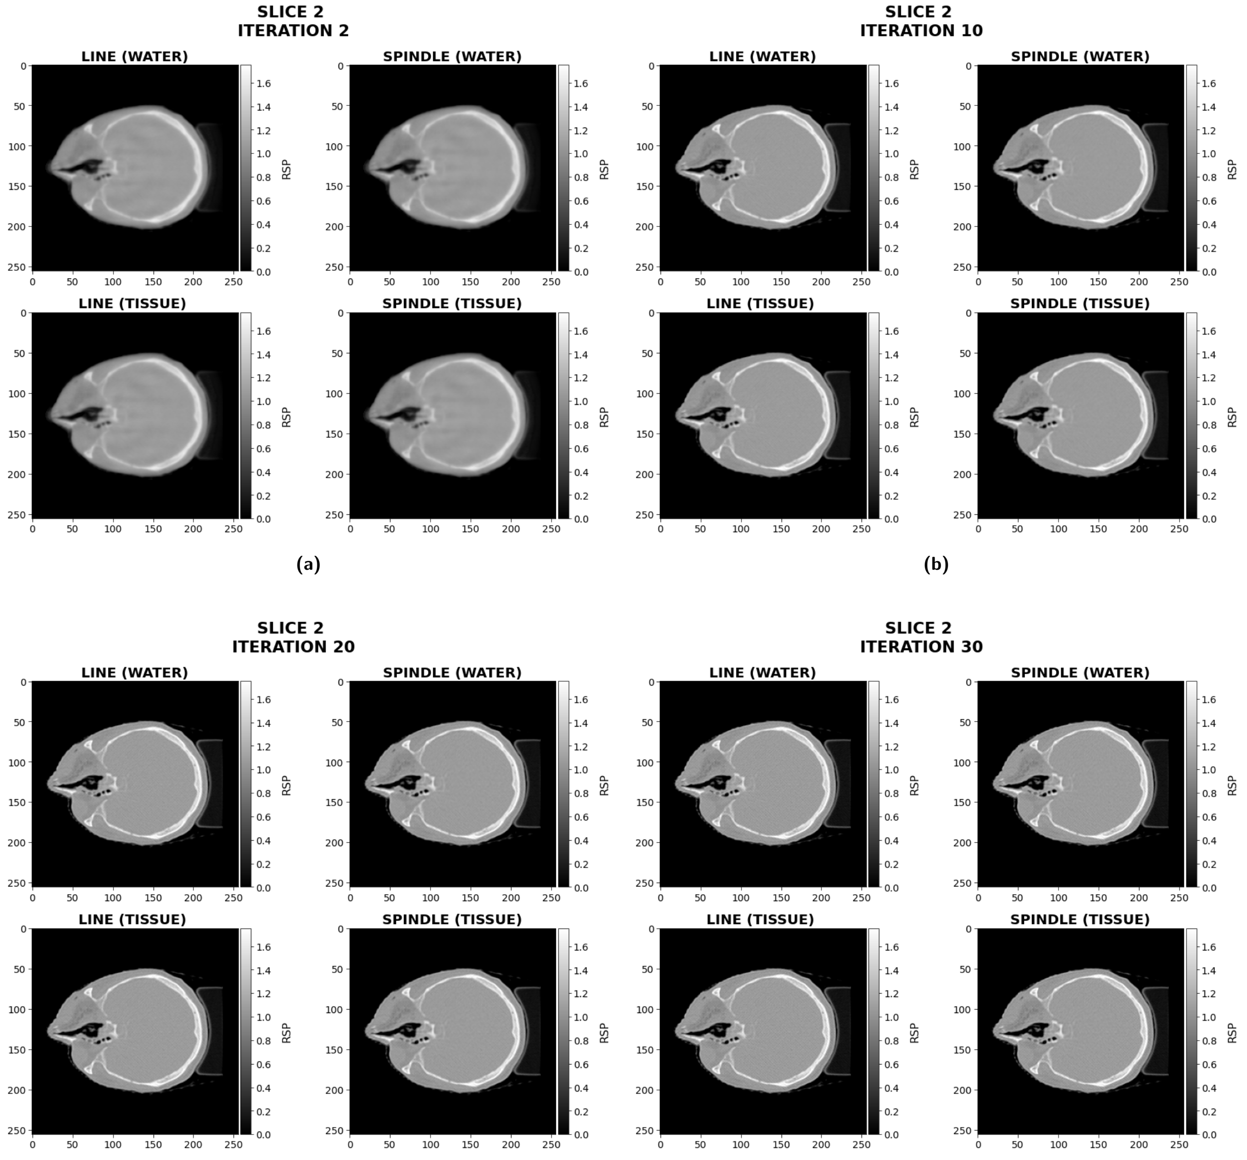

(a) (b) (c) (d)

Figure 3.6: List Mode - Progression of the reconstruction algorithm for slice 2 of patient A. Despite typical blurring effects observable for early iterations (a), the converge of the reconstruction algorithm establishes well-defined images with good spatial resolution in all four models, as shown in (b) and (c). The rise of noise break-up markings is visible towards the last iterations (d).

As shown in Fig. 3.7, differences between the Ground Truth and the reconstructed images are to be taken into account. For this rather uniform slice, the scattering line modeled in tissues has slightly better performances than the water line: a 0.0097% difference establishes the tissue line as the superior model already at iteration 1, peaking with a 0.82% at iteration 12. The scattering spindle, however, fails in proving a better image quality in tissues rather than water: in fact, by iteration 30, the water model has a 2% relative advantage over the tissue one. Nevertheless, it is expected that the introduction of the Gaussian envelop for image reconstruction would lead to a superior model than the scattering line. Even with noise break-up perturbations, the scattering spindle provides a 10.30% and a 8.23% maximum improvement at the last iteration in water and tissues respectively. Furthermore, starting from a 0.18% predominance at iteration 2 of the water line over the tissue spindle, a trend reversal can be seen leading a 8.50% relative improvement by the last iteration.

The RSP relative improvements for this case study are summarized in table 3.3.

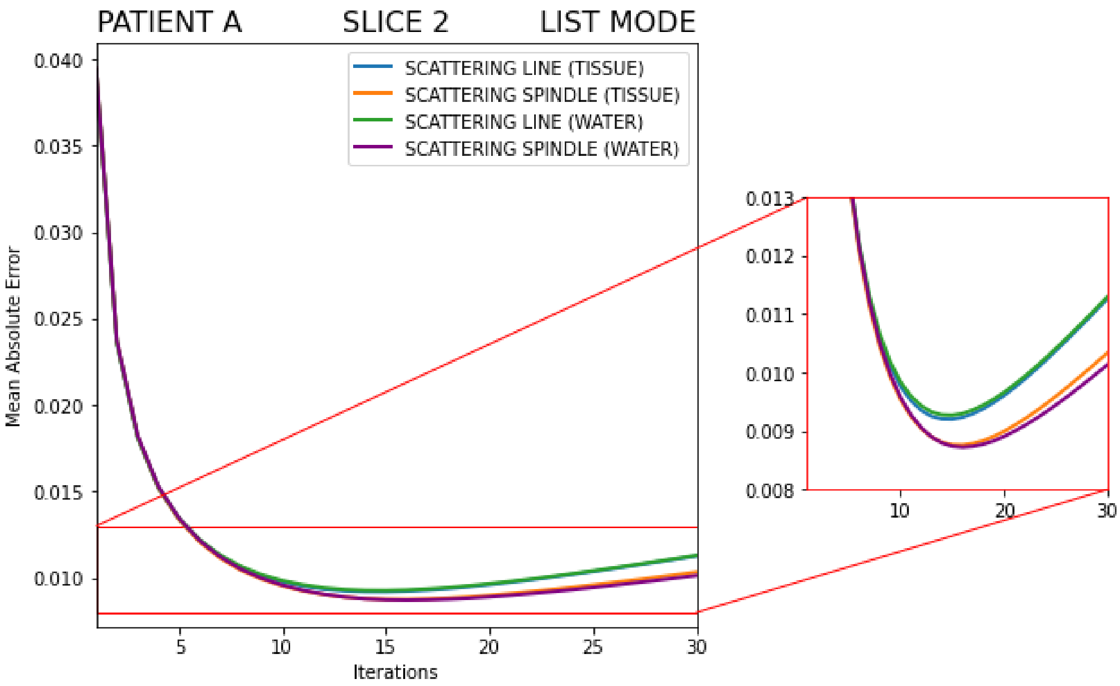

Figure 3.7: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient A. The scattering spindle model is able to outperform the scattering line model already at early iterations and appears to be more stable to noise. However, tissue spindle fails in providing a superior image quality over the water spindle, whereas the tissue line offers a marginal, but quantifiable, improvement over the water line.

Iterations	C1	C2	C3	C4	C5
2	0.01%	0.31%	-0.49%	-0.27%	-0.18%
10	0.78%	0.42%	-2.69%	2.34%	3.10%
20	0.57%	-0.98%	7.83%	6.39%	6.93%
30	0.30%	-2.00%	10.30%	8.23%	8.50%

Table 3.3: List Mode - Relative improvement for all case scenarios for slice 2 of patient A. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

It is interesting to notice that the tissue spindle offers marginal improvement over the water spindle during early iterations, reaching a peak of 0.66% at iteration 4 before starting to be outperformed.

- **Integration-mode**

Concerning integration-mode detector configuration, the reconstructed images for slice 2 are displayed in Fig. 3.8.

The models representing the WET$_{mode}$ of the histogram tend to be less accurate in terms of anatomical details. In fact, at iteration 50 (Fig. 3.8a), the general shape of the slice and the reproduction of bony structures, as well as the extent of the cavities, are more defined for the scattering cones models. Nevertheless, these latter images do suffer from a heavy blurring effect that does not relate to the nominal and the cone model. It is also possible to notice that the WET$_{mode}$ models fail in a precise reconstruction of the anatomical borders which appear to be afflicted by fluctuations and discontinuities. As the number of

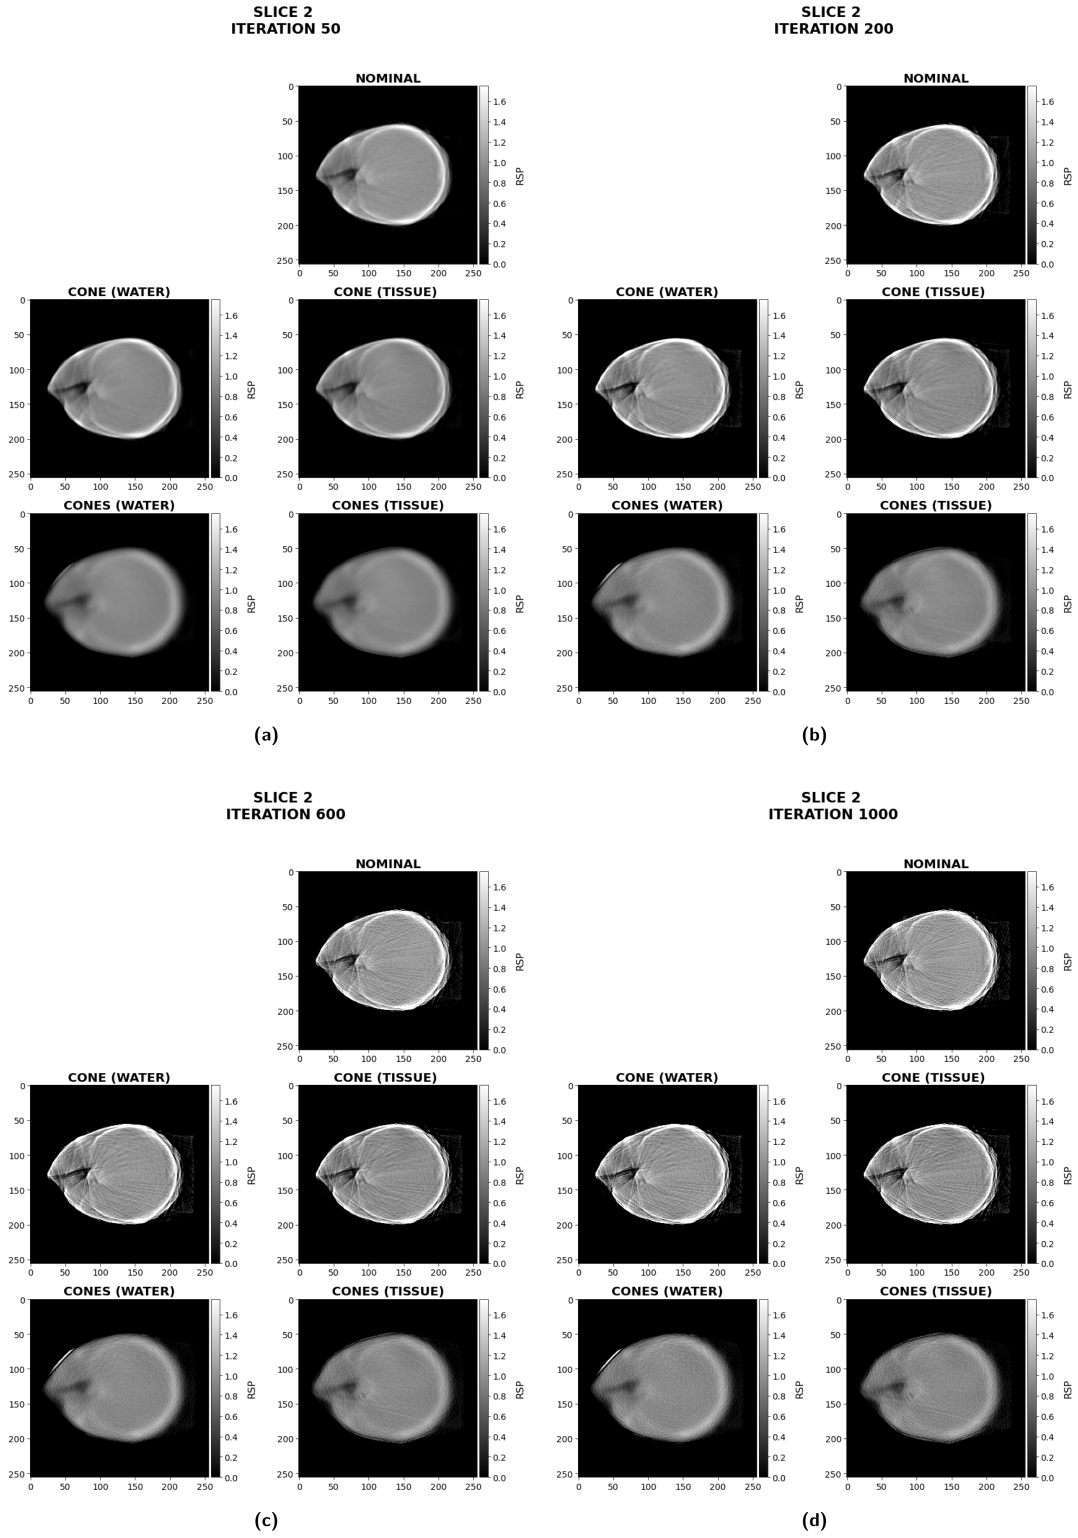

(a)

(b)

(c)

(d)

Figure 3.8: Integration Mode - Progression of the reconstruction algorithm for slice 2 of patient A. Though the images reconstructed according to the scattering cones model are more blurred and smeared, they offer a higher match to Ground Truth in terms of reproduction of the anatomical structures and their borders. Noise break-up effects appear as streak lines in the nominal and cone models, while edge artifacts and graininess is more prominent for the scattering cones model.

iterations progresses, the noise break-up can be observed throughout all tested models. By iteration 1000 (Fig. 3.8d), the nominal and cone models are indeed perturbed by streak artifacts radiating according to precise patterns, although they seem slightly less prominent in the second case than in the first. Noise artifacts appear to be more moderately amplified throughout the reconstruction process for the scattering cones, and it primarily manifests as graininess rather than streaking lines; on the other hand, edge artifacts start to become more relevant already after iteration 200 (Fig. 3.8b).
The comparison of the reconstructed images against the Ground Truth (Fig. 3.9) highlights how the tissue-based models do not allow for the same image quality as the water-based models.

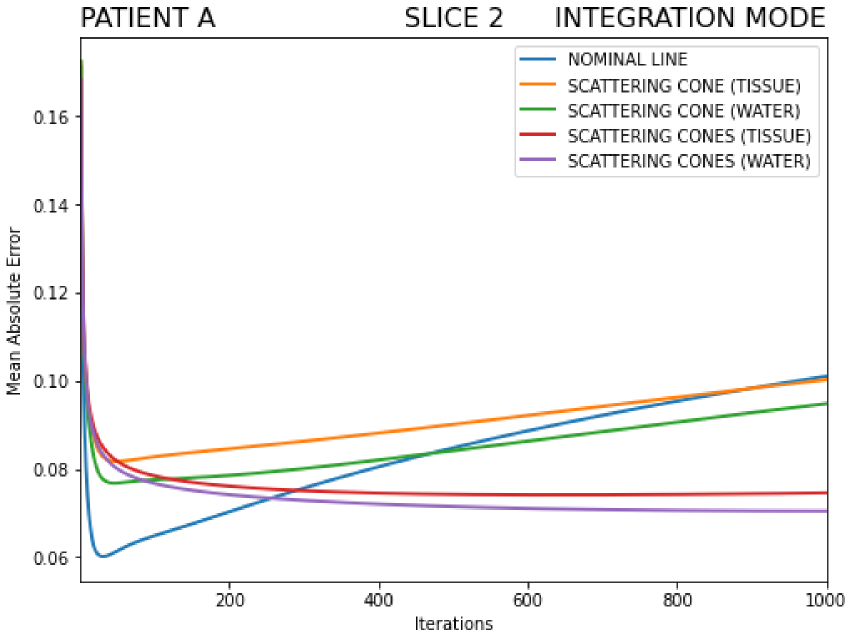

Figure 3.9: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient A. Though the prior knowledge cases seem to fail against the water models, the scattering cone in water is able to outperform the nominal line model after approximately 500 iterations, while the trend reversal is slower in the tissue-based scenario. Moreover, the reconstruction of the images relying on the entire WET spectrum is able to offer a lower relative error in comparison with the WET$_{mode}$ counterparts.

Considering the WET$_{mode}$ cases, the scattering cone in water prevails over the one in tissues, the highest difference being 7.79% at iteration 246. A similar trend, though slightly less pronounced, can be observed for the scattering cones, where the water model outperforms the tissue model with a maximum disparity of 5.82% at the iteration 1000: since the two models are still diverging from one another by the last iteration, this implies that the water cones are also more stable to noise than the tissue cones.
The MAE quantification also illustrates that the scattering model presents some limitations. The nominal line model is superior to the scattering cone model until approximately iteration 500, with a predominance of 32.45% at iteration 18 and 39.89% at iteration 20 in water and tissues respectively. The water cone is however able to outperform the nominal model as the number of iterations progresses, reaching the largest distance of 6.15% at iteration 1000. The tissue cone, instead, appears to be more sensitive to noise effects, thus limiting its behavior against the nominal line to a maximum of 0.76% at iteration 1000.
The use of the entire WET histogram during reconstruction helps in mitigating this problems, lowering the relative error with the Ground Truth and noise sensitivity. Considering a maximum predominance of the scattering cone over the scattering cones in water of around 7.47% at iteration 26, the latter model is able to outweigh the single cone starting from iteration 84 and reach a peak difference of about 25.65% at iteration 1000. The tissue cases behave the same way: the scattering cones overtake the single cone model around iteration 50 onwards and reach the largest distance (around 25.60%) by the last iteration, after the scattering cone showed a 2.95% dominance at iteration 25. The scattering cones are also able to outperform the nominal line: despite an initial maximum difference of 41.07% and 43.66% at iteration 21 in favor of the

nominal line in water and tissues respectively, the scattering cones then present a relative improvement of 30.22% in water and 26.16% in tissue by iteration 1000.
Table 3.4 offers a more compact view of these results for all considered models.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-6.42%	-2.11%	-24.99%	-33.01%	-4.22%	0.005%	-30.26%	-33.01%
200	-7.69%	-2.60%	-11.71%	-20.30%	5.58%	10.04%	-5.47%	-8.22%
600	-6.75%	-4.32%	2.63%	-3.94%	17.68%	19.55%	19.84%	16.38%
1000	-5.74%	-5.82%	6.15%	0.76%	25.65%	25.60%	30.22%	26.16%

Table 3.4: Integration Mode - Relative improvement for all case scenarios for slice 2 of patient A. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

It is worth noting that, though always negative, the difference between the scattering cone in water and tissues seems to start decreasing after iteration 200, which may suggest a slower but increasing instability of the water cone to noise for a large number of iterations.

3.2 Patient B

The cross-section images considered for Patient B are shown in Fig. 3.10.

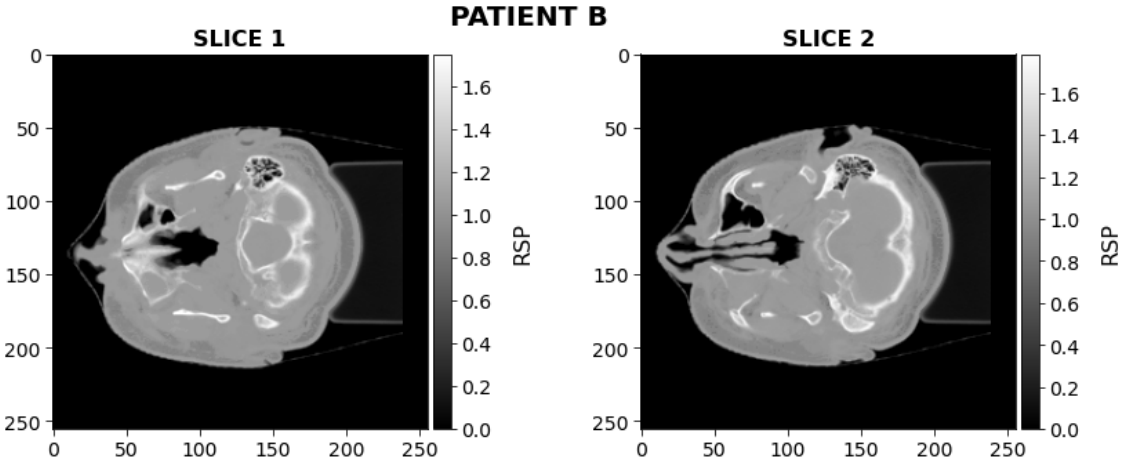

Figure 3.10: Ground Truth B. For patient B, the Ground Truth slices depicted in this figure were reconstructed.

3.2.1 Slice 1

- **List-mode**

Different steps for the reconstruction of slice 1 of patient B are reported in Fig. 3.11. Starting from a blurred rendering (Fig. 3.11a), the OS-SART algorithm quickly provides for a good RSP quality by iteration 10 (Fig. 3.11b). The overall anatomy is in fact well-preserved when compared to the Ground Truth: bony structures are highly detailed and the borders between different tissues are accurately reproduced. Particularly, it is possible to observe an optimal contrast when moving from the central part to the anatomy's border, where the difference between lighter and darker regions is quite evident in marking a distinction between various compositions.

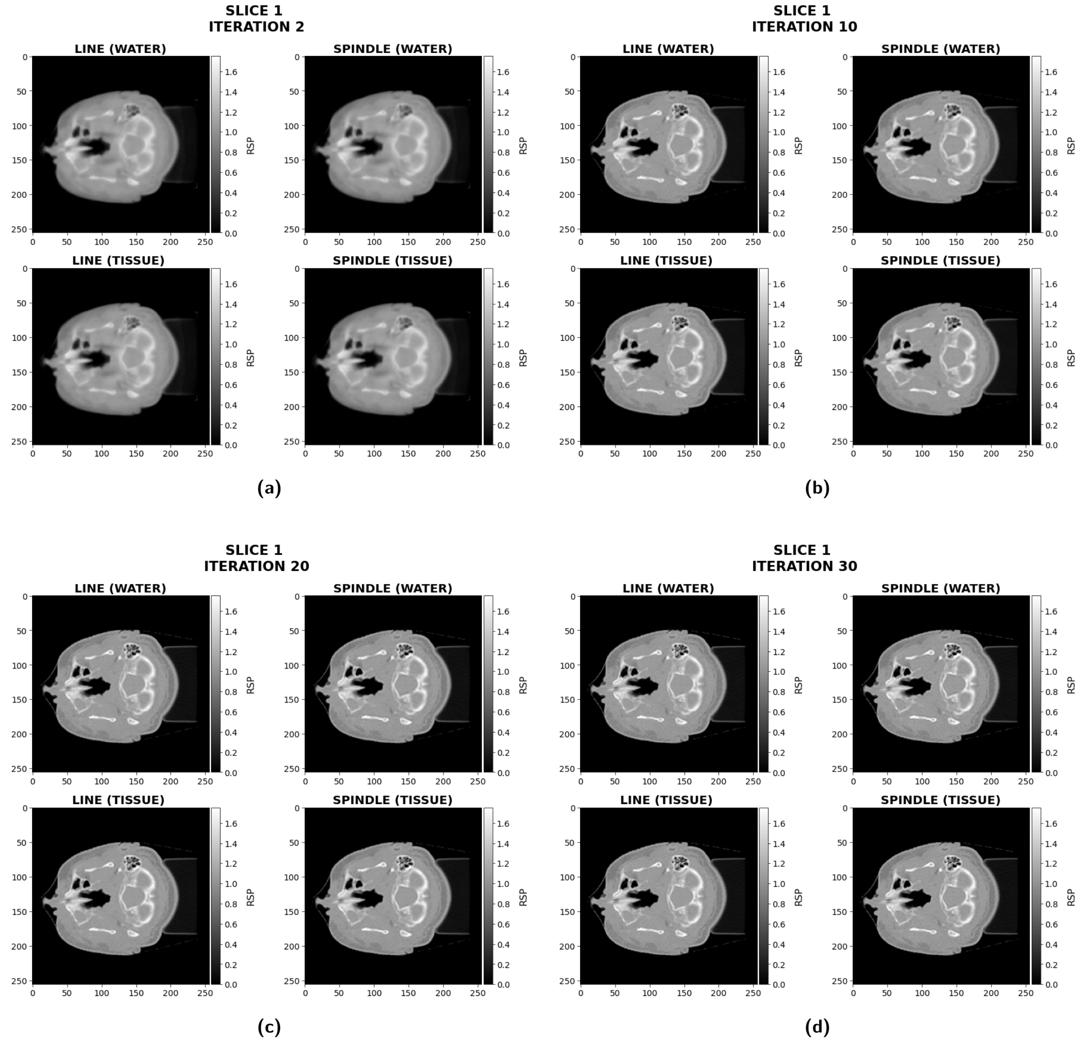

Figure 3.11: List Mode - Progression of the reconstruction algorithm for slice 1 of patient B. Each tested model is able to reproduce the considered anatomy with great accuracy in terms of spatial resolution and contrast details. Even with noise break-up effects, the reconstructed images are still highly informative about the nature of the depicted anatomical structures.

By iteration 20 (Fig. 3.11c), noise amplification starts to take hold since the images seem to appear slightly more pixelated. This is more noticeable at the last iteration (Fig. 3.11d), with streak line artifacts becoming more apparent. However, even with noise break-up, the reconstructed images are able to preserve an acceptable quality and spatial resolution, leaving air cavities, soft tissues and bony structures still entirely recognizable.

The MAE quantification metric applied between the reconstructed images and the Ground Truth references (Fig. 3.12) is able to highlight the behavior of the tested models. On this detailed anatomical section, the scattering line in tissues shows better performances than the one in water, starting with a 0.11% relative improvement already at iteration 1, up to a maximum of 1.60% by iteration 8. As for the scattering spindle, the trend is similar: it is possible to observe a minimum difference of 0.54% at the first iteration and a maximum of 4.07% at iteration 9. In water, the scattering spindle shows a higher error than the scattering line until iteration 17, where a trend reversal can be seen, leading to a relative improvement of 3.85% by the last iteration. On the other hand, the tissue line is predominant over the tissue spindle until iteration 11: after that, the spindle shows better performances with a relative improvement of 4.02% at iteration 30. The scattering spindle in tissues starts with a disadvantage of 0.17% at the first iteration, which then turns into a 4.47% maximum improvement over the water line by the last iteration.

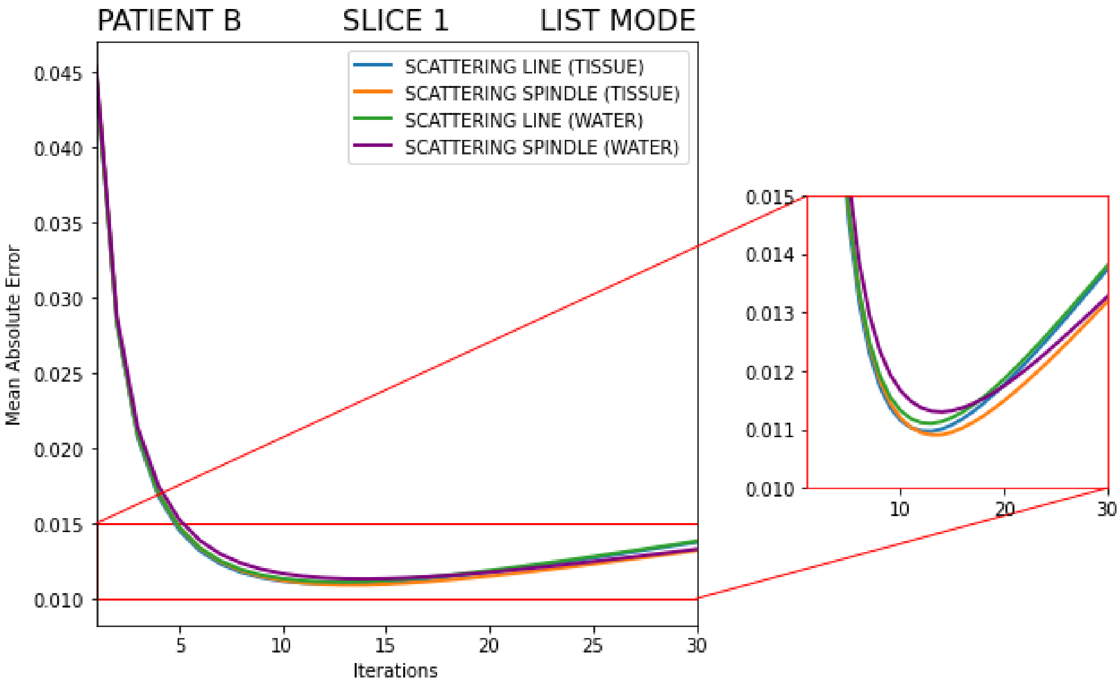

Figure 3.12: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient B. Given the prior knowledge about the scattering behavior of the individual protons, the tissue models are able to provide for a lower error than the water models do. Both in water and tissues, the scattering spindle show more stability to noise than the scattering line, with the tissue spindle leading to an overall better performance than all the other models.

Table 3.5 summarizes all of these findings.

Iterations	C1	C2	C3	C4	C5
2	0.52%	1.33%	-1.48%	-0.65%	-0.13%
10	1.50%	4.03%	-3.01%	-0.37%	1.14%
20	0.84%	2.17%	0.99%	2.32%	3.14%
30	0.46%	0.64%	3.85%	4.02%	4.47%

Table 3.5: List Mode - Relative improvement for all case scenarios for slice 1 of patient B. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

These results highlight how the scattering spindle has better performances against noise break-up than the scattering line does both in water and tissues. Furthermore, the tissue spindle offers a general lower error than all other models tested on this cross-section.

- **Integration-mode**

For integration-mode configuration, the output of the SART at different iterative steps can be seen in Fig. 3.13.

The WET_{mode} models are characterized by lack of detail accuracy and poor spatial resolution when it comes to the definition of anatomical boundaries, with edge artifact already visible by iteration 50 (Fig.

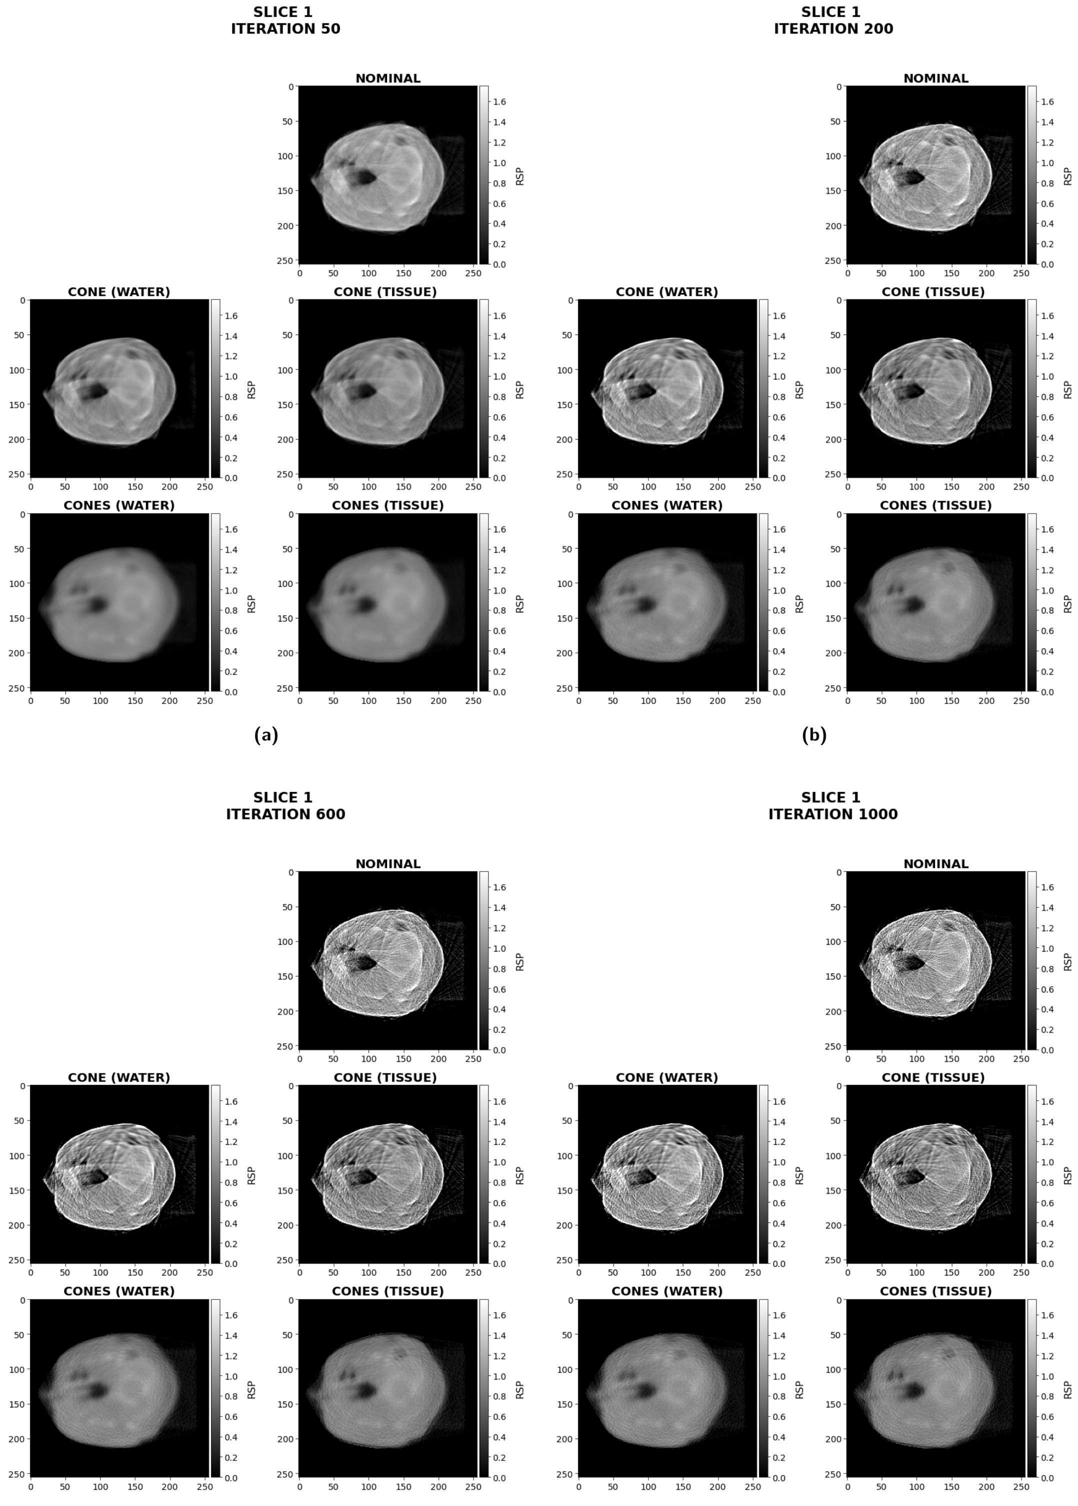

(a) (b) (c) (d)

Figure 3.13: Integration Mode - Progression of the reconstruction algorithm for slice 1 of patient B. The scattering cones images appear less sensitive to noise effects and more capable of reproducing anatomical details. The images reconstructed based on the WET_{mode} can provide information about the general shape of the cross-section, but poor spatial resolution and heavy noise break-up are still limiting factors.

3.13a). At the same time, the scattering cones seem to more efficiently reproduce the anatomy at hand, with more distinct bony structures and well-defined borders, although a blurred appearance can be easily observed. As the reconstruction process progresses, noise break-up effects become more and more prominent, leading to consistent streak lines by iteration 1000, especially for the nominal line (Fig. 3.13d). Though less impactful, graininess and edge artifacts can be pinpointed also for the scattering cones images, even if anatomical geometries are still identifiable.
The MAE quantification between the reconstructed images and the Ground Truth images is shown in Fig. 3.14.

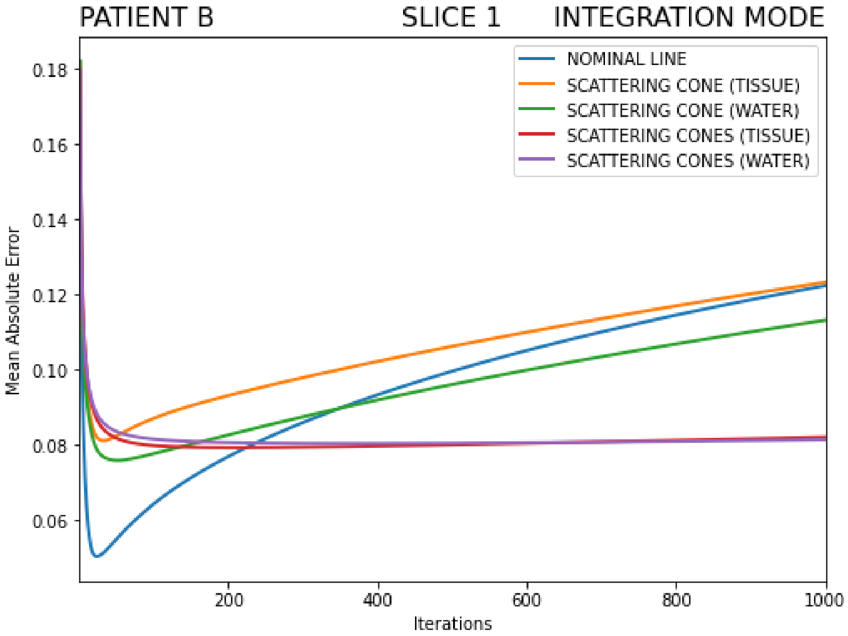

Figure 3.14: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient B. Though the prior knowledge cases seem to fail against the water models, the scattering cone in water is able to outperform the nominal line model after approximately 500 iterations, while the trend reversal is slower in the tissue-based scenario. Moreover, the reconstruction of the images relying on the entire WET spectrum is able to offer a lower relative error in comparison with the WET_{mode} counterparts.

Concerning the WET_{mode} models, the nominal line has the fastest convergence (and, thus, the fastest noise break-up), with a maximum difference from the water cone and the tissue cone of roughly 60% at iteration 18 and 64.98% at iteration 19 respectively. However, the scattering cone in water starts to prevail over the nominal line at iteration 348, leading to a relative improvement of 7.49% by iteration 1000, whereas the scattering cone in tissues fails against both of these models in providing a lower error, with a maximum distance of 12.73% at iteration 160 from the water cone. The scattering cones show better convergence and robustness to noise, implying superior performances than the single cone both in water and tissues: specifically, the relative improvement at iteration 1000 is 28.08% in the first case and 33.44% in the second one. These attributes can be observed also for a comparison with the nominal line: though the nominal line prevail with a maximum difference of 77.53% in water and 75.20% in tissues by iteration 19, the scattering cones is able to reach a 33.47% and 32.95% relative improvement by the last iteration in water and tissues respectively. It is worth noting that, though providing better performances at early iterations (highest improvement of 2.19% at iteration 65), the scattering cones in tissues are outperformed by the water-based scenario already at iteration 677, with a maximum difference of 0.79% by the last iteration.
Table 3.6 summarizes these results.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-8.57%	2.16%	-38.12%	-49.96%	-10.15%	0.74%	-52.13%	-48.85%
200	-12.61%	1.68%	-7.54%	-21.10%	2.48%	14.86%	-4.87%	-3.11%
600	-10.13%	0.22%	4.98%	-4.64%	19.33%	26.91%	23.35%	23.52%
1000	-8.90%	-0.79%	7.49%	-0.75%	28.08%	33.44%	33.47%	32.95%

Table 3.6: Integration Mode - Relative improvement for all case scenarios for slice 1 of patient B. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

3.2.2 Slice 2

- **List-mode**

In Fig. 3.15, the output of the OS-SART algorithm can be observed at different steps for all proposed models. Even before convergence (Fig. 3.15a), the reconstructed images provide a good rendering of the considered anatomy, which can be better observed by iteration 10 (Fig. 3.15b) when blurring effects due to convergence of high frequencies during early iterations are no longer present. The depicted cross-sections offer optimal spatial resolution and image quality: anatomical structures are clearly defined, where details such as complex bony structures and air cavities are clearly detectable. By iteration 20 (Fig. 3.15c), noise effects are slightly noticeable but the overall image quality is well-preserved: the anatomical borders of the different structures remain intact and the variation between different tissue composition is highlighted by an evident contrast. Noise break-up afflicts more heavily the images at iterations 30 (Fig. 3.15d) in terms of streak line artifacts, but the rendered image quality is still considerably acceptable, since no consistent deterioration of the proposed slice is observable and a good contrast still marks the different anatomical regions.

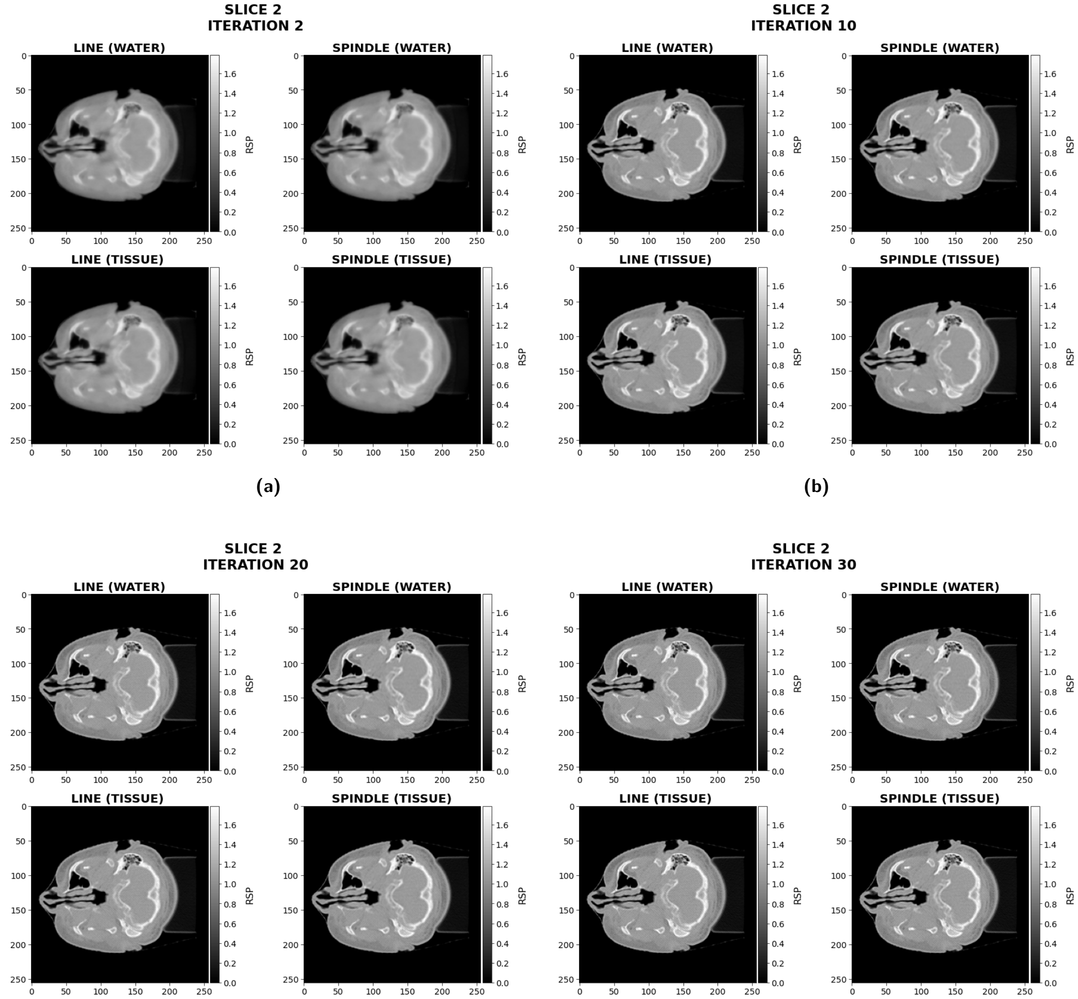

(a) (b) (c) (d)

Figure 3.15: List Mode - Progression of the reconstruction algorithm for slice 2 of patient B. In all tested models, anatomical details and spatial resolution are well-preserved and different tissue composition is rendered by an optimal contrast quality. Noise break-up effects are present towards last iterations, but their contribution is marginal.

A more precise analysis of the behavior of these models is reported in Fig. 3.16, where the MAE quantification was used to study the performances of the reconstructed images against the respective Ground Truth. Although all models are able to provide a low error when compared to the Ground Truth, distinctions can be made. Particularly, the scattering line in tissues offers better performances than the scattering line in water already at iteration 1 with a 0.07% improvement that peaks at iteration 9 with a relative improvement of 1.57%. The scattering spindle in water shows instead a higher error when converging (5.32% difference in favor of the scattering line in water at iteration 8), but better robustness to noise as highlighted by a 2.40% improvement by the last iteration. The scattering spindle in tissues has overall better performances when compared to the other models. In fact, considering a disadvantage of 0.21% at iteration 5 against the water line, this latter model is soon outperformed at iteration 7 with a relative improvement of 0.0028% leading to one of 4.29% by the last iteration. Furthermore, the scattering spindle in tissues has a slower convergence than the tissue line, with a 1.51% disadvantage at iteration 6, only to see a trend reversal leading to 3.67% relative improvement at iteration 30. Finally, a relative improvement of 5.65% already iteration 10 marks the predominance of the tissue spindle over the water one.

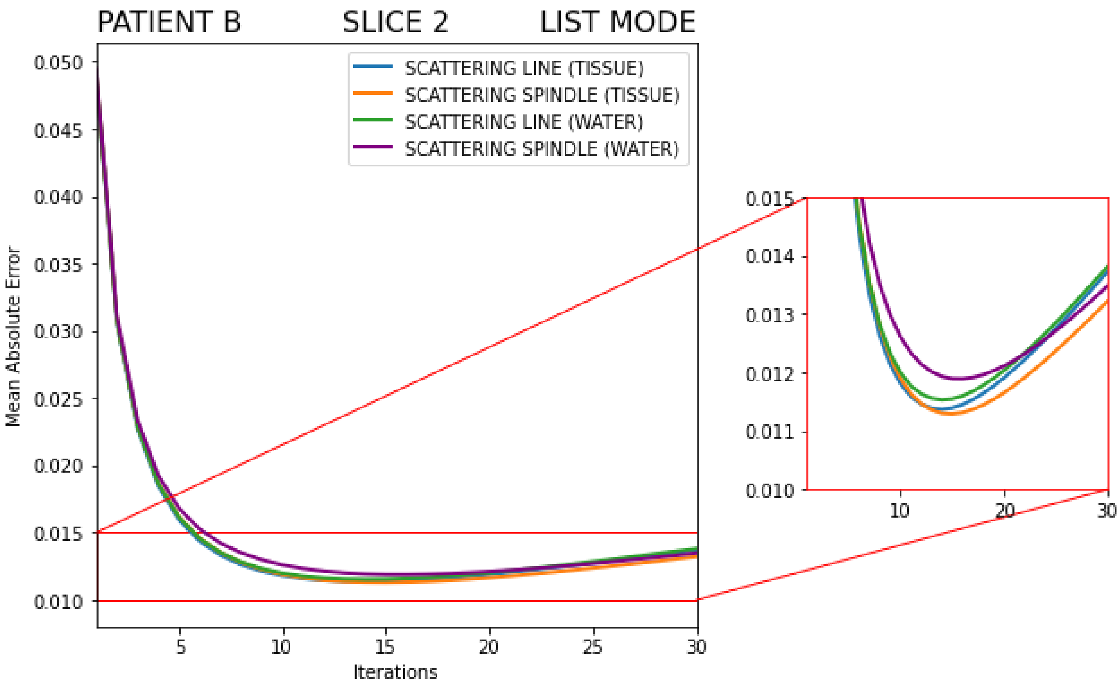

Figure 3.16: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient B. The scattering spindle in tissues has an overall better performance than all other proposed models, with a general lower error and higher stability to noise. The scattering spindle in water has a higher error during convergence, but offers more robustness against noise than the two line models. All models are however capable of reproducing the Ground Truth with optimal accuracy.

These considerations are all summarized in table 3.7.

Iterations	C1	C2	C3	C4	C5
2	0.37%	1.08%	-1.24%	-0.51%	-0.14%
10	1.57%	5.65%	-5.11%	-0.75%	0.83%
20	1.04%	3.79%	-0.66%	2.14%	3.16%
30	0.64%	1.93%	2.40%	3.67%	4.29%

Table 3.7: List Mode - Relative improvement for all case scenarios for slice 2 of patient B. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

- **Integration-mode**

Selected reconstruction steps of the SART algorithm are shown in Fig. 3.17. Already at iteration 50 (Fig. 3.17a), it is clear that simpler models (i.e. nominal line and scattering cone) fail in capturing complex anatomical details, which are only hinted at without the possibility for a specific distinction. Particularly, anatomical borders are not properly defined and air cavities can appear as smoothed out. The scattering cones images are instead able to provide for a better depiction of the considered cross-section: despite the obvious limitations due to blurring effects, the anatomical structures are more clearly identifiable and their borders with air cavities or the more general border between the whole anatomy and the air surrounding it is much more sharper.

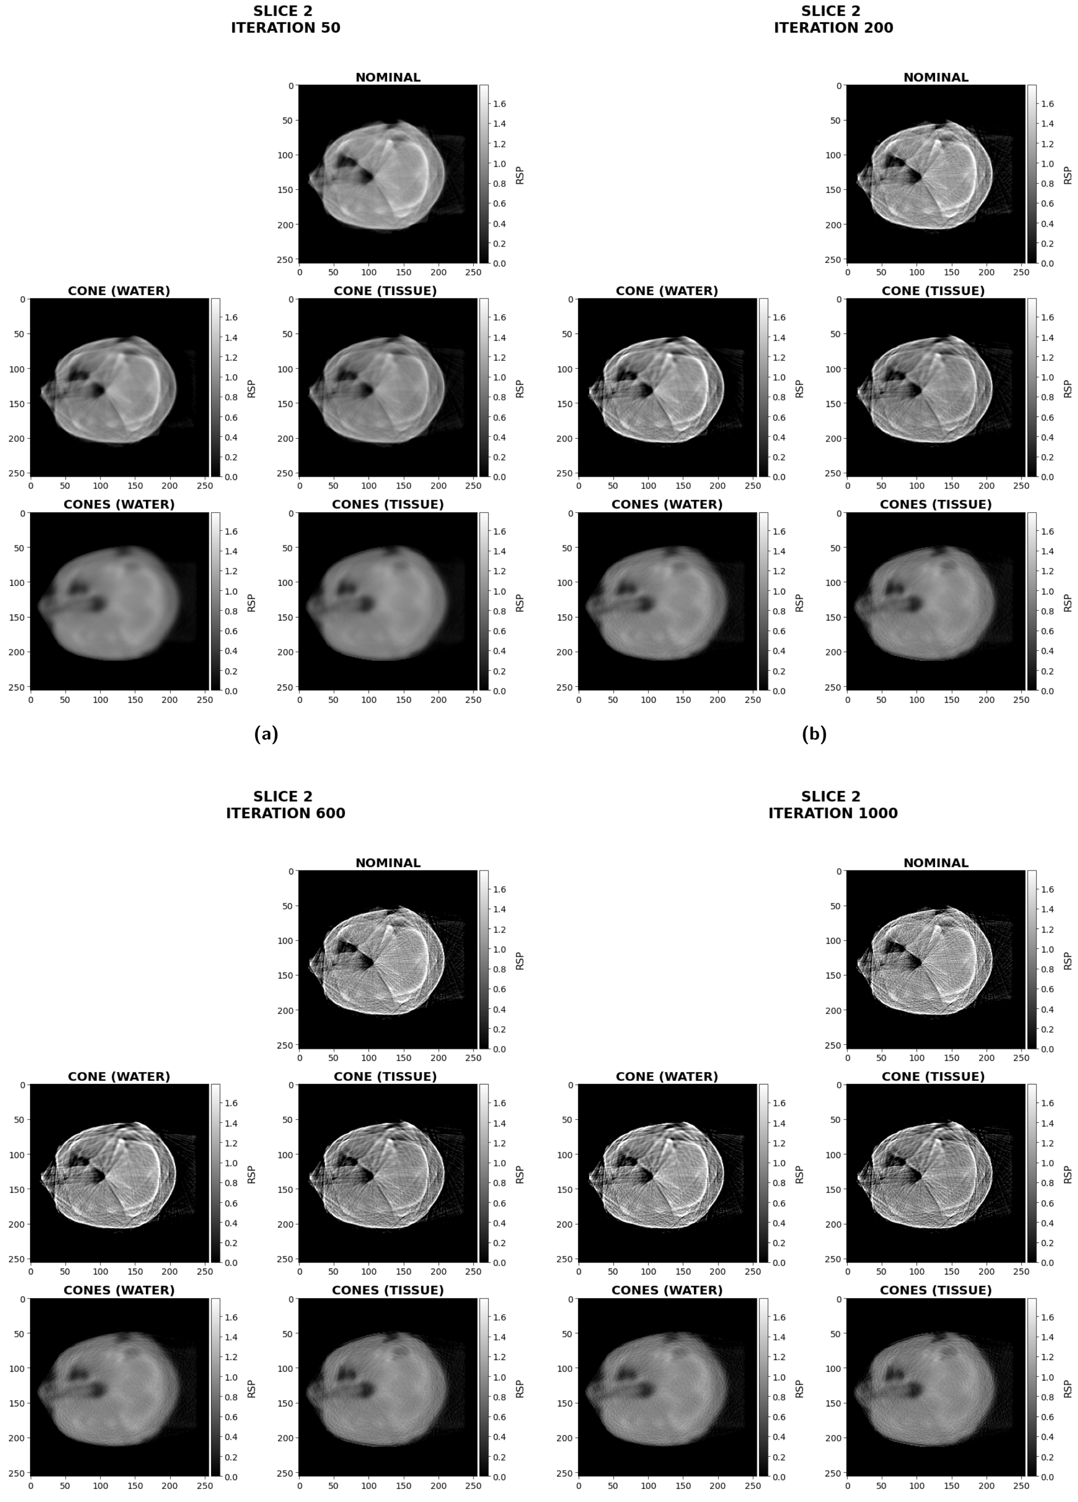

(a) (b)

(c) (d)

Figure 3.17: Integration Mode - Progression of the reconstruction algorithm for slice 2 of patient B. The nominal line and scattering cone images are affected by poor spatial resolution and image quality, while the scattering cones models are able to reproduce more accurately the depicted anatomy, though smoothing phenomena may be observed throughout the entire tomographic reconstruction. Noise break-up heavily affects the reconstructed images, consistently limiting the informative content of the WET$_{mode}$ models more than the scattering cones images.

At iteration 200 (Fig. 3.17b), the WET$_{mode}$ models begin to be consistently affected by noise artifacts in the form of streak lines, while the scattering cones suffer from grainy artifacts which, however, do not compromise entirely image quality since anatomical geometries can still be detected. By the last iterations (Fig. 3.17c and Fig. 3.17d), noise break-up is heavily compromising image quality for the nominal line and the scattering cone images, making it difficult to tell streak line artifacts and anatomical borders apart. Still preserving discrete image quality and contrast, graininess in the scattering cones images is more evident, with edge artifacts appearing especially for the tissue-based scenario.
The MAE quantification of the integration-mode models against the Ground Truth is reported in Fig. 3.18.

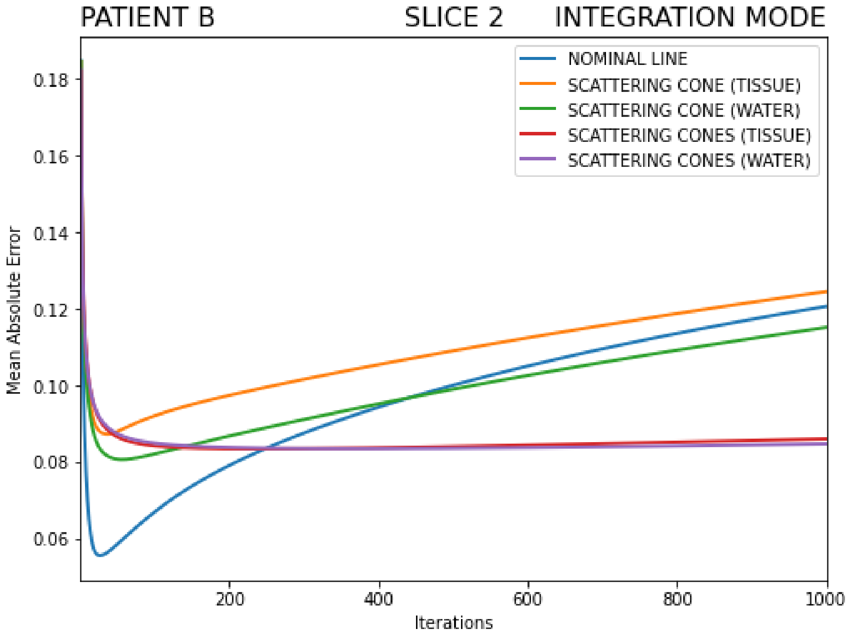

Figure 3.18: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient B. Though a trend reversal can be observed between the nominal line and the scattering cone in water mid-way through the iterative process, the scattering cone in tissues fails in provide a better solution against these two models. The scattering cones offer an overall better image quality and convergence, although they have comparable performances for the water- and tissue-based reconstruction.

When compared to the other more complex models, the nominal line shows the lowest error, but also poor stability against noise. Particularly, the scattering cone in water is able to outperform the nominal line at iteration 442 with a relative improvement of 0.02% up until a 4.52% difference by the last iteration. The scattering cone in tissues, however, fails against both the nominal line and the water cone, with a 3.21% and 8.09% difference respectively at iteration 1000. The scattering cones show a slower but general better convergence. Particularly, the scattering cones in water present a 29.73% and a 26.40% maximum improvement over the nominal line and the water cone respectively by the last iteration, while the scattering cones in tissues outperform the nominal line by a 28.65% error margin and the tissue cone with a maximum improvement of 30.86%, both at iteration 1000. Finally, by starting with a slight predominance during early iterations (0.96% relative improvement at iteration 58), the tissue cones are outperformed by the water cones as the tomographic reconstruction progresses, leading to a 1.54% difference by the last iteration.
A more compact manner of reporting these results can be found in table 3.8.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-8.99%	0.95%	-37.77%	-50.16%	-8.25%	1.63%	-49.14%	-47.72%
200	-12.19%	0.39%	-9.72%	-23.10%	3.44%	14.27%	-5.94%	-5.53%
600	-9.60%	-0.77%	2.34%	-7.03%	18.36%	24.93%	20.27%	19.65%
1000	-8.09%	-1.54%	4.52%	-3.21%	26.40%	30.86%	29.73%	28.65%

Table 3.8: Integration Mode - Relative improvement for all case scenarios for slice 2 of patient B. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

3.3 Patient C

The cross-section images considered for Patient C are shown in Fig. 3.19.

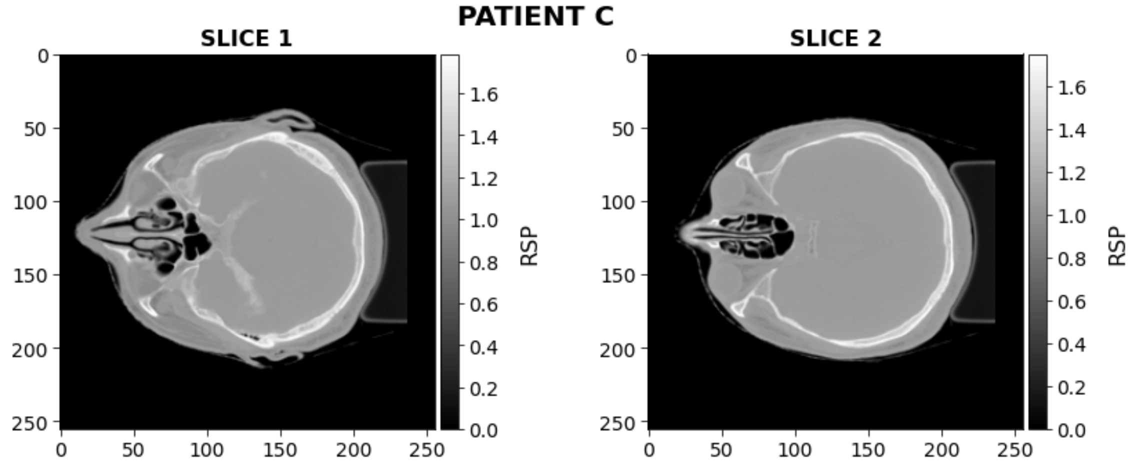

Figure 3.19: Ground Truth C. For patient C, the Ground Truth slices depicted in this figure were reconstructed.

3.3.1 Slice 1

- **List-mode**

The reconstructed images for slice 1 of patient C at sampled stages of the OS-SART algorithm are shown in Fig. 3.20. Starting off with a slightly blurred output due to the convergence of high frequencies at early iterations (Fig. 3.20a), the algorithm quickly approaches an optimal solution. In fact, though the general structure of the head is already clearly visible, the edges of the anatomical features, like bone and tissue boundaries, are more sharply defined by iteration 10 (Fig. 3.20b) in all four proposed models; moreover, the complex geometry of the air cavities is extremely well-preserved and reproduced. At iteration 20 (Fig. 3.20c), the images still offer a good contrast and spatial resolution, however noise artifacts begin to be noticeable, primarily as streak lines. Of course, this is even more evident at the last iteration (Fig. 3.20d), where the images appear more pixelated, though without heavily compromising their accuracy in reproducing the Ground Truth anatomy.

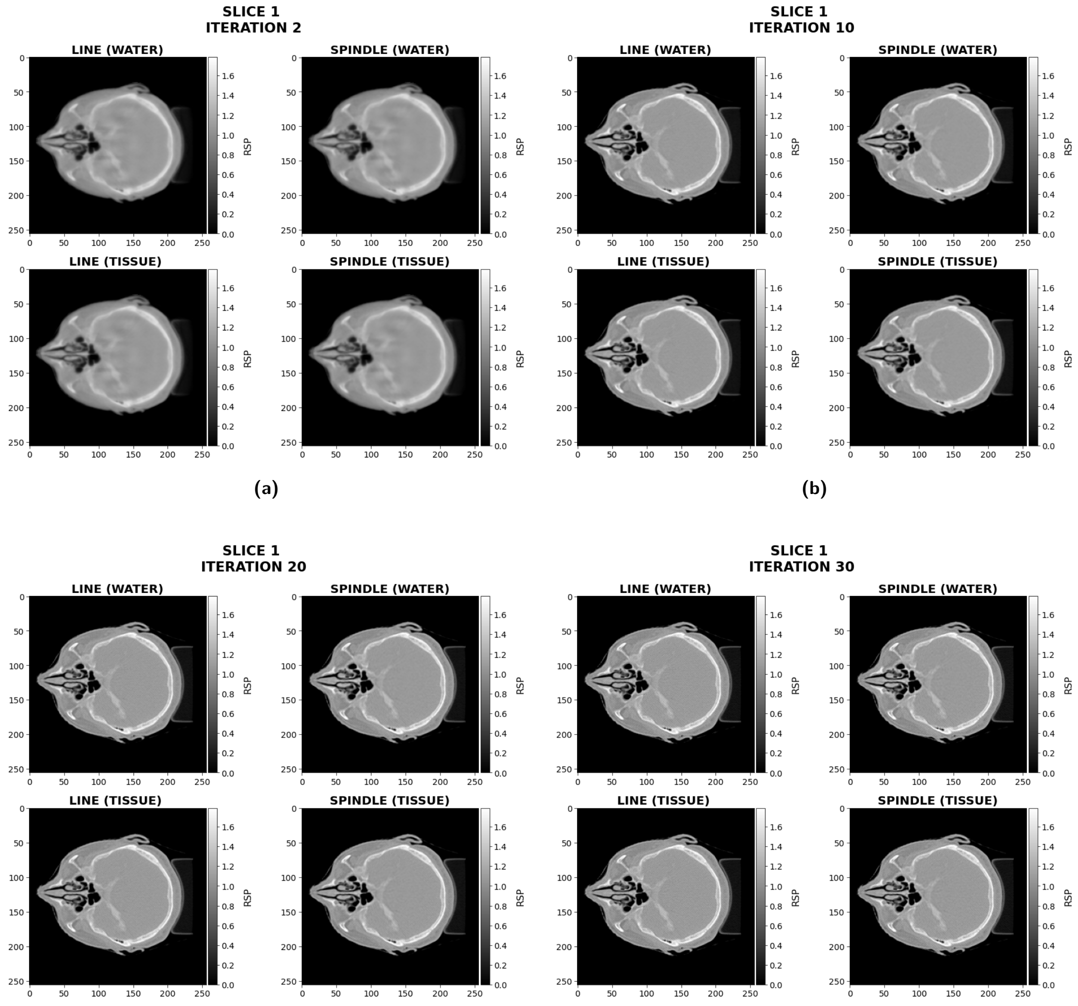

(a) (b)

(c) (d)

Figure 3.20: List Mode - Progression of the reconstruction algorithm for slice 1 of patient C. Although already clear at early iterations, the reproduced anatomy appears well-defined and accurately detailed as the tomographic reconstruction progresses for all the models. Though noise break-up phenomena start to affect the images after the convergence, they are still able to retain an ideal contrast and spatial resolution.

The MAE quantification provides for a better insight on the behavior of the list-mode models for the examined slice (Fig. 3.21). The scattering line in tissues is able to outperform the respective water model with a 0.04% relative improvement already at iteration 1, culminating in a maximum advantage of 1.07% by iteration 11. The scattering spindle in water presents a slower convergence than the water line, as highlighted by a 1.86% difference at iteration 5 in favor of the latter model, but also a greater stability to noise with a relative improvement of 6.48% by the last iteration. The scattering spindle in tissues provides a lower error than all the other models. Particularly, by taking into account a 0.30% difference in favor of the water line at iteration 3, the tissue spindle then reaches a maximum improvement of 7.59% at iteration 30. Furthermore, though a 0.74% difference at iteration 3 marks a slower convergence in favor of the tissue line, the scattering spindle in tissues sees an evident predominance with a maximum relative improvement of 7.11% by the last iteration. Finally, an advantage of 0.27% over the water spindle can already be noticed at iteration 1, leading to a peak improvement of 2.68% at iteration 11.

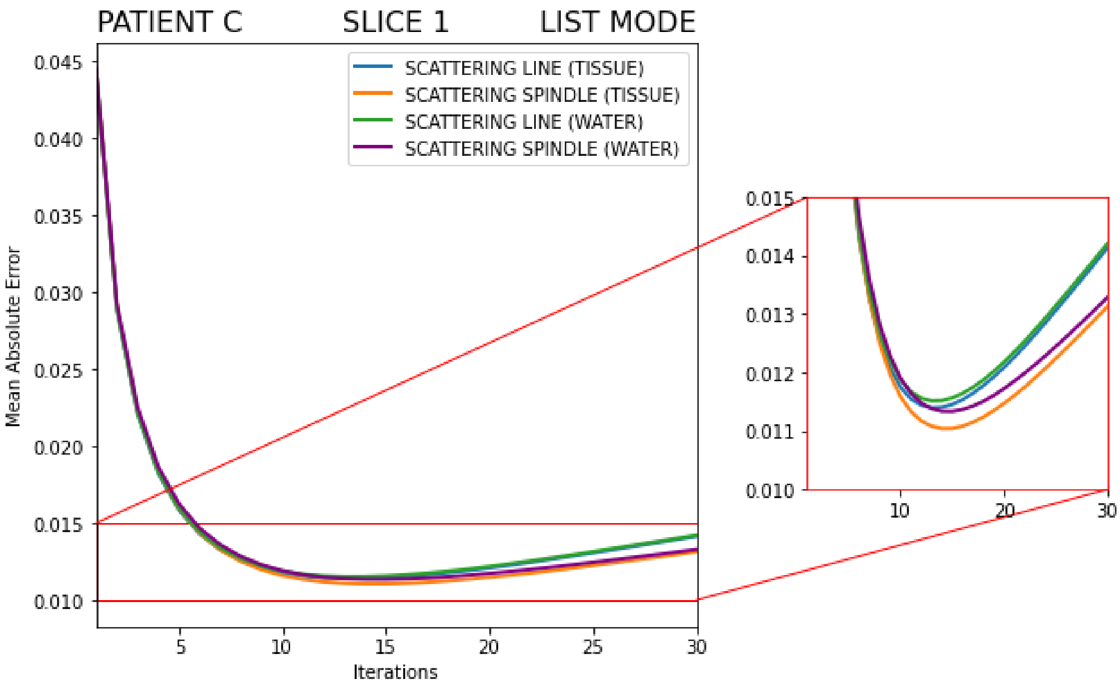

Figure 3.21: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient C. Although all four test models are able to provide optimal performances when compared to the Ground Truth, the tissue-based scenario proves to be more accurate and robust to noise for both the scattering line and the scattering spindle. Particularly, the tissue spindle is the model that is more faithful to the Ground Truth, showing the lowest error even after the convergence point.

All these considerations are summarized in table 3.9.

Iterations	C1	C2	C3	C4	C5
2	0.24%	0.78%	-1.09%	-0.54%	-0.29%
10	1.06%	2.66%	-0.23%	1.38%	2.43%
20	0.76%	2.16%	3.73%	5.09%	5.81%
30	0.52%	1.18%	6.48%	7.11%	7.59%

Table 3.9: List Mode - Relative improvement for all case scenarios for slice 1 of patient C. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

- **Integration-mode**

Fig. 3.22 shows the output of the SART algorithm at different points during the tomographic reconstruction for slice 1 of patient C. Already at iteration 50 (Fig. 3.22a), it is possible to observe how the nominal line and the scattering cone (both in water and tissues) suffer from poor spatial resolution: though the head geometry is recognizable, anatomical details and delimitations between structures are lost, making it challenging to identify different tissue compositions. As the tomographic reconstruction progresses, noise break-up starts to heavily affect the images (Fig. 3.22b): specifically, streak lines are predominant but also edge artifacts, that were only slightly visible at early stages, are now much more evident. This becomes particularly problematic during later iterations, as highlighted by Fig. 3.22d, where anatomical boundaries and streak lines are difficult to tell apart.

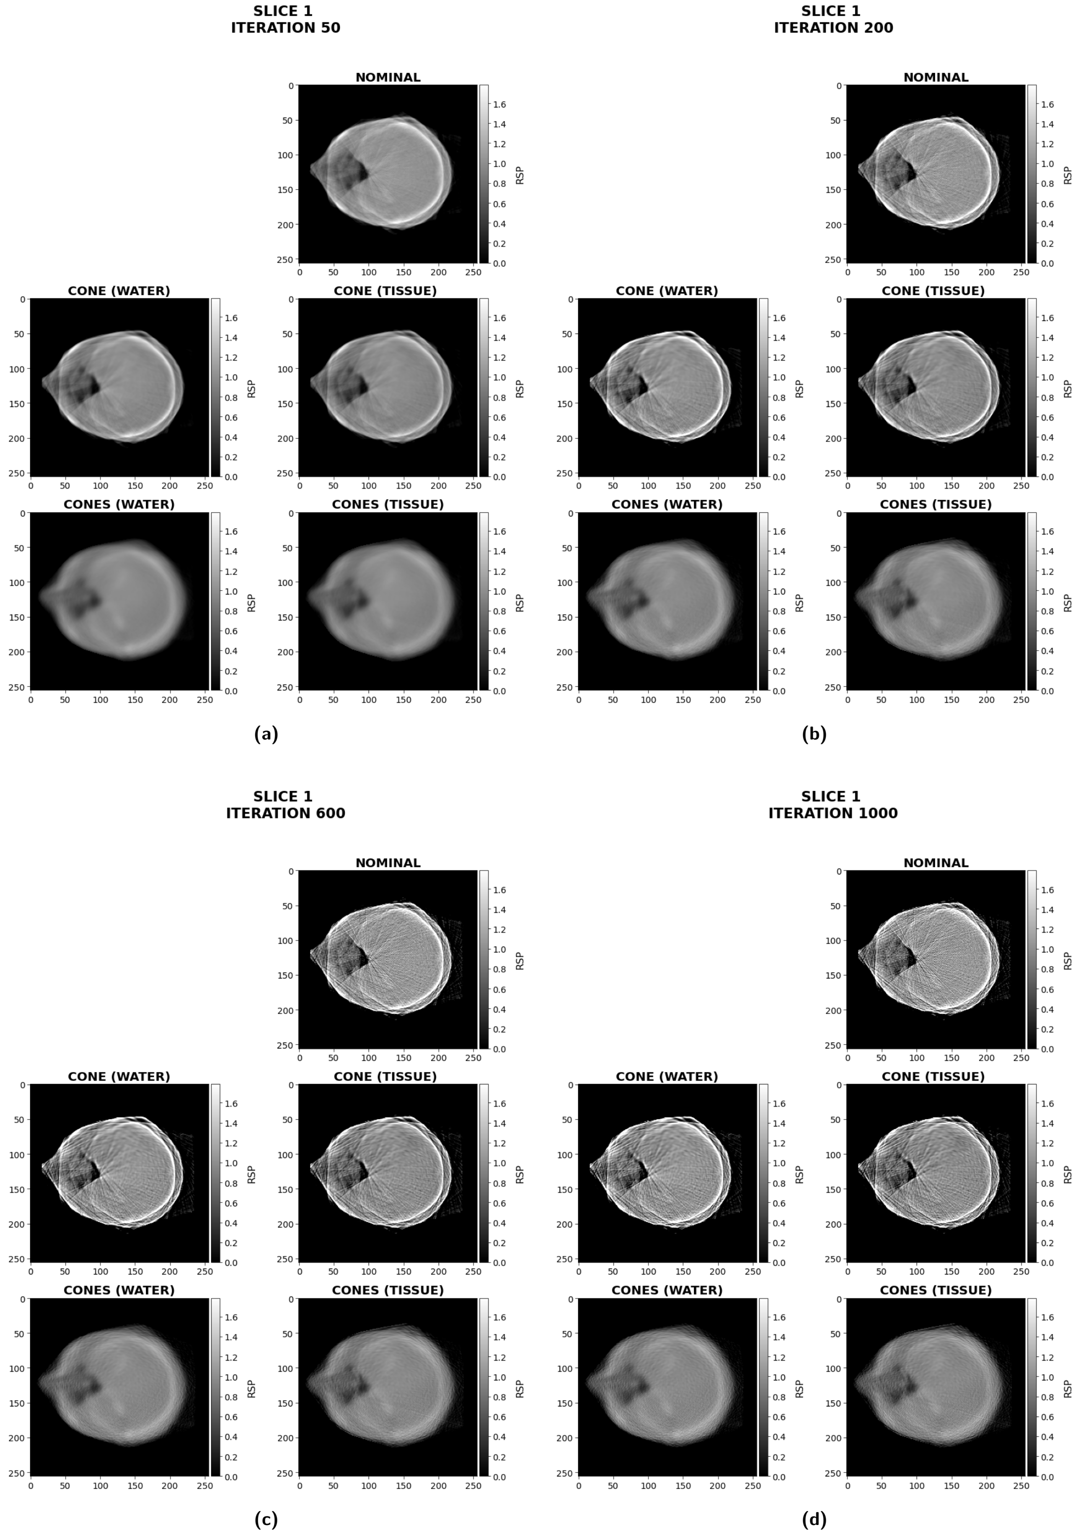

Figure 3.22: Integration Mode - Progression of the reconstruction algorithm for slice 1 of patient C. Due to poor spatial resolution and instability to noise, the WET$_{mode}$ models are limited in reproducing the Ground Truth slice, allowing to retrieve an approximated geometry of the patient anatomy. The scattering cones images are instead more accurate and better suited in preserving the differences in tissue compositions and anatomical structures, although an overall blurring effect can be noticed during the entire tomographic reconstruction.

On the other hand, the reconstructed anatomy appears more preserved when it comes to the scattering cones: both the water- and tissue-based scenario are able to provide much greater accuracy in reproducing the anatomical details of the Ground Truth. In fact, at iteration 50 (Fig. 3.22a), it is possible to make out much more clearly the peculiar shape of the air cavities in the nasal area of the patient than it is with the WET$_{mode}$ models; moreover, the differentiation between soft tissues and bones is definitely sharper and more pronounced. The scattering cones images also seem to be more robust to noise as the iterations progress: in fact, though graininess is easily detectable by iteration 200 (Fig. 3.22b) and then at iteration 600 (Fig. 3.22c), the specific anatomical traits and their borders can still be identified. By the last iteration (Fig. 3.22d), also edge artifacts are more consistent, but noise break-up is not as compromising as it is for the WET$_{mode}$ models. Despite these advantages, image quality is however restrained by a general out-of-focus and blurring effect that can be observed throughout the entire iterative process.
The quantitative comparison between the reconstructed images and the Ground Truth cross-section is reported in Fig. 3.23.

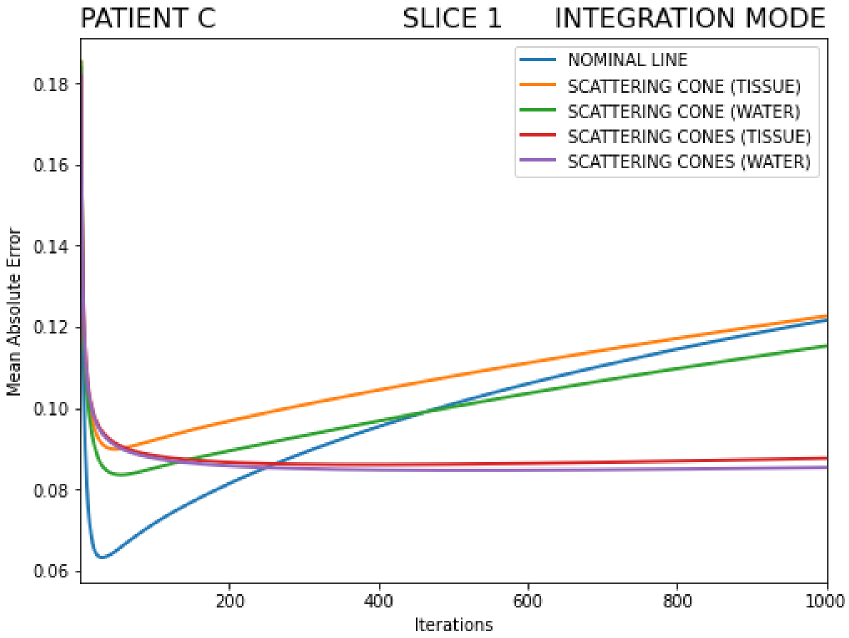

Figure 3.23: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 1 of patient C. The nominal line, the scattering cone in water and the scattering cone in tissues are strongly affected by noise amplification. The scattering cone in water shows however greater robustness to noise when compared to the other two WET$_{mode}$ models, while the tissue cone is not able to outperform either of them. The scattering cones show an overall better image convergence over the other models, but the water-based scenario is still predominant over the tissue one.

Even if the nominal line shows the fastest convergence, it behaves poorly against noise. In fact, it is outperformed by the scattering cone in water already at iteration 463 with a relative improvement of 0.02%, ending up with a maximum difference of 5.23% by the last iteration. However, the scattering cone in tissues presents worse performances when compared to this models, with a difference of 0.85% and 6.42% at iteration 1000 in favor of the nominal line and the water cone respectively. The performances of the scattering cones are instead noticeably different. They show a 50.36% and 51.06% difference with the nominal line by iteration 22 for the water- and tissue-based scenario respectively, but a trend reversal leads to a relative improvement of 29.82% in water and 27.95% in tissues by the last iteration. Moreover, the multiple cones are able to outperform the single cone in water with a 25.94% relative improvement at iteration 1000, while in tissues this same difference amounts to 28.56%. However, the water cones are predominant over the tissue ones, with a maximum advantage of 2.66% by the last iteration.
In table 3.10, the different performances of these models are reported.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-7.54%	-0.63%	-28.44%	-38.13%	-8.25%	-1.29%	-39.04%	-39.91%
200	-8.23%	-1.03%	-9.90%	-18.95%	4.16%	10.53%	-5.33%	-6.42%
600	-7.27%	-1.99%	2.25%	-4.85%	18.26%	22.28%	20.10%	18.51%
1000	-6.42%	-2.66%	5.23%	-0.85%	25.94%	28.56%	29.82%	27.95%

Table 3.10: Integration Mode - Relative improvement for all case scenarios for slice 1 of patient C. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

3.3.2 Slice 2

- **List-mode**

Fig. 3.24 shows sampled stages of the tomographic reconstruction process for slice 2 of patient C. Early iterations see a first convergence of high frequencies causing a slight blurring effect throughout the reconstructed image (Fig. 3.24a), although the anatomical features of the patient's head are already identifiable. By iteration 10 (Fig. 3.24b), the accuracy in reproducing the Ground Truth can be noticed: boundaries between different tissue compositions are clear and even small details (such as the ocular globes and the rich geometry of the air cavities) are well-preserved. However, already at iteration 20 (Fig. 3.24c), noise amplification disturbances can be detected as streak lines begin to appear throughout the considered slice. These artifacts are even more noticeable at the last iteration (Fig. 3.24d), where the images start to be slightly pixaleted but without heavily compromising the overall image quality.

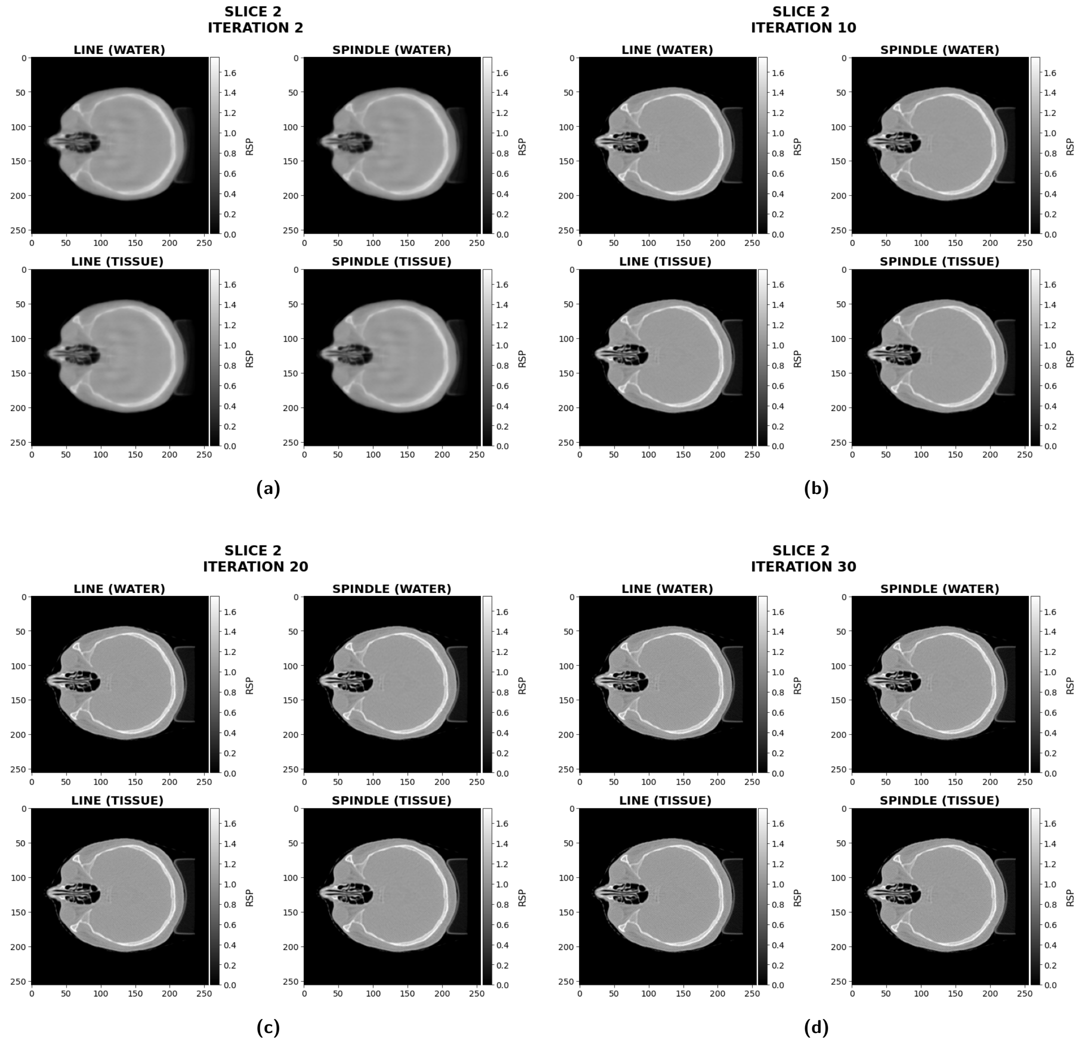

(a) (b)

(c) (d)

Figure 3.24: List Mode - Progression of the reconstruction algorithm for slice 2 of patient C. As the number of iterations progresses, image quality improves in all proposed models in terms of spatial resolution and reproduction of anatomical structures. By the last iterations, noise break-up effects are visible, but the patient's anatomy remains generally well-preserved.

The MAE quantification in Fig. 3.25 provides for a more precise depiction of the performances of these models in reference to the respective Ground Truth slice. The scattering line in tissues is able to outperform the water line already at iteration 1 with a 0.11% relative improvement, which culminates in a 0.91% advantage at iteration 11. The scattering spindle in water sees a difference of 1.12% in favor of the water line at iteration 4 but, by iteration 30, a trend reversal leads to a 8.75% relative improvement. The scattering spindle in tissues represents the model with the best performances. Considering a disadvantage of 0.16% in favor of the scattering line in water at iteration 2, the tissue spindle is then able to reach a maximum relative improvement of 8.86% by the last iteration. Furthermore, though a maximum difference of 0.30% in favor of the scattering line in tissues can be observed at iteration 3, the scattering spindle in tissues is still able to outperform the tissue line reaching a peak advantage of 8.60% at iteration 30. Finally, the tissue spindle is also able to outperform the water spindle already at the first iteration (0.19% relative improvement), peaking at iteration 11 with a 2.25% advantage.

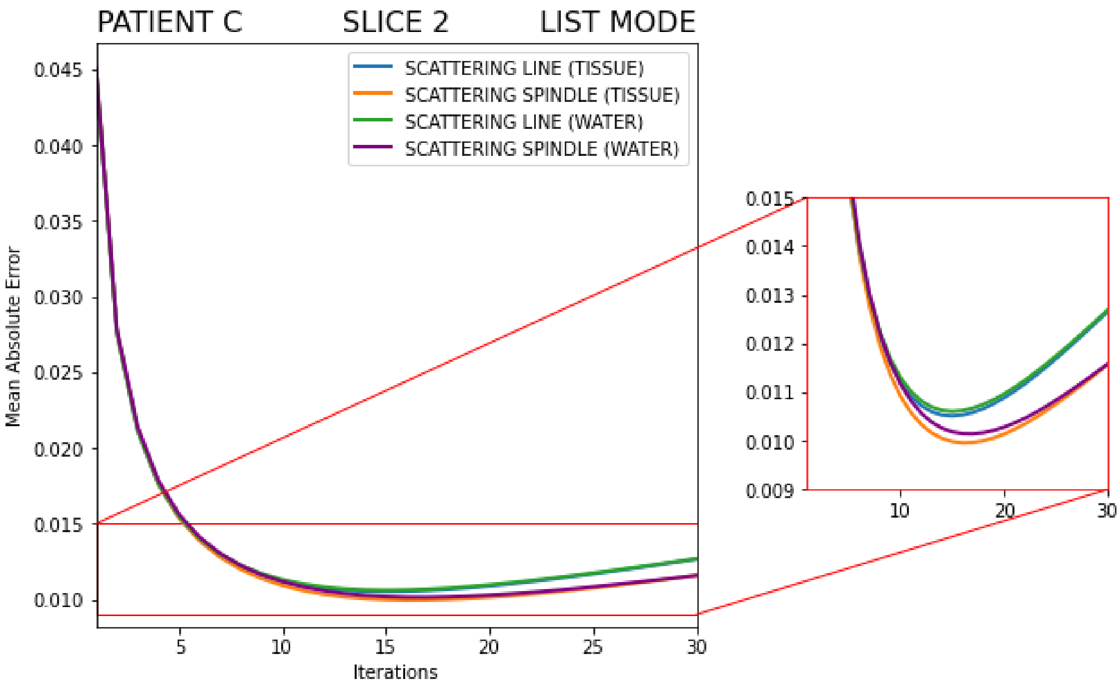

Figure 3.25: List Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient C. The scattering spindle models they are able to reach the lowest error and are also more stable to noise when compared to the respective scattering line models. Moreover, the tissue-based scenario provides for a better image convergence against the water-based one: although marginal for the scattering line, this statement is more significant and evident for the scattering spindle.

The behaviour of the list-mode models for slice 2 of patient C are schematically reported in table 3.11.

Iterations	C1	C2	C3	C4	C5
2	0.11%	0.60%	-0.76%	-0.27%	-0.16%
10	0.89%	2.24%	1.18%	2.53%	3.40%
20	0.72%	1.38%	6.21%	6.84%	7.51%
30	0.28%	0.12%	8.75%	8.60%	8.86%

Table 3.11: List Mode - Relative improvement for all case scenarios for slice 2 of patient C. The percentage improvement of tissues over water for the scattering line model is labeled as C1, while the one for the scattering spindle model is labeled as C2. The percentage improvement of the scattering spindle model over the scattering line model in water is labeled as C3, while the one in tissues is labeled as C4. Column C5 lists the relative improvement of the scattering spindle in tissues over the scattering line in water.

- **Integration-mode**

The images for slice 2 of patient C at different iterations of the reconstruction process are shown in Fig. 3.26. By iteration 50 (Fig. 3.26a), it is evident that the WET$_{mode}$ models have difficulties in reproducing the anatomical structures of the Ground Truth: thought the general shape of the patient slice is recognizable, the head contours are undefined and the boundaries between different regions can be unclear; moreover, the complex geometry of the air cavities in the nasal area is lost. Aside for a constant blurring effect during the tomographic reconstruction, the scattering cones provide for an image quality that is more refined, with anatomical details from the Ground Truth more accurately reproduced.

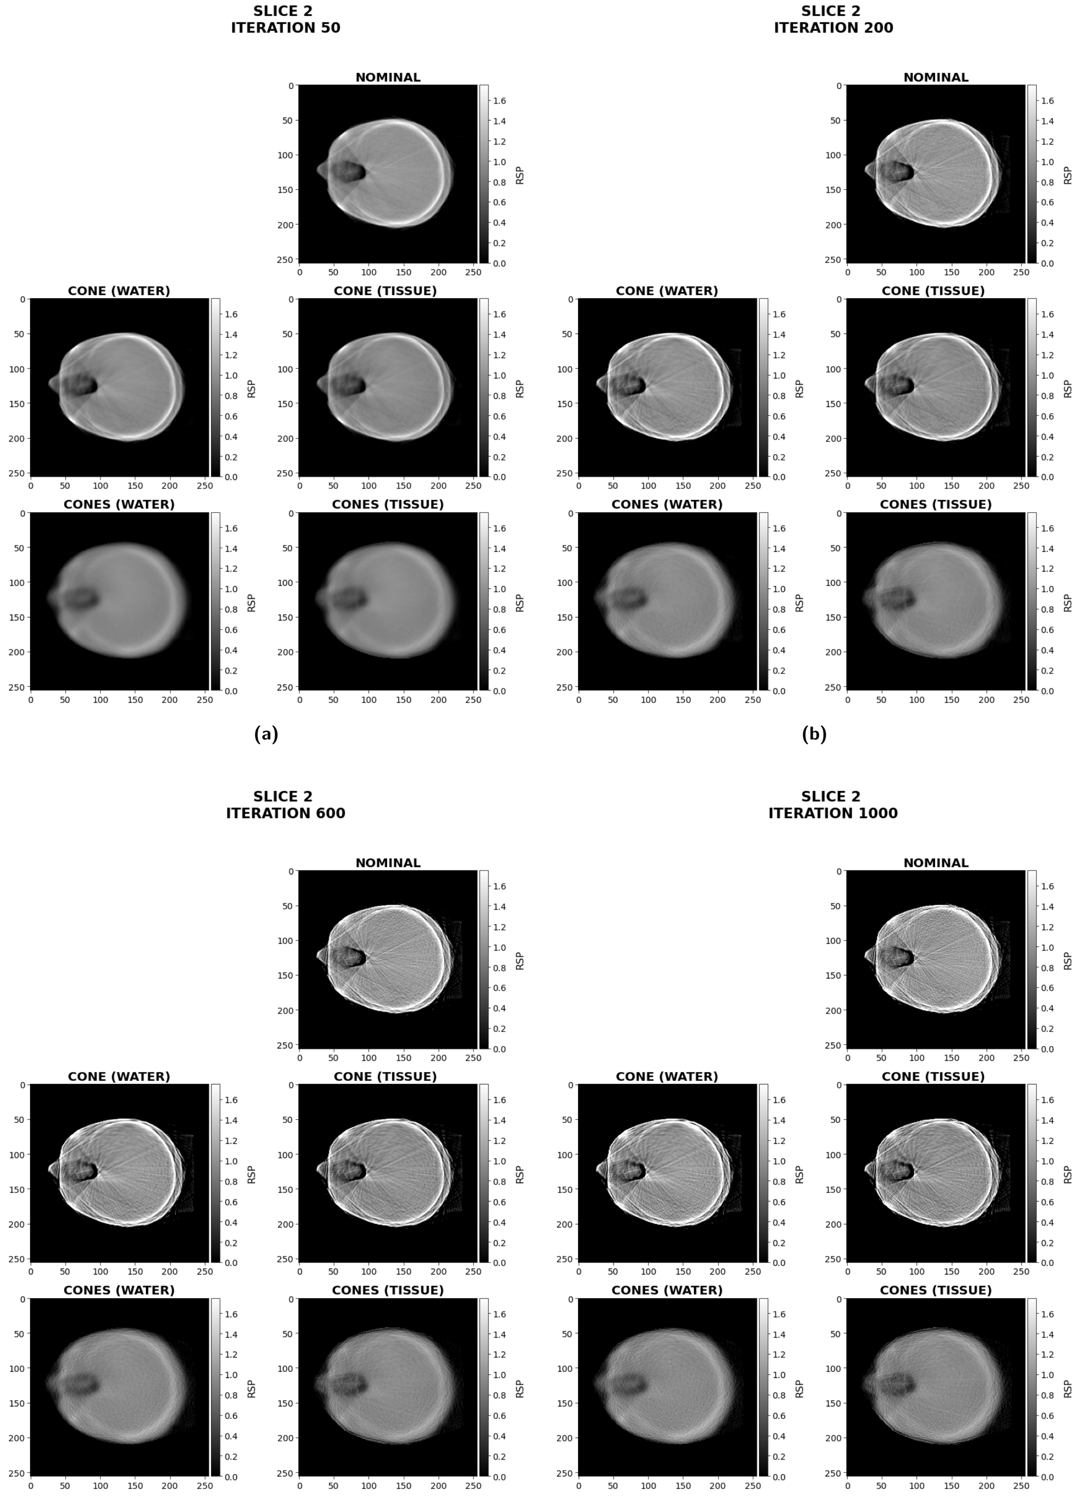

(a) (b) (c) (d)

Figure 3.26: Integration Mode - Progression of the reconstruction algorithm for slice 2 of patient C. Though it is possible to account for heavy differences in tissue composition and the general shape of the considered anatomy, the nominal line and the scattering cone are visibly compromised by poor spatial resolution, low contrast and noise amplification throughout the iterative reconstruction. The scattering cones appear more suited in reproducing the anatomical features of the Ground Truth slice, but a constant out-of-focus effect is still a limiting factor in their accuracy.

Noise amplification is already affecting the reconstructed images by iteration 200 (Fig. 3.26b): specifically, streak line artifacts start to compromise image quality for the nominal line and the scattering cone, while the scattering cones tend to have a more grainy appearance; however, by iteration 600 (Fig. 3.26c), streak lines can also be observed particularly for the scattering cones in tissues. At the last iteration (Fig. 3.26d), noise break-up effects can easily be detected in all images: the WET_{mode} models are once again the most affected, where distinguishing between different anatomical structures and noise artifacts proves challenging; additionally, edge artifacts are observed for the scattering cones.
Comparing the reconstructed images to the Ground Truth using MAE quantification (Fig. 3.27) allows for a more precise evaluation of the performances of the models.

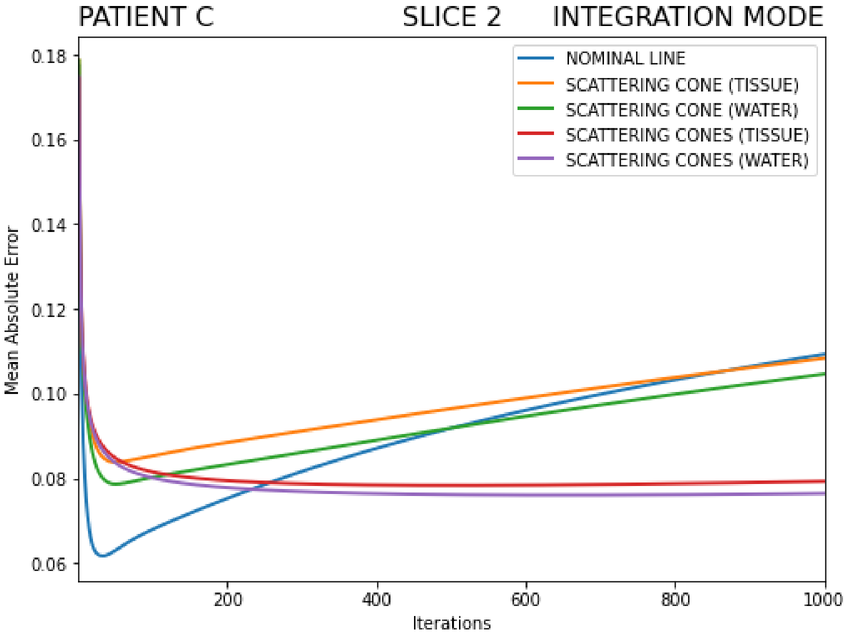

Figure 3.27: Integration Mode - Mean Absolute Error between Ground Truth and reconstructed images for slice 2 of patient C. The nominal line shows a fast convergence but also great instability to noise when compared to the other WET_{mode} models. While the water cone seems to behave more robustly against noise amplification than the nominal line, the scattering cone in tissues is only able to slightly outperform the nominal line towards the last iterations, remaining however at a disadvantage against the scattering cone in water. The scattering cones are instead able to provide for an overall better image convergence than the WET_{mode} models, although the water-based scenario is predominant over the tissue one.

The nominal line seems to reach the lowest error (0.06 at iteration 33), however it is also strongly affected by noise break-up. In fact, the scattering cone in water is able to outperform the nominal line at iteration 503 with a 0.01% advantage, leading to a relative improvement of 4.29% by the last iteration. However, when compared to the other WET_{mode} models, the scattering cone in tissues is able to outperform the nominal line only towards the last iterations (maximum relative improvement of 0.85% at iteration 1000), while if fails against the water cone with a difference of 3.59% at iteration 1000 in favor of the latter. The scattering cones behave differently. In particular, both the water- and the tissue-based scenario are able to outperform the nominal line, with a maximum advantage of 30.05% and 27.47% by the last iteration respectively. Considering specifically the water models, the multiple cones see a maximum disadvantage of 8.27% at iteration 31 against the single cone, but also a relative improvement of 26.92% by the last iteration. Instead, the scattering cones in tissues are able to turn a 3.27% difference in favor of the tissue cone at iteration 27 into a 26.85% relative improvement by iteration 1000. However, the scattering cones in water are predominant over the tissue cones, as assessed by a 3.69% maximum difference at iteration 1000.

Table 3.12 illustrates these trends in a more compact manner at different reconstruction steps.

Iterations	C1	C2	C3	C4	C5	C6	C7	C8
50	-6.53%	-1.38%	-24.86%	-33.01%	-6.35%	-1.20%	-32.78%	-34.61%
200	-6.27%	-2.11%	-10.59%	-17.52%	6.52%	10.18%	-3.38%	-5.55%
600	-4.55%	-3.07%	-1.45%	-3.04%	19.66%	20.80%	20.82%	18.50%
1000	-3.59%	-3.69%	4.29%	0.85%	26.92%	26.85%	30.05%	27.47%

Table 3.12: Integration Mode - Relative improvement for all case scenarios for slice 2 of patient C. The percentage improvement of tissues over water for the scattering cone model is labeled as C1, while the one for the scattering cones model is labeled as C2. The percentage improvement of the scattering cone model over the nominal line model in water is labeled as C3, while the one in tissues is labeled as C4. Analogously, column C5 holds the relative improvement of the scattering cones model over the scattering cone in water, while the respective case in tissue is labeled as C6. Columns C7 and C8 represent the percentage improvement of the scattering cones in water and in tissues over the nominal line respectively.

Chapter 4

Discussion

4.1 Image quality and RSP accuracy

Owing to the statistical limitations described in section 1.7, ion image reconstruction is certainly a delicate procedure. This appears clear when comparing the quality of list-mode and integration-mode images. The approximation of the proton path to a straight line leads to lower spatial resolution and amplified artifacts throughout the iterations, even when introducing the Gaussian scattering cone. Though integration-mode images are helpful in reproducing the general shape and evident anatomical discontinuities in the image, the loss of details remains as an indisputable limiting aspect (Fig. 3.4, Fig. 3.8, Fig. 3.13, Fig. 3.17, Fig. 3.22, Fig. 3.26). The detector configuration itself and how it affects tomographic reconstruction represent key factors in this outcome. The assignment of a single WET value (WET$_{mode}$) to a specific path model (i.e., nominal line or scattering cone) for the update of the image estimate effectively homogenizes the ion path across the object, ignoring the different stopping powers experienced by the ensemble of protons: accurate representations of tissue distribution and small features are thus more difficult to reproduce. Moreover, transition areas (such as bone and soft tissues or anatomy and air) are characterized by pronounced RSP variation: confining these events to one WET occurrence leads to less distinct edges and potentially incorrect segmentation of the anatomical structures. From a clinical point of view, the dominance of the most frequent WET value is also problematic since the reconstruction is biased towards more common tissues: the under-representation of less frequent materials may result in a significant loss of information for imaging purposes. A possible workaround could be the extraction of the mean value from the WET histogram, which would be more sensitive to all the quantities in the distribution. However, this also means that the WET$_{mean}$ used for reconstruction purposes would represent a trade-off among the stopping powers of different tissues: though this would include all tissue types in the anatomy of interest, outliers or extreme fluctuations in the WET distribution may shift the average value, potentially making the reconstructed image less sensitive to tissue distribution and variation. Nevertheless, due to the simplistic nature of the model, the WET$_{mode}$ approach highly suffers from noise amplification phenomena: artifacts such as streak lines, blurring and distortions in the reconstructed image are witnesses to this fact, especially with a high number of iterations (Fig. 3.4d, Fig. 3.8d, Fig. 3.13d, Fig. 3.17d, Fig. 3.22d, Fig. 3.26d). The dedicated reconstruction algorithm for the scattering cones deals with these aspects, accounting for the entire range relevant to the protons making up the pencil beam. In this scenario, the reconstructed images are more efficient in preserving spatial heterogeneity and tissue variation, where soft tissues, bones and air cavities are captured with more contrast and precise boundaries between structures. This also implies an improvement over the WET$_{mode}$ models, since all anatomical structures in the patient's anatomy are accounted for: in terms of treatment planning, this would allow to represent tissues with varying densities and compositions, including tumors or complex organs. Nevertheless, limitations and inaccuracies can still be pointed out. The dedicated algorithm has to compensate for a more processed dataset: the higher dimensionality of the system matrix may induce the SART algorithm to convergence issues and struggle in handling the reconstruction, resulting in a blurring effect that is more impactful when compared to the other integration-mode models. Furthermore, though the Gaussian scattering model is reproduced at each channel, the underlying assumption is still that of straight trajectory, limiting the model's ability to fully capture non-linear paths. Similarly to the single cone models, this approximation can contribute to noise disturbances, even if less destabilizing (Fig. 3.4d, Fig. 3.8d, Fig. 3.13d, Fig. 3.17d, Fig. 3.22d, Fig. 3.26d): it is in fact possible to notice that streak artifacts are not as defined as in the respective single cone images, both in water and tissues. Possible compensation techniques to validate the superiority of the scattering cones model could lie in post-processing steps, such as de-blurring algorithms to sharpen image resolution or edge-preserving filtering to enhance contrast without losing spatial accuracy. Though these techniques should be applied to all proposed models to ensure a fair comparison, it could be especially useful to see if the scattering cones would be able to compete against list-mode images.

Relying on positional information for the MLP computation, list-mode images allow instead for a much more consistent reconstruction, where anatomical structures and gradient transitions are extremely well-preserved. First of all, the benefit of using curved path estimations can be unequivocally assessed by a finer spatial resolution since the model relies on a statistical description of the individual proton events and its interaction with surrounding materials. For this reason, details are captured with great accuracy and even small features are reproduced quite faithfully. This achievement is even more remarkable considering the optimal RSP accuracy not only on relatively homogeneous anatomies (Fig. 3.6, Fig. 3.24) but also for more complex and heterogeneous slices (Fig. 3.2, Fig. 3.11, Fig. 3.15, Fig. 3.20). Not considering some smearing effects given by the convergence of high frequencies during the first iterations, the images appear hardly afflicted by blurring and smoothing attributes: in fact, the possibility of using precise WET quantities associated with single protons allows to avoid averaging effects deriving from the reconstruction based on WET occurrences. Such considerations reflect the interest for list-mode configuration in clinical procedures, since it is able to better differentiate among tissues with similar stopping powers, resulting in an enhanced contrast and sharper structural boundaries in the reconstructed image. Finally, even though noise break-up is certainly a contributing factor, artifacts such as streak lines are overall contained and do not compromise heavily image quality, although a discrete number of iterations was considered and the quantitative impact of a longer reconstruction remains to be investigated.

In general, proton-tracking systems offer a superior image quality due to a more accurate model estimation. On the other hand, integration-mode imaging based on the entire WET histogram has the potential to compete with tomographic reconstruction based on the MLP algorithm. Noise break-up sensitivity for all models proposed in this study could be compensated or somewhat attenuated by implementing superiorization techniques for the reconstruction process (see section 2.3.6). An exemplification of how a Total Variation Superiorization scheme (Penfold et al. [2010]) is able to restrain noise amplification is highlighted in Fig. 4.1 by a less steep curve, after the convergence point, for some models tested on a patient's slice.

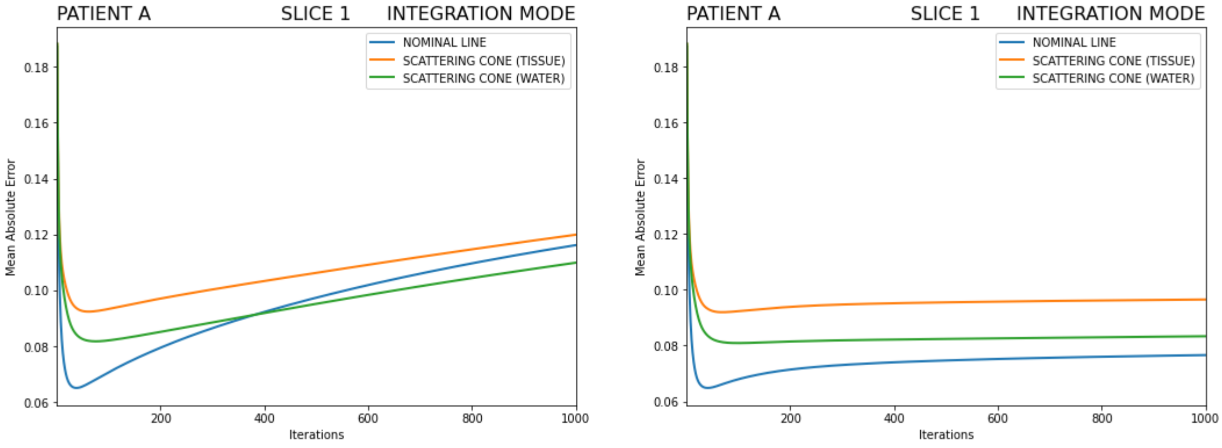

(a) (b)

Figure 4.1: Qualitative example of superiorization technique. After the convergence point, the images reconstructed with Total Variation Superiorization (b) are less sensitive to noise when compared to standard reconstruction algorithms (a).

4.2 Path estimation accuracy

The presented results demonstrate how different path estimations can affect the reconstruction process. The MAE quantification shows an overall lower error when trajectories are estimated based on the MLP algorithm for list-mode configuration than the straight path assumption for integration-mode configuration. This implies that the MLP provides for a more accurate path reconstruction and that its simplifying conditions (such as the small-angle approximation or the Gaussian model for MCS) are actually valid for list-mode data and can be accounted for to a certain extent by 3σ cuts. Though the scattering line model on its own is able to deliver acceptable accuracy even on complex patient anatomies, the intrinsic statistical uncertainty associated with MCS is addressed by the scattering spindle model. Both in water and tissues, the introduction of the Gaussian probability envelop around the ion trajectory leads to a more precise solution when reproducing the cross-section image (Fig. 3.2, Fig. 3.6, Fig. 3.11, Fig. 3.15, Fig. 3.20, Fig. 3.24). However, this comes at the cost of a slower convergence, as can be seen from table 3.1, table 3.3, table 3.5, table 3.7, table 3.9 and table 3.11, where a relative improvement that is negative at early iterations signals a predominance of the line model over the spindle model. Since the reconstruction algorithm has to explore a broader range of potential trajectories and integrate more information compared to the scattering line model, the objective function to be optimized is thus more complex. By the same logic, the spindle model also seems to offer more robustness against noise. Since the scattering line only considers the most probable path, it can be oversensitive to perturbations, leading the reconstruction algorithm to diverge towards a solution that may fit the noise rather than the actual data. The very nature of the scattering spindle makes it so that is more suited to handle uncertainties: by weighting the MLP across a probabilistic cloud, the Gaussian envelop effectively smooths out noisy data, keeping the solution consistent and reducing fluctuations in the reconstruction process. This same effect also acts a solid barrier against outliers, since the Gaussian distribution down-weights their influence as they lie outside the central probability mass, whereas the line-based reconstruction may be strongly skewed by these anomalies.

Understanding the smoothing effect introduced by the Gaussian volumetric distribution is essential to explain the different behavior of the path estimation models in water and tissues. When the ensemble of protons moves across pure water, the line integral is given by a cumulative sum of unit RSP. On the other hand, accounting for the anatomy's actual stopping powers when assembling the system matrix translates to a trajectory that sees instead different RSP values as it moves along its nominal path. Depending on the nature of the traversed material, this can greatly bias the resulting WET quantity: ideally, soft tissues and air cavities would generate a lower WET than water, while bony structures would lead to higher WET values. Since the Gaussian scattering model relies on information provided by the RSP map, it is intuitive that higher WET values are a synonym of a larger probability envelop: thus, the water-based models are generally more exposed to the smoothing properties provided by the Gaussian volumetric distribution. All of this stands to reason when comparing the reconstructed anatomies. Slice 2 for patient A shows a rather uniform and homogeneous anatomy, so the ions will be moving across an overall lower stopping power in tissues than in water. The scattering spindle in water presents a lower error than the one in tissues, particularly towards the last iterations since the larger Gaussian envelop makes the noise break-up toll less consistent. Interestingly enough, the scattering line does not behave the same way, but Fig. 3.7 and table 3.3 actually identify an improvement, though marginal, of the tissue model over the water one. When the the trajectories are estimated based on prior knowledge from the X-ray CT files, the scattering model is matter-of-factly more precise, since it offers an actual depiction of how protons would interact upon traversing a material with a specific stopping power. The scattering line only accounts for the MLP, so the OS-SART algorithm is not biased by any probabilistic averaging by the system matrix: the tissue model, being more realistic, is *de facto* capable of better reproducing the Ground Truth image. Cross-sections such as slice 1 for patient A, slice 1 and slice 2 for patient B or slice 1 and slice 2 for patient C demonstrate what happens when considering more complex anatomies. Not only the superiority of a tissue-based reconstruction is proven once again by the scattering line, but the larger Gaussian distribution associated with higher stopping powers also induces better results when comparing the tissue spindle to the water one (Fig. 3.3, Fig. 3.12, Fig. 3.16, Fig. 3.21, Fig. 3.25).

Similar conclusions can be drawn for integration-mode reconstruction. As can be seen from Fig. 3.5, Fig. 3.9 and Fig. 3.14, Fig. 3.18, Fig. 3.23 and Fig. 3.27, the nominal line benefits from a faster convergence than the water and tissue cone, also reaching a lower error at early iterations. The nominal line assumes a deterministic path for the pencil beam, making the reconstruction problem essentially more straightforward: though it may not represent the optimal solution, the SART algorithm is able to match more easily the projection data with the image estimate, since the model does not account for scattering. However, obvious limitations are implied in this case, since the nominal line is more sensitive to noise effects and error accumulation over the course of the iterations. On the other hand, the scattering cone is influenced by the probability envelop. The intra-reconstruction smoothing effect is in fact responsible for a slower convergence but also for a higher robustness to noise: this is witnessed by the fact that the water cone is able to outperform the nominal line as the reconstruction process progresses, capturing more precisely the actual distribution of the ion path and minimizing reconstruction fluctuations. Going one step further, the behavior of the scattering cone in tissues may appear otherwise contradictory: despite the implementation of the Gaussian envelop, this model tends to surpass very slowly the nominal line (which shows better performances at early iterations, though this tendency is quickly reversed due to increasing noise break-up), but also shows a higher error than the water cone does. The previously stated considerations about the scattering model can be accredited for this scenario as well but it is worth noting that, unlike list-mode images, this behavior seems to regularly occur in each slice, disregarding the depicted anatomy: the intrinsically flawed estimation of the ion trajectory is here explicit in explaining why the difference between the water and the tissue model is much more evident in integration-mode than list-mode configuration. For proton-tracking systems, the MLP algorithm allows to retrieve a more realistic trajectory that is able to support the additional information given by the introduction of the tissue spindle. On the other hand, when the nominal line is overloaded by data coming from a scattering distribution implemented with prior knowledge on the patient's anatomy, the complexity of the model is such that the integration-mode reconstruction fails in a correct representation of the underlying physics: this results in a systematic underestimation of the WET line integral, leading to a smaller Gaussian cone whose smoothing contribution is less reliable than the one in water.

Finally, as already anticipated in section 4.1, the reconstruction of the images based on the scattering cones is more demanding, since each column of the system matrix is now representative of the entire WET spectrum, meaning that each iteration of the dedicated tomographic algorithm has to handle more data: this results in a slower but better image convergence. In fact, while the optimal solution is reached during early iterations for the other models, the scattering cones tend to have their minimum mid-way across the iterative process but, already after iteration 200, their relative error is lower than all the other models, both in water and tissues. This behavior is preserved even towards the last iterations, indicating also a greater stability to noise break-up. From these results, it is clear that relying on the complete range dataset leads to a more robust performance where the actual distribution of the ion paths deviating from the nominal trajectory is more accurately captured. The trade-off between convergence speed and reconstruction accuracy is therefore properly justified, as the scattering cones are easily able to outperform the respective cone models, providing for a superior physical representation and clinical relevance. Regardless, the intra-reconstruction filtering is still a playing character when it comes to water/tissue comparison. Even if the probability envelop is declined on different WET components, their total contribution is what actually makes up the system matrix adopted for the SART algorithm: the individual cones on the different channels are still affected by the imprecise physical perspective of a proton-integrating system, contributing to a less stable model in tissues rather than in water even in a reconstruction scenario that is theoretically less biased by the ion radiography. However, as previously stated, the anatomy at play is a key component to keep in mind. While it is true that integration-mode configuration cannot benefit from the same path estimation accuracy provided by list-mode configuration, it is also undeniable that the scattering cones represent a more precise model than the single scattering cone. On rather uniform slices (such as slice 2 of patient A) or slices whose anatomy is strongly characterized by large or randomly distributed air cavities (such as slice 1 of patient A, slice 2 of patient B or slice 1 and slice 2 of patient C), it is understandable that the overall WET associated to the scattering cones in tissues may be lower or at least comparable to the one in water, leading to errors lower in water (Fig. 3.5, Fig. 3.9, Fig. 3.23, Fig. 3.27) or somewhat similar performances (Fig. 3.18) respectively. Nonetheless, as seen for list-mode images, a higher complexity of the patient's anatomy may lead to an overall higher WET along the ion trajectory. Based on this, it is reasonable that the aforementioned scenario would be inverted for slice 1 of patient B (Fig. 3.14). In this case, although the water-based scenario appears to behave slightly better against noise towards the last iterations, the tissue cones are able to withstand the implementation of the probabilistic envelop during early iterations, showing a lower error. This seems to suggest a similar trend to the list-mode models, where a proper estimation of the ion path and its scattering behavior combined with a sufficiently structured anatomy may lead to a compensation of the larger intra-reconstruction smoothing effect carried out by the water models.

The reasoning so far is on its own comprehensive in explaining why the proposed models show a certain behavior when declined in different scenarios, but a parallel point of view may be found. In fact, a key factor may hide not only in the anatomical structure of the cross-sections examined in this work, but more generally in the specific anatomical district taken into consideration. The *head & neck* region can roughly be approximated to concentric circles, where the inhomogeneities (i.e. variations in tissue properties) tend to be distributed along the radial direction: in other words, when moving towards or away from the center of the skull, lateral differences in tissue distribution (such as between white matter, gray matter and bone) are more important than longitudinal ones. In light of this, it is important to recall that, for prior-knowledge reconstruction, the stopping power information necessary to implement the Gaussian scattering model is retrieved along the proton or pencil beam's nominal direction in the RSP map for list-mode (see section 2.3.4) and integration-mode configuration (see section 2.4.3) respectively. This implies that the accuracy of a tissue-based scenario is properly valid as long as the longitudinal inhomogeneities are dominant: if lateral inhomogeneities are instead more significant, the physical assumptions underlying the probabilistic envelop itself may become innacurate. In a water-based scenario, these issues are not relevant since a reconstruction with no prior knowledge homogenizes the whole anatomy at play to unit RSPs. Of course, once again these effects have to be combined with the accuracy of the specific path estimation model. For proton-tracking systems, the stopping power data extracted from the RSP map is then used locally to reconstruct the trajectory (and the Gaussian envelop) point by point, so the uncertainty in the scattering behavior of the individual proton is already minimized: cases in which the scattering tissue model performs worse the the water one are in fact rare and may be mostly attributed to the particular structure of the actual slice to be reconstructed, as discussed above. However, proton-integrating systems cannot benefit from a local-wise estimation of the ion trajectory: the predominance of lateral discontinuities on a straight line model, where the scattering power is derived from longitudinal data and implemented as a Gaussian cone along the nominal direction, may very well explain why the tissue cone fails against not only the water cone but, in some cases (Fig. 3.5, Fig. 3.18, Fig. 3.23) also against the nominal line. To a certain extent, the scattering cones may behave according to this very logic, although once again the higher precision of this model and the specific anatomy of the different slices are also to be taken into account for those cases where the tissue cones prevail over the water ones. Although reasonable, the influence of lateral and longitudinal inhomogeneities on path estimation has not been explored as the consideration of different anatomical regions characterized by different patterns of inhomogeneities (i.e. thorax and pelvis) was beyond the purposes of this study. In particular, the pelvic region shows a different geometry: unlike the head, its shape is compressed along the anterior-posterior axis and stretched along the lateral axis. This implies that the main inhomogeneities are distributed longitudinally, i.e. they are primarily observed front-to-back rather than side-to-side. In this context, the scattering model could be able to capture more accurately the tissue distribution and the proton scattering power embedded in the Gaussian envelop, possibly improving the scattering cone outcome.

4.3 Computational aspects

While not the main focus of this project, a brief overview on the computational cost implied by the proposed project workflow is due. Particularly, three main steps are to be considered: data production, system matrix assembling and iterative reconstruction.

Since the MLP algorithm allows to resolve the single proton trajectory, list-mode data are heavily processed. Specifically, the generation of the pencil beam matrices was parallelized on cluster resources by projection angle: the running time per projection angle amounted to roughly 12 hours. Next, the assembly of the 36 ordered subsets took about one day, as it did the reconstruction process itself; this last two steps were performed on local servers. These properties were shared by all four list-mode scenarios.

Because of a simpler model, the computational toll for integration-mode configuration is reduced. In particular, the runtime needed to create the system matrices for all pencil beams across all projection angles added up to 4 hours on a single local machine, in terms of nominal line, scattering cone in water and scattering cone in tissues separately. Then, the assembly of the single set and the image reconstruction together took around 30 minutes. The path estimation for the scattering cones had instead to be parallelized on

cluster resources following the same framework implemented for list-mode data. Here, the single projection angle runtime was over 2 days in tissues, while it took only 12 hours in water: this prominent discrepancy is mostly due to the multiple interpolation steps that are more time-consuming when implemented on the entire calibration curve in tissues rather than just one single value needed for the binary RSP map in water. Finally, since the contribution of each cone had to be weighted in a probabilistic manner, the combined runtime of the system matrix assembly and image reconstruction amounted to around 3 hours, for both the water- and tissue-based scenario. The last two steps were once again conducted locally.
Table 4.1 summarizes the computational aspects of this work.

ALGORITHM STEPS	RUNTIME								
	Line (Water)	Line (Tissue)	Spindle (Water)	Spindle (Tissue)	Nominal	Cone (Water)	Cone (Tissue)	Cones (Water)	Cones (Tissue)
System matrix generation	∼ 12 hours	∼ 12 hours	∼ 12 hours	∼ 12 hours	∼ 4 hours	∼ 4 hours	∼ 4 hours	∼ 12 hours	over 2 days
Subsets assembly	∼ 12 hours	∼ 12 hours	∼ 12 hours	∼ 12 hours	∼ 10 min	∼ 10 min	∼ 10 min	∼ 3 hours	∼ 3 hours
Image reconstruction	∼ 24 hours	∼ 24 hours	∼ 26 hours	∼ 24 hours	∼ 20 min	∼ 20 min	∼ 20 min	∼ 20 min	∼ 20 min

Table 4.1: Processing times for all image reconstruction steps in list-mode and integration-mode models. Though list-mode configuration offers an overall better performance and image quality compared to integration-mode, the computational cost can be quite antagonizing. The scattering cones may demand longer path estimation times, but the image reconstruction process is faster than the list-mode one.

From these findings, it appears obvious that, where list-mode configuration shines in terms of image quality and reconstruction accuracy, this comes at the cost of computational speed: within this frame, the analysis and production of large quantities of data could prove quite resource- and time-consuming. A compromise worth exploring could be reducing the number of protons per pencil beam: however, besides the loss of spatial resolution due to less information being captured, a lower proton statistics may help in a faster MLP computation, yet this would also imply a slower reconstruction since the algorithm would take longer to match the image estimate to the projection data and get to convergence. The scattering cones provide for an elegant compromise: the need for parallelized resources is actively compensated by a less demanding reconstruction. Especially relying on the water model, whose performances are able to compete with the tissue-based scenario (see section 4.1 and 4.2), the overall reconstruction process would roughly take half the time required for list-mode images.

Chapter 5

Conclusion

Ion beam therapy has the potential to become a landmark in cancer treatment. In fact, the physical properties and biological advantages of charged particles may represent a suitable solution when it comes to aggressive tumors or complex anatomical structures. Rather than relying on X-ray data for treatment planning, ion computed tomography and ion radiography can provide a significant development towards this goal, as they offer a more accurate description of ion physics. However, Multiple Coulomb Scattering events as the ions move across the patient can lead to poor spatial resolution and loss of information in the reconstructed image. In this work, the scattering behavior of protons across three different patients was assessed. The images were reconstructed according to two different detector configurations, while either assuming no prior knowledge on the scattering model (i.e. protons moving just across water) or with such prior knowledge (i.e. protons moving across actual tissues).

In list-mode configuration, the scattering line allows to retrieve good image quality, but the introduction of a Gaussian probabilistic envelop around the ion trajectory for the scattering spindle leads to an overall better convergence and robustness to noise. For proton-integrating systems, the Gaussian volumetric distribution for the scattering cone provided advantages over the straight nominal path of the pencil beam only for the water-based scenario, while the image reconstructed in tissues showed a higher error. A novel approach where all WET heterogeneities along the ion beam trajectory are accounted for was explored. The scattering cones led to slower but better convergence across the reconstruction process, although blurring effects are extremely evident in the images. The introduction of the Gaussian envelop induces an intra-reconstruction smoothing effect which favors image convergence. For simpler models (i.e., nominal line and single scattering cone), the Gaussian distribution destabilizes the underlying physics to be represented, leading to worse performances of the tissue models over the water ones. For more accurate models, such as the scattering cones or the list-mode ones, this effect is instead observable on simpler anatomies.

Since the reconstruction based on prior knowledge about the scattering behavior is theoretically more precise, future developments could be exploring the introduction of intra-reconstruction filtering mechanisms for these models in order to validate them over the water-based scenario. The investigation of a different anatomical region (i.e. pelvic district) could also be considered in order to assess the different effects of radial and longitudinal inhomogeneities on the Gaussian scattering model. Moreover, though proton-tracking systems stand right now the as a superior configuration, the reconstruction processes based on the scattering cones can represent a plausible alternative: future studies may investigate the possibility of introducing post-processing steps to improve image quality and resolution, making the reconstructed images comparable to list-mode performances. This may help in the clinical consolidation of proton imaging with a simpler and less expensive detector configuration.

Finally, the models proposed in this project may find interesting applications for machine learning purposes. The adoption of artificial intelligence for tomographic image reconstruction is rapidly spreading (Wang et al. [2020]): in particular, physics-informed networks are gaining relevance due to their hybrid nature between data-driven and model-based approaches. Within this project, the forward-projection model embedded in model-based tomographic image reconstruction was taken to the maximal accuracy for both detector configurations: this work could therefore constitute a step forward in data-driven approaches for tomographic image reconstruction embedding the physics model of the data. In this regard, latest works from LMU

Department of Experimental Physics - Medical Physics aim to combine X-ray and ion imaging in order to tackle the intrinsic inaccuracies of the analytical forward-projection model and the HU-RSP semi-empirical calibration: deep-learning architectures have been proposed for the inference of the ion CT image based on the X-ray CT image and sparse ion radiographies, where the physical model could be embedded within the network (Butz et al. [2024]) or given as input in form of back-projection data (Tinelli et al. [2024]).

Bibliography

Gillian C. Barnett, Catherin M. L. West, Alison M. Dunning, Rebecca M. Elliot, Charlotte E. Coles, Paul D. P. Pharoah, and Neil G. Burnet. Normal tissue reactions to radiotherapy: towards tailoring treatment dose be genotype. *Nature Reviews Cancer*, 9(2):134–142, 2009.

Rajaminickam Baskar, Kuo Ann Lee, Richard Yeo, and Kheng-Wei Yeoh. Cancer and Radiation Therapy: Current Advances and Future Directions. *International Journal of Medical Sciences*, 9(3):193–199, 2012.

Adrian C. Begg, Fiona A. Stewart, and Conchita Vens. Strategies to improve radiotherapy with targeted drugs. *Nature Reviews Cancer*, 11(4):239–253, 2011.

Niels Bohr. On the decrease of velocity of swiftly moving electrified particles in passing through matter. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 30(178):581–612, 1915.

M. D. Brooke and S. N. Penfold. An inhomogeneous most likely path formalism for proton computed tomography. *Physica Medica*, 70:184–195, 2020.

Alexander Burker, Thomas Bergauer, Albert Hirtl, Christian Irmler, Stefanie Kaser, B. Knausl, Florian Pitters, and Felix Ulrich-Pur. Single particle tracking uncertainties in ion imaging. *arXiv: Medical Physics*, 2020.

I. Butz, H. A. Loarca, A. Schiavi, V. Patera, O. Oektem, G. Kutyniok, K. Parodi, and C. Gianoli. Data-driven stopping power calibration of the treatment planning x-ray CT using sparse-view ion radiographies. Poster presented at the European Congress of Medical Physics, 2024.

Ravi A. Chandra, Florence K. Keane, Francine E.M. Voncken, and Charles R. Thomas Jr. Contemporary radiotherapy: present and future. *The Lancet*, 398:171–184, 2021.

Omar Chibani. Energy-loss straggling algorithms for Monte Carlo electron transport. *Medical physics*, 29 (10):2374 – 2383, 2002.

C. A. Collins-Fekete, E. Bär, L. Volz, H. Bouchard, L. Beaulieu, and J. Seco. Extension of the Fermi-Eyges most-likely path in heterogeneous medium with prior knowledge information. *Physics in Medicine & Biology*, 62(24):9207–9219, 2017a.

C. A. Collins-Fekete, L. Volz, S. K. Portillo, L. Beaulieu, and J. Seco. A theoretical framework to predict the most likely ion path in particle imaging. *Physics in Medicine & Biology*, 62(5):1777–1790, 2017b.

Charles-Antoine Collins-Fekete, Nikolaos Dikaios, Gary Royle, and Philip M. Evans. Statistical limitations in proton imaging. *Physics in Medice & Biology*, 65(8), 2020.

Laura A. Dawson and Cythia Ménard. Imaging in Radiation Oncology: A Perspective. *The Oncologist*, 15: 338–349, 2010.

Joaquim Teixeira de Assis, Olga Yevseyeva, Ivan Evseev, Ricardo Tadeu Lopes, Hugo reuters Schelin, Márgio Cezar Loss Klock, Sergei A. Paschuk, Reinhard W. Schulte, and David C. Williams. Proton computed tomography as a tool for proton therapy planning: preliminary computer simulations and comparisons with X-ray CT basics. *X-Ray Spectrometry*, 34:481–492, 2005.

Shrinivas D. Desai and Linganagouda Kulkarni. A Quantitative Comparative Study of Analytical and Iterative Reconstruction Techniques. *International Journal of Image Processing*, 4, 2010.

Marco Durante and Harald Paganetti. Nuclear Physics in particle therapy: a review. *Reports on Progress in Physics*, 79(9):59pp, 2016.

Marco Durante and Jay S.Loeffler. Charged particles in radiation oncology. *Nature Reviews Clinical Oncology*, 7:37–43, 2010.

A. Embriaco, E. V. Bellinzona, and A. Fontana, A. Rotondi. On Molière and Fermi-Eyges scattering theories in hadrontherapy. *Physics in Medicine & Biology*, 62(15):6290–6303, 2017.

Leonard Eyges. Multiple Scattering with Energy Loss. *Physical Review Journals*, 74:1534–1535, 1948.

Serena Gianfaldoni, Roberto Gianfaldoni, Uwe Wollina, Jacopo Lotti, Georgi Tchernev, and Torello Lotti. An Overview on Radiotherapy: From Its History to Its Current Applications in Dermatology. *Open Access Macedonian Journal of Medical Sciences*, 5(4):521–525, 2017.

C. Gianoli, M. Göppel, S. Meyer, P. Palaniappan, M. Rädler, F. Kamp, C. Belka, M. Riboldi, and K. Parodi. Patient-specific CT calibration based on ion radiography for different detector configurations in ^{1}H, ^{4}He and ^{12}C ion pencil beam scanning. *Physics in medicine and biology*, 65(24), 2020.

Bernard Gottschalk, Andreas Koehler, Robert J. Scheider, Janet M. Sisterson, and Miles Wagner. Multiple Coulomb Scattering of 160 MeV protons. *Nuclear Instruments & Methods in Pysics Research Section B-beam Interactions With Materials and Atoms*, 74:467–490, 1993.

Cai Grau, Marco Durante, Dietmat Georg, Johannes A. Langendijk, and Damien C. Weber. Particle Therapy in Europe. *Molecular Oncology*, 14(7):1492–1499, 2020.

Virgil L. Highland. Some Practical Remarks on Multiple Scattering. *Nuclear Instruments and Methods*, 129: 497, 1975.

Hyun Do Hu and Seonghoon Kim. History of Radiation Therapy Technology. *Progress in Medical Physics*, 31(3), 2020.

R. F. Hurley, R. W. Schulte, V. A. Bashkirov, A. J. Wroe, A. Ghebremedhin, H. F. W. Sadrozinski, V. Rykalin, G. Coutrakon, P. Koss, and B. Patyal. Water-equivalent path length calibration of a prototype proton CT scanner. *Medical Physics*, 39(5):2438–2446, 2012.

Oliver Jäkel. Heavy ion radiotherapy. *New Technologies in Radiation Oncology. Medical Radiology*, pages 365–377, 2006.

A. M. Koehler. Proton Radiography. *Science*, 160(3825):303–304, 1968.

Lennart R. Koetzier, Domenico Mastodicasa, Timothy P. Szczykutowicz, Niels R. van der Werf, Adam S. Wang, Veit Sandfort, Aart J. van der Molen, Dominic Fleischmann, and Martin J. Willemink. Deep Learning Image Reconstruction for CT: Technical Principles and Clinical Prospects. *Radiology*, 306(3), 2023.

Aafke Christine Kraan. Range Verification methods in particle therapy: underlying physics and Monte Carlo modelling. *Frontiers in Oncology*, 5(150), 2015.

Nils Krah, Denis Dauvergne, Jean Michel Létangan, and Simon Rit Étienne Testa. Relative stopping power resolution in time-of-flight proton CT. *Physics in Medicine & Biology*, 67(16), 2022.

Lev Davidovich Landau. On the energy of fast particles by ionization. *Journal of Physics*, 8(4):201–205, 1944.

S. Lohmann, R. Holeňák, P. L. Grande, and D. Primetzhofer. Trajectory dependence of electronic energy-loss straggling at keV ion energies. *Phys. Rev. B*, 107:085110, Feb 2023.

G. R. Lynch and O. I Dahl. Approximations to multiple coulomb scattering. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 58(1):6–10, 1991.

Yoshitaka Matsumoto, Nobuyoshi Fukumitsu, Hitoshi Ishikawa, Kei Nakai, and Hideyuki Sakurai. A Critical Review of Radiation Therapy: From Particle Beam Therapy (Proton, Carbon, and BNCT) to Beyond. *Journal of Personalized Medicine*, 11(825), 2021.

Stephen Joseph McMahon. The linear quadratic model: usage, interpretation and challenges. *Physics in Medice and Biology*, 64, 2019.

Sebastian Meyer. *On the Clinical Potential of Ion Computed Tomography with Different Detector Systems and Ion Species*. PhD thesis, Ludwig-Maximilians-Universität, Munich, Germany, 2019.

Sebastian Meyer, Chiara Gianoli, Lorena Magallanes, Benedikt Kopp, Thomas Tessonnier, Guillaume Landry, George Dedes, Bernd Voss, and Katia Parodi. Comparative Monte Carlo study on the performance of integration- and list-mode detector configurations for carbon ion computed tomography. *Physics in Medicine & Biology*, 62(3):1096 – 1112, 2017.

Sebastian Meyer, Marco Pinto, Katia Parodi, and Chiara Gianoli. The impact of path estimates in iterative ion CT reconstructions for clinical-like cases. *Physics in medicine and biology*, 69(9), 2021.

Gert Molière. Theorie der Streuung schneller geladener Teilchen II. Mehrfach-und Vielfachstreuung. *Zeitschrift Naturforschung Teil A*, 3:78, 1948.

H. Muraishi, K. Nishimura, S. Abe, H. Satoh, S. Hara, H. Hara, Y. Takahashi, T. Mogaki, R. Kawai, K. Yokoyama, N. Yasuda, T. Tomida, Y. Ohno, and T. Kanai. Evaluation of Spatial Resolution for Heavy Ion CT System Based on the Measurement of Residual Range Distribution With HIMAC. *IEEE Transactions on Nuclear Science*, 56(5):2714–2721, 2009.

Wayne D. Newhauser and Rui Zhang. The physics of proton therapy. *Physics in Medicine & Biology*, 60(8): 155–209, 2015.

Carmen Seller Oria, Sebastian Meyer, Elisabetta De Bernardi, Katia Parodi, and Chiara Gianoli. A Dedicated Tomographic Image Reconstruction Algorithm for Integration-Mode Detector Configuration in Ion Imaging. *2018 IEEE Nuclear Science Symposium and Medical Imaging Conference Proceedings (NSS/MIC)*, pages 1–3, 2018.

Katia Parodi, Alfredo Ferrari, Florian Sommerer, and Harald Paganetti. Clinical CT-based calculations of dose and positron emitter distributions in proton therapy using the FLUKA Monte Carlo code. *Physics in Medicine & Biology*, 52:3369 – 3387, 2007.

S. N. Penfold, R. W. Schulte, Y. Censor, and A. B. Rosenfeld. Total variation superiorization schemes in proton computed tomography image reconstruction. *Medical Physics*, 37(11):5887–5895, 2010.

Scott Penfold and Yair Censor. Techniques in Iterative Proton CT Image Reconstruction. *Sensing and Imaging*, 16(19), 2015.

Gisele C. Pereira, Melanie Traughber, and Raymond F. Muzic Jr. The Role of Imaging in Radiation Therapy Planning: Past, Present and Future. *BioMed Research International*, 2014:9 pages, 2014.

G. Poludniowski, N. M. Allinson, and P. M. Evans. Proton radiography and tomography with application to proton therapy. *The British journal of radiology*, 88(1053), 2015.

Aisya Puspasari and Rafiq Sulistyo. The importance of imaging in radiation therapy: a review article. *International Journal of Radiology and Imaging*, 1(2):17–25, 2022.

I. Rinaldi, S. Brons, J. Gordon, R. Panse, B. Voss, O. Jäkel, and K. Parodi. Experimental characterization of a prototype detector system for carbon ion radiography and tomography. *Physics in Medicine & Biology*, 58(3):413–427, 2013.

A. Schiavi, M. Senzacqua, S. Pioli, A. Mairani, G. Magro, S. Molinelli, M. Ciocca, G. Battistoni, and V. Patera. Fred: a GPU-accelerated fast-Monte Carlo code for rapid treatment plan recalculation in ion beam therapy. *Physics in Medicine & Biology*, 62(18):7482–7504, 2017.

Uwe Schneider and Eros Pedroni. Proton radiography as a tool for quality control in proton therapy. *Medical Physics*, 22(4):353–363, 1995.

W. Schneider, T. Bortfeld, and W. Schlegel. Correlation between CT numbers and tissue parameters needed for Monte Carlo simulations of clinical dose distribution. *Physics in Medicine & Biology*, 42(2):459–478, 2000.

R. Schofield, L. King, U. Tayal, I. Castellano, J. Stirrup, F. Pontana, J. Earls, and E. Nicol. Image Reconstruction: Part 1 - understanding filtered back projection, noise and image acquisition. *Journal of cardiovascular computed tomography*, 14(3):219 – 225, 2020.

R.W. Schulte, S. N. Penfold, J. T. Tafas, and K. E Schubert. A maximum likelihood proton path formalism for application in proton computed tomography. *Medical Physics*, 35(11):4849–4856, 2008.

Jarle Rambo Sølie, Lennart Volz, Helge Egil Seime Pettersen, Pierluigi Piersimoni, Odd Harald Odland, Dieter Röhrich, Håvard Helstrup, Thomas Peitzmann, Kjetil Ullaland, Monika Varga-Kofarago, Shruti Mehendale, Ola Slettevoll Grøttvik, Viljar Nilsen Eikeland, Ilker Meric, and Joao Seco. Image quality of list-mode proton imaging without front trackers. *Physics in medicine and biology*, 65(13), 2020.

Julia Telsemeyer, Oliver Jäkel, and Mária Martišíková. Quantitative carbon ion beam radiography and tomography with a flat-panel detector. *Physics in Medicine and Biology*, 57(23), 2012.

T. Tessonier, T. Marcelos, A. Mairani, S. Brons, and K. Parodi. Phase Space Generation for Proton and Carbon Ion Beams for External Users' Applications at the Heidelberg Ion Therapy Center. *Frontiers in Oncology*, 5(297), 2016.

C. Tinelli, I. Butz, J. Titze, C. Paganelli, G. Baroni, K. Parodi, and C. Gianoli. Deep learning for x-ray CT calibration based on proton radiographies in pre-clinical proton therapy. Poster presented at the European Congress of Medical Physics, 2024.

Felix Ulrich-Pur, Thomas Bergauer, Alexander Burker, Albert Hirtl, Christian Irmler, Stefanie Kaser, Florian Pitters, and Simon Rit. Feasibility of a proton CT system based on 4D-tracking and residual energy determination via time-of-flight. *Physics in Medicine and Biology*, 67(9), 2022.

P. V. Vavilov. Ionization losses of high-energy heavy particles. *Soviet Physics*, 5:749–751, 1957.

Dongxu Wang. A critical appraisal of the clinical utility of proton therapy in oncology. *Medical Devices: Evidence and Research*, 8:439–446, 2015.

G. Wang, J. C. Ye, and B. De Man. Deep learning for tomographic image reconstruction. *Nature Machine Intelligence*, 2(12):737–748, 2020.