



**HAL**  
open science

# Deep Learning-based Methods for Radiotherapy Dose Optimization

Sonia Martinot

► **To cite this version:**

Sonia Martinot. Deep Learning-based Methods for Radiotherapy Dose Optimization. Medical Imaging. Université Paris-Saclay, 2023. English. NNT : 2023UPAST226 . tel-04606852

**HAL Id: tel-04606852**

**<https://theses.hal.science/tel-04606852v1>**

Submitted on 10 Jun 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Deep learning-based methods for radiotherapy dose optimization

*Apprentissage profond pour l'optimisation de dose en  
radiothérapie*

**Thèse de doctorat de l'université Paris-Saclay**

École doctorale n°573, Interfaces  
Spécialité de doctorat : Mathématiques Appliquées  
Graduate School : Sciences de l'ingénierie et des systèmes  
Réfèrent : CentraleSupélec

Thèse préparée dans l'unité de recherche **Mathématiques et Informatique pour la Complexité et les Systèmes (CentraleSupélec, Université Paris-Saclay)** et à **Therapanacea**, sous la direction de **Nikos PARAGIOS**, Professeur à l'Université Paris-Saclay et Président-Directeur Générale de Therapanacea, la co-supervision de **Maria VAKALOPOULOU**, Maîtresse de conférence à CentraleSupélec, et le co-encadrement de **Charlotte ROBERT**, Maîtresse de conférence à l'Institut Gustave Roussy.

**Thèse soutenue à Paris-Saclay, le 15 décembre 2023, par**

**Sonia MARTINOT**

## **Composition du jury**

Membres du jury avec voix délibérative

<b>Daniella THORNWARTH</b> Professeur, Eberhard Karls Universität at Tübingen	Présidente
<b>Vincent LEPETIT</b> Professeur, ENPC ParisTech	Rapporteur & Examineur
<b>David SARRUT</b> Directeur de recherche, CNRS - INSA Lyon	Rapporteur & Examineur
<b>Eric DEUTSCH</b> Professeur, Institut Gustave Roussy	Examineur

**Titre :** Apprentissage profond pour l'optimisation de dose en radiothérapie.

**Mots clés :** Deep Learning, Monte-Carlo, radiothérapie

**Résumé :** La radiothérapie est un pilier du traitement moderne du cancer. La simulation de la dose de radiation avant son administration est une étape primordiale pour garantir la sécurité et la qualité du plan proposé. La méthode Monte-Carlo (MC), reconnue pour simuler le transport des particules et les interactions avec la matière, s'impose comme la référence en termes de précision dans cette tâche. Cependant, ses exigences en termes de calcul entravent son intégration dans le flux de travail clinique.

Le travail présenté dans cette thèse se concentre sur l'exploitation du potentiel de l'apprentissage profond pour accélérer les simulations MC des distributions de dose en radiothérapie. Notre nouvelle approche implique le traitement de séquences de simulations MC de faible précision à l'aide de modèles d'apprentissage profond. Étant donné les importantes données d'entraînement requises pour l'optimisation des modèles d'apprentissage profond, nous avons créé un jeu de données comprenant des simulations MC de réels plans de radiothérapie du monde réel, à l'aide d'un supercalculateur.

Notre recherche explore des architectures à la fois récurrentes et entièrement convolutives, que nous avons adaptées pour gérer des séquences de distributions de dose en 3D. Des études d'ablation ont mis en évidence les limitations inhérentes découlant de la petite taille de notre ensemble de données en 3D. Ce défi nous a conduit à explorer les capacités de nos modèles récurrents dans le cadre d'entraînement faiblement supervisé. L'évaluation quan-

titative et qualitative de nos modèles a montré des performances compétitives comparés aux méthodes classiques et de pointe.

Néanmoins, le calcul de la séquence d'entrée de simulation de dose MC à faible précision requiert toujours des ressources informatiques conséquentes. Pour contourner cette exigence entravante, nous avons élaboré des modèles qui extraient des informations essentielles des simulations MC à faible précision en entrée, sans nécessiter la séquence complète lors de l'inférence. Notre modèle repose sur des représentations vectorielles du nombre de photons simulés lors de la simulation MC des séquences de dose en entrée pour prédire les cartes de dose de haute précision. Cette approche a montré des performances améliorées tout en réduisant la charge informatique.

Enfin, nous présentons une nouvelle classe de fonctions de coût pour optimiser les modèles d'apprentissage profond et compléter davantage notre stratégie d'accélération de la méthode MC. Notre nouveau critère d'optimisation, basé sur le taux de passage de l'indice gamma, s'aligne sur les normes d'évaluation clinique des distributions de dose. Nous proposons une approximation efficace, tensorielle et différentiable de cette métrique pour permettre des entraînements rapides de modèles d'apprentissage profond. De plus, nous fournissons une recette d'entraînement sur mesure qui renforce le potentiel de cette nouvelle fonction de perte innovante. Nos expériences soulignent les avantages significatifs de l'optimisation des modèles avec notre fonction de perte.

**Title :** Deep learning-based methods for radiotherapy dose optimization. **Keywords :** deep learning, generation, Monte-Carlo, radiotherapy.

**Abstract :**

Radiation therapy is a cornerstone of modern cancer treatment. Simulating the radiation dose prior to delivery is a paramount step to ensure the safety and quality of the proposed plan. The Monte-Carlo (MC) method, renowned for simulating particle transport and interactions with matter, stands as the gold standard for this task in terms of precision. However, its computational demands hinder its adoption into clinical radiotherapy workflow.

The work presented in this thesis focuses on leveraging the potential of deep learning approaches to accelerate Monte-Carlo computations of radiotherapy dose distributions. Our novel approach involves processing sequences of low-precision MC simulations using deep learning models. Given the extensive training data required for optimizing deep learning models, we created a dedicated dataset consisting of MC simulations of real-world radiotherapy plans, executed using a supercomputer.

Our research delves into recurrent and fully convolutional deep learning architectures, which we adapted to handle sequences of 3D dose distributions. Ablation studies highlighted the inherent limitations stemming from the small size of our 3D dataset. This challenge led us to explore the capabilities of our recurrent framework within a weakly-supervised setting. Quantitative and qualitative evaluation of

our proposed models demonstrated competitive performance when benchmarked against classical and state-of-the-art methods.

Nonetheless, computing the input low-precision MC dose simulation sequence still results in an unwanted computational bottleneck. To circumvent this impeding requirement, we devised models that extract essential information from the low-precision input MC dose simulations without requiring the full sequence during inference. Our proposed model relies on positional embeddings of the number of photons simulated during MC simulation of the input dose sequences to predict the high precision dose maps. This approach displayed improved performances while reducing the computational overhead.

Finally, we introduce a new class of loss functions to optimize deep learning models and further complete our acceleration strategy. Our novel criterion, based on the gamma index passing rate, aligns with clinical evaluation standards for dose distributions. We propose an efficient, tensorized and differentiable approximation of this metric to enable fast trainings of deep learning models. Additionally, we provide a tailored training recipe that enhances the potential of this innovative loss function. Our experiments underline the significant benefits of optimizing deep learning models with our clinical-based loss function.

# Contents

<b>List of Figures</b>	<b>3</b>
<b>List of Tables</b>	<b>7</b>
<b>1 Introduction</b>	<b>11</b>
1.1 Problem statement . . . . .	12
1.2 Radiotherapy . . . . .	12
1.3 Classical methods for dose computation . . . . .	23
1.4 Monte-Carlo method . . . . .	27
1.5 Radiotherapy in the era of Artificial Intelligence . . . . .	31
1.6 Thesis overview . . . . .	36
<b>2 Monte-Carlo Dataset</b>	<b>43</b>
2.1 Patient cohort . . . . .	44
2.2 Simulation of dose distributions . . . . .	45
2.3 Experimental setup . . . . .	53
2.4 Conclusion . . . . .	59
<b>3 Denoising Monte-Carlo simulations</b>	<b>61</b>
3.1 Introduction . . . . .	62
3.2 ConvLSTMs for 2D denoising . . . . .	65
3.3 3D-ConvLSTM . . . . .	70
3.4 Ablation studies . . . . .	76
3.5 Weakly-Supervised ConvLSTMs . . . . .	79
3.6 Conclusion . . . . .	82

<b>4</b>	<b>Accelerating Monte-Carlo with intermediate representations</b>	<b>83</b>
4.1	Introduction . . . . .	83
4.2	Cascaded Denoising . . . . .	84
4.3	Deep Particle Embeddings . . . . .	87
4.4	Conclusion . . . . .	91
<b>5</b>	<b>Deep Gamma index-based Loss functions</b>	<b>93</b>
5.1	Introduction . . . . .	94
5.2	Related work . . . . .	95
5.3	Definitions . . . . .	96
5.4	Methods . . . . .	97
5.5	Experimental design . . . . .	99
5.6	Results . . . . .	102
5.7	Conclusion . . . . .	106
<b>6</b>	<b>Discussion and Future work</b>	<b>107</b>
6.1	Main findings . . . . .	107
6.2	Limitations and future work . . . . .	110
	<b>Remerciements (Acknowledgements)</b>	<b>117</b>



# List of Figures

1.1	Radiation therapy treatment workflow. . . . .	13
1.2	Inelastic collisions of an electron with an atom. . . . .	17
1.3	Principal photon-matter interactions. . . . .	19
1.4	Schematics of the head of a LINAC (Fiak et al. 2021). . . . .	21
1.5	Comparison of dose delivery IMRT versus VMAT . . . . .	23
1.6	Comparison of classical methods for dose simulation in photon beam radiotherapy (De Martino et al. 2021). . . . .	25
1.7	Applications of AI in the radiation therapy workflow (Huynh et al. 2020). . . . .	31
1.8	CT scan with overlapped delineation of organs involved in radiation therapy (Martini et al. 2020). . . . .	33
1.9	Example of a DL-based dose prediction pipeline (X. Chen et al. 2023). . . . .	35
1.10	Patients' CT and corresponding normalized dose distributions com- puted with Monte-Carlo at various numbers of simulated particles. . . . .	37
2.1	Available data for each patient in the cohort. . . . .	44
2.2	Visualizations of the OpenGATE modelization of the lower part of a LINAC's head. . . . .	47
2.3	Spatial and energy distributions of the particles comprised in the phase space. . . . .	49
2.4	Supercomputer computation strategy of MC simulations. . . . .	50
2.5	Example of a patient's CT, corresponding MC simulations at various noise levels and associated uncertainty maps. . . . .	52
2.6	Violin plots of uncertainty values accross dataset vs number of sim- ulated particles. . . . .	52

2.7	Creation of intermediate noisy simulations from the computed sub-simulations. . . . .	54
3.1	Aggregation procedure of the sub-simulations to create input dose sequence. . . . .	63
3.2	Global studied learning pipeline. . . . .	64
3.3	Architecture of an LSTM cell. . . . .	65
3.4	Training and inference pipelines of the 2D ConvLSTM-based model. . . . .	67
3.5	Dose distributions generated by the model and NLM. . . . .	69
3.6	Visualization of gamma index matrix for the NLM and our model's denoised dose distributions. . . . .	69
3.7	Proposed architecture with stacked 3D ConvLSTM cells. . . . .	71
3.8	3D Unet enhanced with 3D ConvLSTM cells in the skip connections. . . . .	72
3.9	Comparison of dose distributions generated by 3D ConvLSTMs and benchmarked methods. . . . .	75
3.10	Comparison of dose profiles and DVHs generated by 3D ConvLSTMs and benchmarked methods. . . . .	75
3.11	Noise2Noise training pipeline applied to our sequence denoising task. . . . .	79
3.12	Visual comparison of denoised volumes by each model and the ground truth. . . . .	81
4.1	Pipeline for end-to-end dose denoising with cascaded UNets. . . . .	85
4.2	Architecture of the Deep Particle Embedding Model and denoising pipeline. . . . .	89
4.3	Denoising results - dose prediction and gamma index matrix - of a 1e9 simulation with the deep particle embedding model. . . . .	90
5.1	Pipeline of training using sigmoid-GPR-based loss functions. . . . .	100
5.2	Boxplots representing the evaluation metrics achieved by trained models for each case in the test set depending on the loss function used for training . . . . .	104
5.3	Qualitative comparison of the dose distributions and corresponding gamma index matrices depending on the loss function chosen during training. . . . .	105

5.4	Boxplots of execution times of our proposed approach, of the SSIM and the exhaustive search method on 3D dose distributions. . . . .	106
-----	--	-----



# List of Tables

2.1	Number of available dose volumes per patient in the dataset, depending on the number of simulated particles. . . . .	55
3.1	Quantitative evaluation of the dose generated by our model and the NLM method. . . . .	68
3.2	Evaluation metrics for the performance of the models on the test set.	74
3.3	Impact of the sequence length on quantitative performances of 3D ConvLSTMs results. . . . .	77
3.4	Lightweight 3DConvLSTM: Impact of CT addition to the input of the model. . . . .	77
3.5	Impact of complexity reduction of 3D ConvLSTMs on quantitative metrics. . . . .	77
3.6	Quantitative results per anatomy group on the test set for the lightweight 3D ConvLSTM model. . . . .	78
3.7	Quantitative results for the different evaluated models. . . . .	81
4.1	GPR (3mm/3%) values computed on the test set in areas within 20% of the maximum dose. . . . .	86
4.2	Evaluation metrics for the performance of the models on the test set.	91
5.1	Evaluation metrics over the dose distributions comprised in the test set. Different benchmarks over the considered loss functions for different metrics are highlighted with their mean and standard deviation. With bold we indicate the best performing methods per metric. . . . .	103

5.2	Speed comparison of metrics computed over 2D and 3D dose distributions. . . . .	103
-----	---	-----

# Acronyms

**AI** Artificial Intelligence.

**CCC** Collapse Cone Convolution.

**ConvLSTM** Convolutional Long-Short Term Memory.

**CPU** Central Processing Unit.

**CT** Computed Tomography.

**GPR** Gamma-index Passing Rate.

**GPU** Graphical Processing Unit.

**Gy** Gray.

**IMRT** Intensity-Modulated Arc Therapy.

**LBTE** Linear Boltzmann Transport Equation.

**LINAC** Linear Accelerator.

**MC** Monte-Carlo.

**MLC** Multi-Leaf Collimator.

**MRI** Magnetic Resonance Imaging.

**MSE** Mean Squared Error.

**N2N** Noise2Noise.

**NLM** Non Local Means.

**OAR** Organ-at-Risk.

**ORL** Oto-Rhino-Laryngology.

**PB** Pencil Beam.

**PTV** Planning Target Volume.

**RT Dose** Radiation Therapy Dose.

**RT Struct** Radiation Therapy Structure Set.

**RT Plan** Radiation Therapy Plan.

**SSIM** Structural Similarity Index Measure.

**Sv** Sievert.

**VMAT** Volumetric Modulated Arc Therapy.

# Chapter 1

## Introduction

### Contents

---

<b>1.1</b>	<b>Problem statement</b>	<b>12</b>
<b>1.2</b>	<b>Radiotherapy</b>	<b>12</b>
1.2.1	Treatment Process	12
1.2.2	Physics of Radiotherapy	16
1.2.3	Linear Accelerator & Dose Delivery	20
<b>1.3</b>	<b>Classical methods for dose computation</b>	<b>23</b>
1.3.1	Basics of dose computation	23
1.3.2	Non-MC dose computation engines	24
<b>1.4</b>	<b>Monte-Carlo method</b>	<b>27</b>
1.4.1	Mathematical Framework	27
1.4.2	Monte-Carlo for Radiotherapy	28
<b>1.5</b>	<b>Radiotherapy in the era of Artificial Intelligence</b>	<b>31</b>
1.5.1	AI and Radiotherapy	31
1.5.2	Deep Learning and Monte-Carlo dose computation	35
<b>1.6</b>	<b>Thesis overview</b>	<b>36</b>
1.6.1	New perspective on Monte-Carlo simulations	38
1.6.2	Monte-Carlo dataset of dose distributions	38
1.6.3	Designing a learning loss function for Monte-Carlo	39
1.6.4	Overview of the contributions	39

---

## 1.1 Problem statement

Deep Learning, a subfield of Artificial Intelligence (AI), stands as a transformative force that has revolutionized a multitude of industries through its remarkable capability to discern complex data patterns and relationships. In this era of technological advancements, the integration of Deep Learning techniques into the realm of medical physics and radiation oncology presents a promising frontier. Its potential lies in simplifying and enhancing the intricate processes associated with radiotherapy treatments.

In recent decades, radiotherapy has risen as one of the most effective cancer treatments, offering a localized approach to deliver precise doses of ionizing radiation to malignant tissues, while sparing surrounding healthy tissue. The linchpin of this precision lies in the accuracy of radiotherapy treatment planning, which is critical for optimizing tumor control and minimizing collateral damage. To this end, simulating the irradiation dose before actual delivery is a key quality assurance step in treatment planning. The Monte-Carlo method (Kase et al. 1978; Duderstadt et al. 1979), renowned for simulating particle transport, has become a cornerstone of radiotherapy dose calculations, offering a comprehensive and detailed representation of radiation transport and interactions. However, it comes with a significant computational intensity, which hampers its potential for real-time or near-real-time dose calculations, thereby limiting its immediate clinical applicability.

The focus of this doctoral thesis embarks on a journey to explore and harness the power of Deep Learning to expedite Monte-Carlo simulations in the context of radiotherapy dose calculation. In this introduction, we delve into the intricacies of the radiotherapy treatment process, reiterate the fundamental mathematical framework of Monte-Carlo simulation within the context of dose generation, and provide a comprehensive overview of the multifaceted challenges addressed and the contributions made within this thesis. By leveraging the strengths of Deep Learning in addressing these challenges, we aim to open up new horizons in the field of radiation oncology.

## 1.2 Radiotherapy

### 1.2.1 Treatment Process

The radiotherapy treatment process (DeVita Jr et al. 2019) is a highly specialized procedure designed to deliver precise doses of radiation to cancerous tumors while minimizing damage to

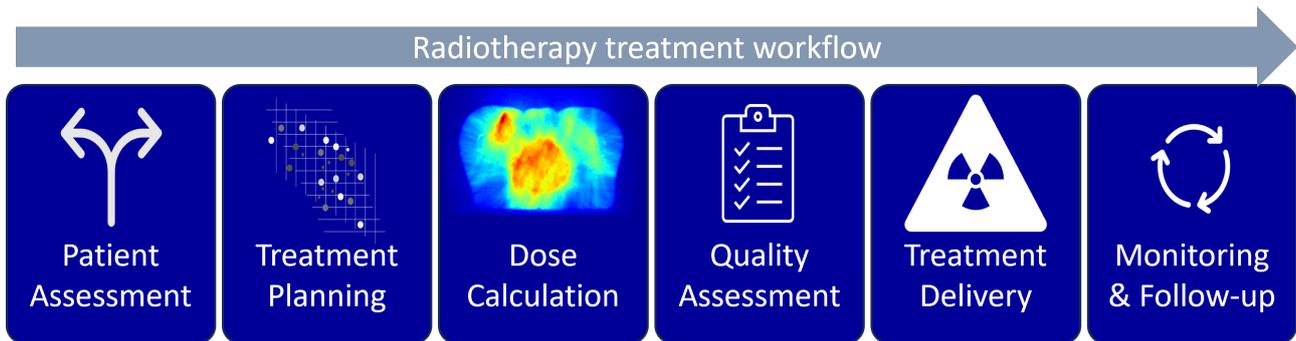


Figure 1.1: **Radiation therapy treatment workflow.** The image gives an overview of the workflow a patient undergoes in radiation therapy. The process starts by a consultation that determines the course of treatment with radiotherapy, followed by a treatment planning step that involves the acquisition of medical images and subsequent design of a treatment plan. As dose calculation step is then required to simulate the plan and perform a quality assurance review prior to delivery. After treatment, the patient is monitored and receives follow-up care.

surrounding healthy tissues. The process involves multidisciplinary collaboration among radiation oncologists, medical physicists, radiation therapists, dosimetrists, and other healthcare professionals. This thorough approach aims to deliver effective cancer treatment while prioritizing patient safety and quality of life. Depicted in Figure 1.1, we give a detailed description of the radiotherapy treatment process:

**Consultation and Assessment:** The treatment process begins with an initial consultation between the patient and the radiation oncologist. During this consultation, the oncologist reviews the patient’s medical history, diagnostic imaging, biopsy results, and other relevant information. The oncologist discusses the potential benefits, risks, and expected outcomes of radiotherapy treatment. If radiotherapy is deemed appropriate, a personalized treatment plan is formulated.

**Simulation and Treatment Planning:** In this stage, the patient undergoes a simulation session. The patient is positioned on the treatment table in the same manner they will be during actual treatment. Specialized imaging techniques are used to precisely define the tumor’s location and shape. These techniques comprise Computed Tomography (CT) (Oldendorf 1978) scans and Magnetic Resonance Imaging (MRI). CT scans are medical imaging tests that create pictures inside selected areas of the body using ionizing radiation. MRI uses strong magnetic fields and radio waves to yield detailed images of the insides of a body (Liang et al. 2000). CT provides high spatial resolution, making it excellent for visualizing anatomical structures

and especially bones. It is particularly useful for quick imaging, making it suitable for patients who may have difficulty remaining still for extended periods. CT is widely available, cost-effective, and has techniques to reduce metal artifacts, which is beneficial for patients with metal implants. However, CT has limitations in soft tissue contrast, making it challenging to distinguish between certain soft tissues. On the other hand, MRI offers superior soft tissue contrast, making it ideal for delineating tumors and identifying normal structures (Schmidt et al. 2015). It provides multi-parametric imaging, including functional and metabolic data, through various techniques. MRI does not use ionizing radiation, a significant advantage when repeated imaging is necessary, and is preferred for imaging certain organs, such as the brain and pelvic organs. Despite these advantages, MRI has drawbacks, including longer imaging times, sensitivity to metal causing artifacts, and limited bone visualization. Additionally, MRI machines may be less widely available and more expensive than CT scanners. Thus, the choice between CT and MRI depends on the specific clinical requirements, the type of information needed, and the characteristics of the tissues being imaged (V. Khoo et al. 1999; Vincent S Khoo et al. 2000; Chandarana et al. 2018). Using these medical images, the oncology team also identifies adjacent healthy organs and tissues to be spared from radiation. The information obtained from CT or MRI scans is used to create a three-dimensional model of the patient's anatomy, which is then employed to design a detailed treatment plan that outlines the dose, angles, and beams to be used during the treatment.

**Dose Calculation:** The density information provided by CT scans is crucial for accurately calculating the radiation dose that will be delivered to the tumor and surrounding tissues. Different tissues absorb and scatter radiation differently, so having accurate density information is essential for precise dose calculations. With the information from the simulation of the dose, advanced treatment planning software is used to calculate the optimal radiation dose distribution. The goal is to ensure that the tumor receives the prescribed dose while minimizing radiation exposure to healthy tissues. Computing the dose from the prescription allows to design the best possible plan. Consequently, the quality of the plan depends on the accuracy of the dose simulation (Nelms et al. 2013; Masi et al. 2013). Moreover, when a treatment plan has been defined, meaning that all machine parameters have been set, dose simulation is paramount to check that the plan is sound. Machine parameters encompass the position of beam shapers in the linear accelerators, intensity of the beam, rate at which the dose is delivered, speed of rotation of the machine around the patient, to name a few. Hence, the high complexity of the treatment calls for precise dose simulation tools to ensure the safety and accuracy of the

parametrization.

**Quality Assurance:** Before treatment begins, the treatment plan undergoes a rigorous quality assurance process. These guidelines cover various aspects of radiotherapy, including equipment commissioning, procedures, and safety (D. I. Thwaites et al. 2005; Merwe et al. 2017; Kutcher et al. 1994). The plan is reviewed by medical physicists and radiation therapists to verify its accuracy and safety. This step is essential to prevent errors and ensure that the calculated dose is delivered precisely as intended.

**Treatment Delivery:** During the treatment sessions, the patient lies on the treatment table in the exact position determined during simulation. Highly specialized equipment, such as linear accelerators, delivers the prescribed radiation beams according to the treatment plan. Modern radiotherapy techniques, such as intensity-modulated radiation therapy and image-guided radiation therapy, allow for precise targeting and adjustment of the radiation beams based on daily imaging.

**Monitoring and Follow-Up:** Throughout the treatment course, the patient's response to treatment is carefully monitored. Radiation therapists, medical physicists, and radiation oncologists work together to ensure that the treatment parameters are accurately delivered. Periodic imaging and clinical assessments are conducted to evaluate the tumor's response and any potential side effects.

**Completion of Treatment:** Once the prescribed number of treatment sessions is completed, the patient finishes the radiotherapy course. The treatment team may provide recommendations for managing any lingering side effects and schedule follow-up appointments to assess the treatment's effectiveness.

**Post-Treatment Follow-Up:** After treatment completion, the patient continues to be monitored during follow-up visits. Imaging and clinical evaluations help track the tumor's response and any potential recurrence. Long-term follow-up is essential to ensure the treatment's success and address any late-emerging side effects.

## 1.2.2 Physics of Radiotherapy

In radiotherapy, ionizing radiation is used to treat cancer by delivering high doses of radiation to tumor cells. In this section, we describe the fundamental mechanisms underpinning ionizing radiation.

**Ionizing Radiation:** Ionizing radiation (Effects of Atomic Radiation et al. 1982; Mettler 1985) is a type of electromagnetic radiation such as X-rays and gamma rays, or particle radiation. This radiation carries enough energy to dislodge tightly bound electrons from atoms or molecules. This results in the formation of particles also called ions. These charged particles can be electrons, protons, alpha particles or heavy ions. Particle ionizing radiation includes alpha particles, beta particles, and neutrons. Alpha particles are made up of two protons and two neutrons, beta particles are high-energy electrons or positrons, and neutrons are uncharged particles found in atomic nuclei. On the other hand, electromagnetic ionizing radiation encompasses X-rays and gamma rays. X-rays are produced by accelerating electrons to collide with a target material, while gamma rays originate from the radioactive decay of certain atomic nuclei. Both X-rays and gamma rays are high-energy photons. Photons are neutral particles therefore this type of radiation is referred to as indirectly ionizing radiation (Rosenberg 2008).

Ionizing radiation in the context of radiotherapy is quantified using units such as the Gray (Gy), with the convention that  $1\text{Gy} = 1\text{ J/kg}$ , and the Sievert (Sv). The gray measures the absorbed dose of ionizing radiation, while the sievert takes into account the biological effect of the radiation, incorporating factors such as the type of radiation and the tissue affected. Outside of any medical context, ionizing radiation can damage biological tissues by causing ionization and disruption of chemical bonds within cells. This can lead to various effects, including DNA damage, cell death, and increased risk of cancer. The severity of these effects depends on the type of radiation, the dose received, and the sensitivity of the exposed tissue.

**Electron interactions with matter:** Ionizing electromagnetic and particle radiation occur when an accelerated electron enters the neighbourhood of an atom nucleus. The electron then undergoes Coulomb interactions with the atom, comprising of interactions with atomic orbital electrons of the atom or directly with the atom nucleus. Through these interactions, the electron may either lose its kinetic energy through collisions, resulting in radiation, or change its direction which is called scattering. Collisions can occur with the orbital electrons of the atom or with the nucleus and have various effects whether the collision is elastic or inelastic. During an elastic collision, the electron is merely deflected from its original course but suffers

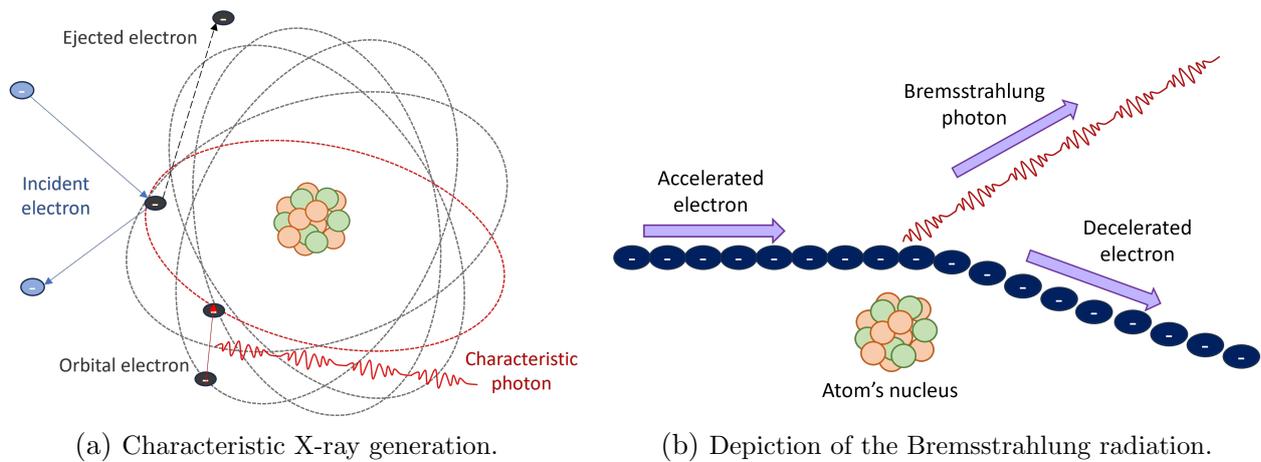


Figure 1.2: **Inelastic collisions of an electron with an atom.** This figure provides schematics explaining the generation of characteristic photons caused by atomic de-excitation and of bremsstrahlung photons caused by electron deceleration.

no energy loss. However, in an inelastic collision, the electron loses part of its kinetic energy which corresponds to the emission of radiation. Figure 1.2 sum up the mechanism behind the two types of inelastic collisions that an electron can encounter with an atom (Podgorsak 2005).

During an inelastic collision with orbital electrons of the atom, the incident electron can knock inner-shell electrons from the atom, and is deflected from its path and loses kinetic energy as well. These Coulomb interactions result in either atomic ionization followed by atomic excitation. As depicted in Figure 1.2a, the orbital electron is ejected, provided the energy of the incident electron is greater than the binding energy of orbital electron. The empty space left is then filled by another outer shell electron, emitting a single X-ray photon, also called a characteristic photon, with an energy level equivalent to the energy level difference between the outer and inner shell electron involved in the transition.

During inelastic collisions with the atom's nucleus, the loss in kinetic energy is translated in the form of bremsstrahlung radiation (Blumenthal et al. 1970). Bremsstrahlung, a German term meaning "braking radiation," occurs when fast-moving electrons close to the speed of light interact with the nuclei of atoms within a target material. As displayed in Figure 1.2b, when high-energy electrons are decelerated by the electric field of atomic nuclei in a target material, they lose part of their kinetic energy in the form of a bremsstrahlung photon, which includes X-rays and gamma rays. The energy of the emitted gamma rays depends on the energy lost by the electrons during this interaction. The process can be described and quantified, making gamma rays a type of electromagnetic radiation with very high energy which falls at the highest

energy end of the electromagnetic spectrum. The energy  $E_\gamma$  of the emitted photon, i.e. of the gamma ray, can be calculated using the following formula:

$$E_\gamma = \hbar f$$

Where:

- $E_\gamma$  is the energy of the emitted photon (in joules).
- $\hbar$  is Planck's constant -  $\hbar = 6.626 \times 10^{-34}$  J.s.
- $f$  is the frequency of the emitted photon in hertz.

The frequency  $f$  of the emitted photon depends on the energy lost by the electron during the interaction with the target material. The higher the energy loss, the higher the frequency, and therefore the energy of the emitted gamma ray. The energy lost by an electron  $\Delta E_\gamma$  as it is decelerated by the electric field of a nucleus can be calculated using the following formula:

$$\Delta E_\gamma = \frac{e^2}{2r}$$

Where:

- $\Delta E_\gamma$  is expressed in joules.
- $e$  is the elementary charge of the electron and  $e = 1.602 \times 10^{-19}$  C.
- $r$  is the distance of closest approach between the electron and the nucleus in meters.

The frequency of the emitted photon is related to the energy loss of the electron as follows:

$$f = \frac{\Delta E_\gamma}{\hbar}$$

The result of these interactions is the emission of X-ray photons, both characteristic and bremsstrahlung, from the target material. These X-rays emerge in all directions from the target, forming a broad spectrum of X-ray energies, with some of them being very high-energy X-rays. It's important to note that the majority of the X-rays generated in radiation therapy are produced through the bremsstrahlung process. Thus, the target material usually comprises atoms with high atomic numbers, which have a lower electron density, to favor collisions between the incident electron and the nuclei to emit bremsstrahlung radiation.

**Ionizing photon radiation:** In the context of radiotherapy in this thesis, we focus on radiation from photon beams. Ionizing photon beams are the preferred technique, as they allow to target deep-seated tumors (Mayles et al. 2007). In penetrating an absorbing medium, photons can experience various stochastic interactions with the atoms of the medium which can comprise whole atoms, molecules, atomic orbital electrons or nuclei. Photons can randomly undergo several interactions. These interactions create new secondary ionizing particles which can be charged particles like electrons or uncharged particles such as secondary photons. Secondary charged particles deposit their energies near the location of the interaction. Following an interaction, the photon faces two outcomes. The photon can be absorbed completely by the medium while the rest of its energy is transferred to secondary charged particles. Alternatively, the photon is scattered, meaning that the secondary particles are necessarily another photon which can be accompanied by a secondary charged particle. The principal interactions observed at megavoltage energy range are pair production, the photoelectric effect and Compton scattering and are displayed in Figure 1.3.

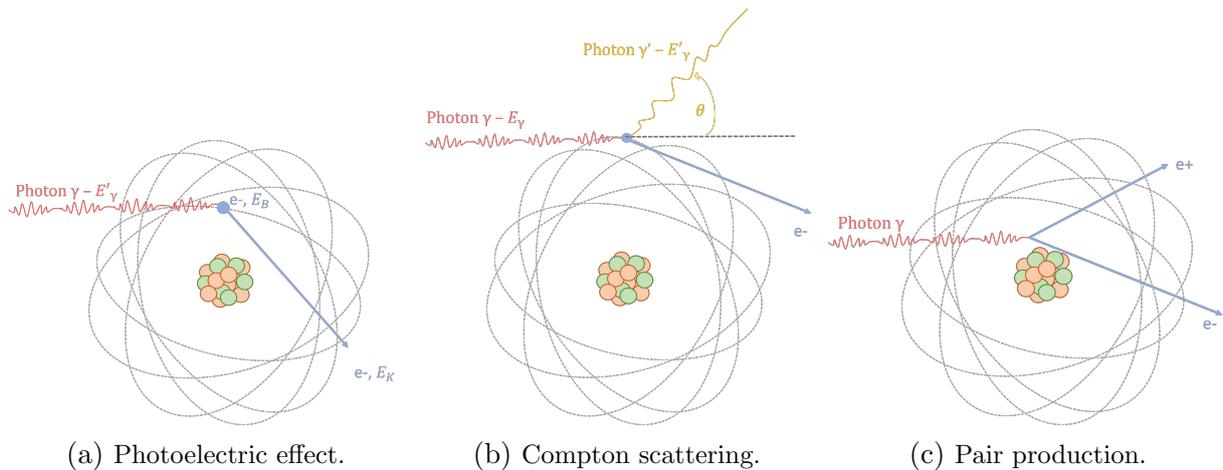


Figure 1.3: **Principal photon-matter interactions.** This figure details key interactions between photons and matter to better understand the mechanisms leading to radiation dose deposition. Each of these three interactions yields secondary ionizing particles that will in turn also interact with matter and deposit energy.

Photon-orbital electron interactions depend on the closeness of the electron to the nucleus. When considering a tightly bound inner shell electron, the photon undergoes the photoelectric interaction. As depicted in Figure 1.3a, the photon gets absorbed by the inner shell electron and all of its energy is transferred to the electron. As a result, the photon disappears and part of the photon's energy that was transferred to the electron overcomes the binding energy, thus ejecting the electron from the inner shell. The remaining of the energy becomes the kinetic

energy of the scattered electron. The vacancy in the inner shell is then filled by an electron from a higher energy shell which can create a characteristic x-ray photon. Let  $E_K$  and  $E_B$  be the kinetic energy and the binding energy of the ejected orbital electron, then  $E_K = E_\gamma - E_B$ . In general, the photoelectric effect dominates at low photon energy.

When interacting with a loosely bound outer shell electron with negligible binding energy, a photon can undergo Compton scattering: the photon transfers some of its energy to the electron and is deflected, while the electron is ejected. The amount of energy transferred to the scattered outer shell electron depends on the scattering angle  $\theta$ , as indicated in Figure 1.3b. The more energetic the incident photon, the more forward, i.e.  $\theta$  small, is the scatter. In general, the Compton effect dominates at intermediate photon energies. Let  $E_\gamma$  and  $E'_\gamma$  be the energies of respectively the incident photon and the deflected photon, then the energy of the deflected photon follows:

$$E'_\gamma = \frac{E_\gamma}{1 + \frac{E_\gamma}{511\text{keV}} \times (1 - \cos\theta)}$$

When a photon interacts with nuclei by getting close to it as depicted in Figure 1.3c, it undergoes either photodisintegration or interacts with the electrostatic field of the nucleus. In the latter case, the pair production phenomenon occurs, which creates an electron-positron pair near the nucleus. In general, pair production is most frequent a high photon energy (Rosenberg 2008).

In photon beam radiation therapy, the radiation dose is deposited in the patient's body following the principal photon interactions described above. These interactions attenuate the photon beam following Lambert-Beer law (Beer et al. 1852):

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

Where  $I(x)$  is the intensity of the radiation beam at depth  $x$ ,  $I_0$  is the initial intensity and  $\mu$  is the linear attenuation coefficient of the traversed medium. When penetrating a patient's body, the beam is thus attenuated depending on various factors, such as tissue density, radiation beam characteristics and patient anatomy.

### 1.2.3 Linear Accelerator & Dose Delivery

**Linear Accelerator (LINAC):** LINACs (D. Thwaites et al. 2006) are used in various radiation therapy techniques. They play a critical role in the treatment of cancer by delivering high-energy radiation precisely to the tumor, either to shrink it before surgery, as a primary treatment, or to relieve symptoms in palliative care. The goal is to maximize the therapeutic

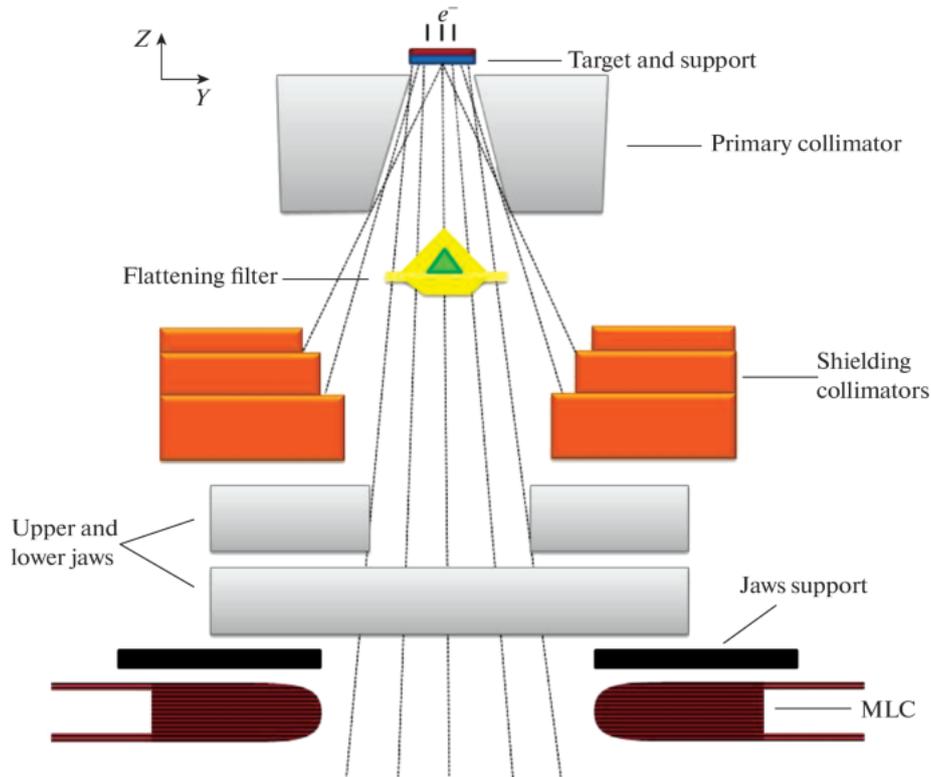


Figure 1.4: **Schematics of the head of a LINAC (Fiak et al. 2021)**. The figure provides a simplified view of the elements that compose a LINAC's head, from which the radiation is generated. The bremsstrahlung radiation and other ionizing particles are generated when electrons reach the target composed of a heavy metal. Collimators, jaws and the multi-leaf collimator are element that participate in shaping the radiation beam to the desired shape.

effect on the tumor while minimizing harm to normal tissues and organs.

A LINAC functions by producing high-energy electron beams for radiation therapy. We distinguish treatment using electron beams for superficial tumors from photon beams that better target deep tumors within the patient (Mayles et al. 2007). To generate a photon beam, the electrons enter an accelerating waveguide, in which radio-frequency waves are used to accelerate the electrons. A waveguide is typically a metallic structure that guides the electron beam. Strong magnetic fields are used to steer and focus the electron beam as it travels through the LINAC and ensure precise control of the beam's position and direction. Using the magnetic fields, the electrons are aimed at a heavy metal target like tungsten to generate X-rays via bremsstrahlung radiation, which we developed in subsection 1.2.2. High energies are necessary in order for the bremsstrahlung emission direction to follow the electron's initial course. Typical high energies start from 6 MeV. The emitted X-rays follow a spectrum of

energy that presents high peaks as a result of characteristic x-rays emissions combined with bremsstrahlung radiation. In order to compensate for the natural intensity variation of the X-ray beam produced by the metal target, a flattening filter is placed in the LINAC's head as depicted in yellow in Figure 1.4 to ensure uniform dose distribution across the treatment field.

The radiation beam is precisely shaped with collimators and beam shapers to precisely target the cancer cells, while minimizing exposure to healthy tissue. Upper and lower jaws, displayed in Figure 1.4 as grey rectangles are mechanical devices positioned near the beam exit. They allow adjustment of the radiation beam's width and height. By moving the jaws, the beam's shape is tailored to match the tumor's dimensions. The Multi-Leaf Collimator (MLC), in dark red in Figure 1.4, is a critical part of the LINAC and is usually located close to the patient. It consists of numerous individual "leaves" that can move independently. The positions of these leaves can be adjusted to create complex and customizable shapes for the radiation beam. By dynamically repositioning the leaves during treatment, the MLC allows precise sculpting of the radiation dose, conforming it to the tumor's shape and minimizing exposure to surrounding healthy tissues. This dynamic control ensures accuracy and safety in radiation therapy.

**Dose delivery:** Within the realm of treatment planning, two sophisticated strategies, Intensity-Modulated Arc Therapy (IMRT) (Bortfeld 2006) and Volumetric Modulated Arc Therapy (VMAT) (Otto 2008), have emerged as leading contenders. IMRT manipulates radiation intensity across multiple static beams or arcs with intensity modulation, sculpting the radiation dose to conform precisely to the tumor's shape. Conversely, VMAT introduces an additional dimension of flexibility by orchestrating continuous gantry motion coupled with continuous modulation, optimizing both treatment efficiency and precision.

Figure 1.5 compares the dose deposited in the case of IMRT and VMAT delivery types. Thanks to the continuous movement of the rotating beam, the VMAT approach brings additional flexibility in the shape of the beam allow to reduce drastically irradiation outside of the tumor. VMAT tends to be faster and more efficient in terms of treatment delivery (Quan et al. 2012), making it the preferred choice for many clinical scenarios, especially when treatment time is a critical factor.

These diverse approaches coupled with the progressive sophistication of LINACs, empower clinicians to tailor radiotherapy treatments with an unprecedented level of customization, ensuring maximal therapeutic benefit while minimizing potential side effects. As radiotherapy techniques become increasingly sophisticated, the demand for personalized, adaptive treatments has grown significantly. Real-time dose simulations empower clinicians to monitor and

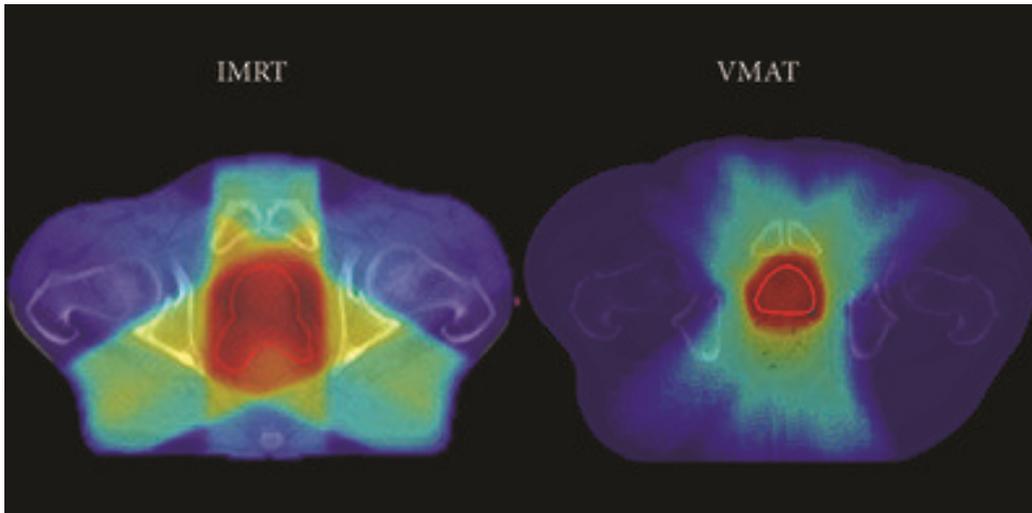


Figure 1.5: **Comparison of dose delivery IMRT versus VMAT**(Vanneste et al. 2016). This figure provides a qualitative comparison of the radiation dose delivered in the case of IMRT and VMAT. The figure highlights the homogeneity of the dose delivered by a VMAT plan in contrast with IMRT where the beams deliver a less complex dose. The dose delivered by VMAT is more precise and tailored to spare surrounding tissues.

respond to dynamic changes in a patient’s anatomy during treatment, ensuring that the radiation precisely targets the tumor while sparing healthy tissues. This adaptive approach enhances treatment efficacy and minimizes side effects.

## 1.3 Classical methods for dose computation

In subsections 1.2.2 and 1.2.3, we described the physical mechanisms underpinning radiation dose deposition and presented the main means of delivery - LINAC, VMAT - of that dose in the context of external photon beam radiation therapy. In this section, we present the basics of dose computation and describe the principal dose computation methods that are not Monte-Carlo.

### 1.3.1 Basics of dose computation

**Dose:** Let us define the radiation dose. Let  $m$  be the mass of the penetrated medium and  $\epsilon$  is the net energy transferred to a unit volume when a particles enters and leaves it. Then the dose is the energy deposited per unit mass, i.e.:

$$D = \frac{d\epsilon}{dm}$$

**Fluence:** Let  $N$  be the number of particles and  $x$  the unit area. Let  $\phi(x, \epsilon) = \frac{dN}{dx}$ , expressed in  $mm^{-2}$  be the number of particles crossing a unit area, also called fluence. We can write fluence as differential in energy:  $\phi_E = \frac{d\phi}{dE}$ , in which case we derive the total fluence by integrating over  $E$ . We define energy fluence which, with  $E$  the energy carried by the incident particle, is defined by:

$$S(x, E) = \frac{d(EN)}{dx}$$

**Attenuation coefficient:** We define  $\mu$  the attenuation coefficient of the traversed medium of density  $\rho$ . Let  $N_A$  the Avogadro's number,  $A$  the relative atomic mass of the target medium, then  $\mu = \frac{N_A}{A}\rho$  is the probability of interactions in a medium per unit area. The ratio  $\frac{\mu}{\rho}$  is called the mass attenuation coefficient and does not depend on  $\rho$ .

**TERMA:** The total energy released per unit mass (TERMA) describes the total energy imparted to the medium by electrons in a certain region of interest. It is a quantity used to quantify the energy deposition by electrons in a given volume of tissue or other medium irradiated with incident photons of energy spectrum  $E \in [0, E_{max}]$  and is written as:

$$T = \int_0^{E_{max}} T_E dE \quad \text{with} \quad T_E = S(x, E) \frac{\mu}{\rho}$$

**Spread kernel:** Spread kernels (Ahnesjö et al. 1987) are mathematical functions that describe how radiation beams spread and deposit energy in a medium. Spread kernels  $k(x, E)$  are typically derived from extensive measurements and calculations, often using Monte Carlo simulations, to accurately characterize and account for the behavior of radiation, including scattering, photoelectric electrons, pair production particles and bremsstrahlung photons.

**Total absorbed dose:** Given the prior definitions, we can define the total absorbed dose in a medium as a convolution at a given location  $x$ :

$$D(x) = \frac{1}{\rho(x)} \int_E \int_V T_E \rho(r) k(r, E) dr dE$$

### 1.3.2 Non-MC dose computation engines

In this subsection, we provide a non-exhaustive review of the main dose computation engines that are not Monte-Carlo based. In order to validate any treatment planning using these

techniques, dose calculation is paramount to ensure the treatment plans effectively delivers the correct dose. Because of the complexity of the treatment plans, both in terms of particle-matter interactions and machine settings, computing the corresponding deposited dose can be challenging. To that end, several dose calculation engines have been developed and have several key components:

- Account for the attenuation of the radiation beam, characterized by the electron density derived from the patient’s CT.
- Consider scattering and secondary radiation production.
- Model the radiation beam by considering the beam energy, shape and modifiers (collimators and MLC).

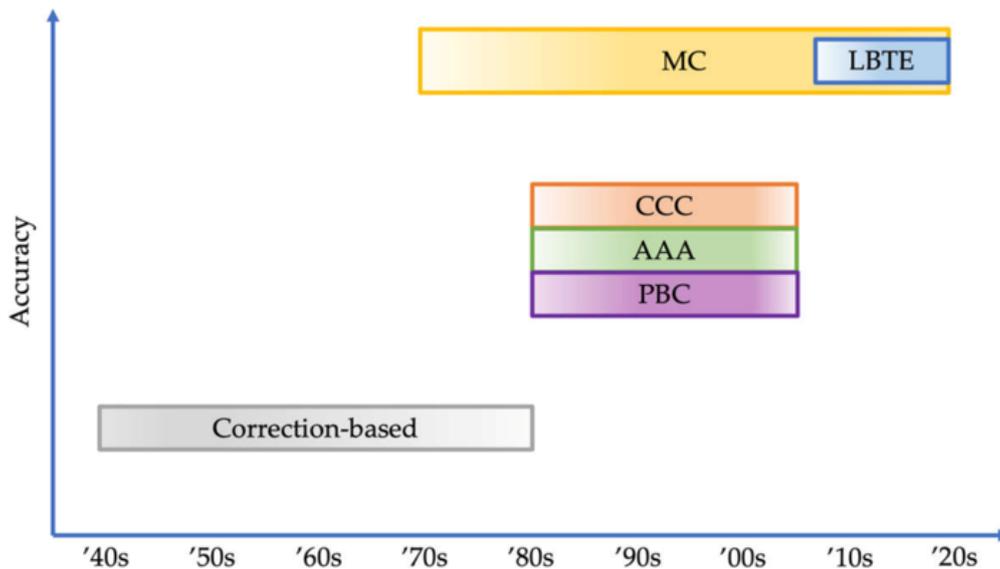


Figure 1.6: **Comparison of classical methods for dose simulation in photon beam radiotherapy (De Martino et al. 2021).** This figure provides a general comparison across time of the classical methods used to simulate radiation dose distributions. The figure highlights the supremacy of the MC method in terms of accuracy.

**Correction-based methods:** Introduced in the 1940’s up to the 1970’s, correction-based methods for dose calculation use empirical data and algorithms to apply corrections to dose calculations to account for various physical effects, such as tissue heterogeneities, scatter, and beam modeling. They interpolate or extrapolate dose values derived from fundamental measurements in water, such as the percentage depth dose, for various field sizes at a specific

source-to-surface distance. This approach incorporates the use of essential concepts like tissue-air ratio and tissue-phantom ratio, which make correction-based methods especially successful in homogeneous media (Cunningham 1972; Khan et al. 1973; Clarkson 1941).

**Pencil Beam (PB):** A "pencil beam" is a term used in the context of radiation therapy and medical physics to describe a highly focused and narrow beam of radiation that is used to deliver a precise dose of radiation to a specific target within a patient's body. More precisely, the PB kernel (Mohan et al. 1986) is computed by integrating the spread kernel along the propagation axis of the incident beam or using MC simulations. The PB has the advantage of increased calculation speed but suffers from uncertainty in the case of heterogeneous medium. Several improvements of the PB kernel have been proposed, notably the anisotropic analytical algorithm (AAA) (Tillikainen et al. 2008).

**Collapse Cone Convolution (CCC):** The CCC algorithm (Anders Ahnesjö 1989) computes the total absorbed dose, i.e. the convolution by considering that energy released by particles propagates and is attenuated in a conical volume. CCC takes into account variations in tissue density, composition, and electron density. It also applies correction factors to the convolution result to account for inaccuracies in regions with tissue density variations, such as in the presence of lung or bone. These correction factors help enhance accuracy in heterogeneous environments.

**Linear Boltzmann Transport Equation (LBTE):** The LBTE (Failla et al. 2010; Vassiliev et al. 2010) is a mathematical equation used in the field of radiative transfer and particle transport. In contrast with MC simulations which are stochastic methods, LBTEs are deterministic equations that provide a mathematical framework for modeling the transport of radiation particles. They describe the behavior of particles in terms of distribution functions and are based on solving partial differential equations that can be solved using numerical methods. In the context of radiation physics and radiation therapy, LBTE-based methods are computational techniques that use the Linear Boltzmann Transport Equation to model and simulate the behavior of radiation particles as they interact with matter, including human tissue. These methods are used to predict the transport and deposition of radiation in various applications, including radiation therapy, radiology, and nuclear engineering.

In summary, several methods exist and are used in treatment planning to simulate the radia-

tion dose deposited in the patient. Many works compared dose calculation methods (Richmond et al. 2021; Vangvichith et al. 2019; Bosse et al. 2020; Hasenbalg et al. 2007; Gray et al. 2009; Han et al. 2011; Çath et al. 2013), both in terms of accuracy and of computational speed. Nevertheless, in all these works and as depicted in Figure 1.6, one method remains unrivaled and the gold standard in terms of precision: the MC method.

## 1.4 Monte-Carlo method

In this section, we first recall the general mathematical framework of the Monte-Carlo method and then develop its use in dose simulation and highlight the main challenges this thesis tries to address.

### 1.4.1 Mathematical Framework

The Monte-Carlo (MC) (Metropolis 1987; Raeside 1976) method is a powerful computational technique widely used for simulating complex systems and solving mathematical problems that may be challenging to solve analytically. Its name is inspired by the famous Monte-Carlo Casino in Monaco, known for its element of chance, which reflects the probabilistic nature of the method. The MC method is particularly valuable in situations where deterministic solutions are either impractical or unavailable due to the complexity of the system being studied. One of the key applications of the MC method is in simulating radiation transport for radiotherapy dose calculations.

The problem is formulated in terms of mathematical equations and rules governing the behavior of the system. In the context of radiation transport and radiotherapy dose calculation, this involves modeling the interactions of particles (e.g., photons or electrons) with matter. Let us consider a general problem expressed as an integral, where  $f(x)$  is a function that we want to integrate over the interval  $[a, b]$  to find the value of  $I$ :

$$I = \int_a^b f(x) dx$$

**Random Sampling:** Random numbers are generated to simulate the behavior of the system. These random numbers guide the trajectory and interactions of particles as they travel through the medium. The random sampling captures the inherent uncertainty and stochastic nature of particle interactions. In the MC method, random numbers are generated to simulate the

behavior of the system. Let's assume that  $X_i$  are independent random variables uniformly distributed over the interval  $[a, b]$ . Then, the MC estimate of the integral  $I$  is given by:

$$I \approx \frac{b-a}{N} \sum_{i=1}^N f(X_i)$$

where  $N$  is the number of random samples.

### 1.4.2 Monte-Carlo for Radiotherapy

In the field of radiotherapy, the MC method plays a pivotal role in accurately calculating the distribution of radiation dose deposited in a patient's body. This is crucial for treatment planning, where precision is paramount to ensure effective tumor control while minimizing damage to healthy tissues. When ionizing radiation is delivered to a patient, it interacts with the biological tissues along its path. These interactions result in energy deposition, which contributes to the absorbed dose and ultimately impacts the biological response. In the context of radiotherapy dose simulation, the MC method is applied to simulate the behavior of ionizing radiation as it interacts with biological tissues.

**Problem formulation:** Computing the total dose deposited during treatment involves modeling the interactions of individual particles (such as photons or electrons) as they traverse the patient's body. To simulate particle interactions, random numbers are generated to determine the particle's trajectory. For instance, the probability distribution for scattering angles can be sampled to determine the deflection of a particle's path due to scattering interactions. Thus, the dose effectively deposited at a certain point can be represented as the integral:

$$D(x, E) = \int_0^\infty S(x, E') \cdot P(E' \rightarrow E) dE'$$

Here,  $D(x, E)$  is the dose deposited at a specific point  $x$  and energy  $E$ ,  $S(x, E')$  is the energy fluence at point  $x$  and energy  $E'$ , and  $P(E' \rightarrow E)$  is the probability that an incident particle with energy  $E'$  deposits energy  $E$  in the medium. However, accurately predicting the radiation dose distribution is a complex task due to the varying densities and compositions of tissues, as well as the intricate patterns of particle interactions. MC simulations offer a comprehensive approach to modeling these interactions. By simulating the trajectories of individual particles (e.g., photons or electrons) and tracking their interactions with matter, the MC method provides a detailed representation of how radiation is deposited in the patient's body. This approach considers

factors such as scattering, absorption, and secondary particle production, which contribute to the overall dose distribution. Through the statistical analysis of a large number of simulated particles, the MC method produces an estimation of the dose distribution. This distribution informs radiation oncologists about the expected radiation dose received by different tissues, aiding in the precise design of treatment plans that optimize the therapeutic outcome while minimizing side effects. By simulating the trajectories of many particles and tracking their interactions, the MC method provides an estimate of the dose distribution. This is obtained by summing the energy deposited by each simulated particle at various points in the patient's body.

In conclusion, the MC method's probabilistic approach to simulating complex systems finds a valuable application in the field of radiotherapy dose calculation. By accurately modeling radiation transport and interactions, MC simulations contribute to the enhancement of treatment planning and the overall quality of care for cancer patients undergoing radiation therapy.

**Uncertainty:** The uncertainty of the deposited dose in a MC simulation is influenced by various factors, including the statistical properties of the simulation, the complexity of the geometry and interactions, and the quality of the simulation setup. The uncertainty is typically expressed as the standard deviation or the standard error of the mean of the calculated dose values from multiple simulated particles. This uncertainty is related to the number of simulated particles  $N$  and the inherent randomness of particle interactions that we described in section 1.2.2.

Let  $D_i$  be the dose value deposited by the  $i$ -th simulated particle in a considered voxel. We define the mean dose  $\mu_D$  deposited by all simulated particles in that voxel:

$$\mu_D = \frac{1}{N} \sum_{i=0}^N D_i$$

Then the standard deviation  $\sigma_D$  of the dose values obtained from a Monte Carlo simulation can be calculated using the formula for the sample standard deviation:

$$\sigma_D = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (D_i - \mu_D)^2} = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (D_i^2 - \mu_D^2)} \quad (1.1)$$

Formula 1.1 considers the differences between each individual dose value  $D_i$  of the  $i^{th}$  particle and the mean dose  $\mu_D$  of all the simulated dose values, and then averages the squared

differences. Please note that  $N$  in the denominator is subtracted by 1 ( $N - 1$ ) due to using the sample standard deviation formula. When we have the entire population of dose values, you would use  $N$  instead of  $N - 1$  for the denominator. This formula helps quantify the spread or dispersion of the dose values around the mean dose, giving an idea of the variability in the simulation results. Keep in mind that the larger the number of simulated particles  $N$ , the more accurate the standard deviation estimation is likely to be. Indeed, as it appears in formula 1.1, increasing the number of simulated particles  $N$  reduces the uncertainty, but at a diminishing rate. Additionally, other sources of uncertainty, such as modeling approximations and data uncertainties, may contribute to the overall uncertainty (Paganetti 2012; Landry et al. 2010) and have a negative impact on subsequent treatment planning (Ma et al. 2005).

**Computational Challenge:** Recent advancements underscore the pressing necessity to seamlessly integrate two essential aspects of treatment planning and delivery: real-time dose simulations for adaptation and the unparalleled precision offered by MC computations. Within this domain, MC methods remain unrivaled, eclipsing conventional algorithms like pencil beam and collapsed cone convolution engines in terms of precision. The MC method is grounded in the probabilistic simulation of the interactions of myriad particles with matter. Therefore, it provides an exceptionally faithful representation of dose distribution. This level of precision is particularly critical when treating tumors located near critical structures or in situations where the patient’s anatomy may change significantly over the course of treatment.

However, the formidable computational demands associated with simulating sufficient particles to yield high-precision dose maps, devoid of inherent MC noise, prevents integration into clinical workflows. Moreover, the exigencies of contemporary treatment techniques amplify the computational complexity of radiation dose calculations. For instance, IMRT restricts the use of only a few discrete gantry angles, whereas VMAT employs continuous gantry movement around the tumor center, optimizing patient care. However, VMAT’s benefits are counterbalanced by the need for increased simulated particles to counteract the inherent noise in MC simulations. As the number of simulated particles escalates, the noise affecting the true dose diminishes. Nevertheless, the full-scope particle transport modeling inherent to MC methods exacts an exorbitant computational toll, constraining their widespread clinical adoption, even though several works have tried to improve the computational requirements by offering simplifications of the method (Udagedara et al. 2015; Nowak 2018; Raffuzzi et al. 2022). Recent research has made strides in mitigating this issue, capitalizing on hardware acceleration and leveraging efficient Graphical Processing Unit (GPU) implementations (Hissoiny et al. 2011;

Neph et al. 2019b; Badal et al. 2009; Martinsen et al. 2009).

Nonetheless, the practical implementation of MC methodologies within clinical settings has been persistently hindered by its computational intensity. This temporal constraint poses a significant challenge in real-world applications where prompt and efficient decision-making is of paramount importance.

## 1.5 Radiotherapy in the era of Artificial Intelligence

### 1.5.1 AI and Radiotherapy

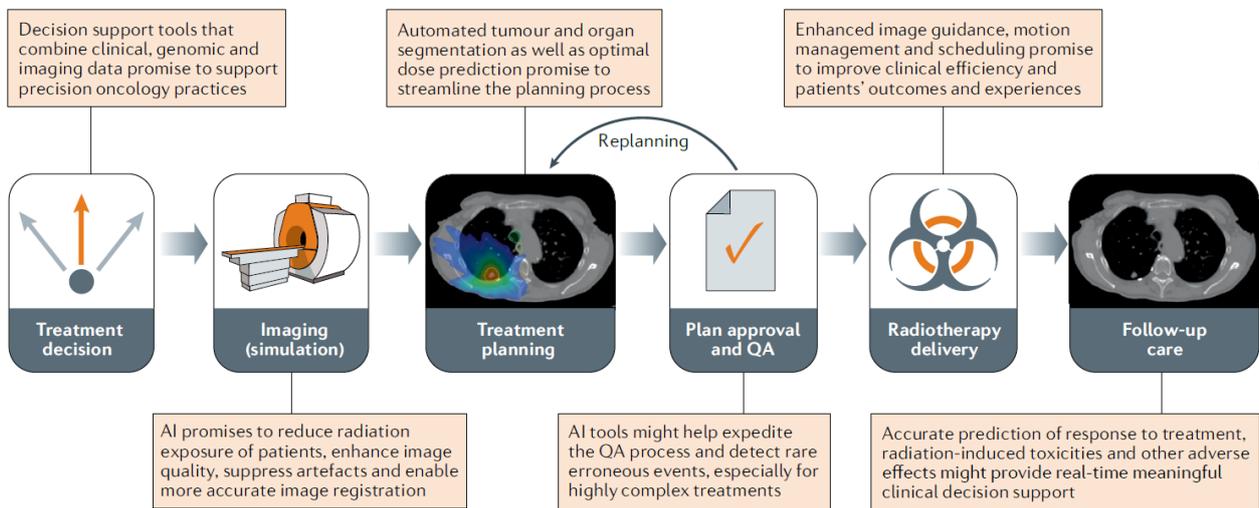


Figure 1.7: **Applications of AI in the radiation therapy workflow (Huynh et al. 2020).** This figure gives a broad overview of the different roles AI can play at each step of the radiotherapy treatment workflow. AI could facilitate activities such as decision-making, acquisition of medical images, delineation of organs, dose computations and quality assessment.

The emergence of artificial intelligence and deep learning as a powerful computational paradigm has engendered shifts across a multitude of domains, providing avenues for innovative strategies to circumvent these limitations. AI has the potential to significantly enhance the radiotherapy treatment workflow in several ways.

**Consultation and Prescription:** The clinical workflow in radiation therapy commences with patient evaluation, where the radiation oncologist conducts a comprehensive review, encompassing the patient's symptoms, medical history, physical examination, pathological and

genomic data, diagnostic studies, prognostic indicators, comorbidities, and radiation-related toxicity risks. Based on the synthesis of these diverse data sources, the radiation oncologist formulates a treatment plan. An emerging challenge in this process is the overwhelming accumulation of data, far beyond what humans can swiftly process. AI-based methods, capable of automatically extracting pertinent clinical insights, are becoming pivotal in creating decision support tools for clinicians at the initial point of patient care. AI techniques, from medical image assessment and extraction meaningful information from biology reports (Savova et al. 2017), have demonstrated early potential in guiding treatment choices and patient management (Jochems et al. 2017; Oberije et al. 2015). AI models have shown promise in enhancing prognosis and forecasting treatment outcomes (Oberije et al. 2015), although they have yet to be integrated into routine clinical practice.

Before the treatment planning stage, the radiation oncologist establishes the prescribed radiation dose for the tumor and defines dose constraints for adjacent organs. These decisions rely on national standards and clinical trial evidence. However, variations in tumor biology can lead to significant differences in radiation sensitivity, even among the same cancer type. Additionally, due to the tumor's geometrical relationship with surrounding organs, achieving the desired dose can be challenging, often realized late in the planning process. AI platforms have the potential to introduce a new level of personalization to radiotherapy by predicting the tumor's radiation sensitivity and determining the optimal achievable dose for a specific treatment plan based on tumor and organ contours (Lou et al. 2019).

**Organ delineation and segmentation:** In radiation therapy, one of the most time-consuming yet crucial task is the manual segmentation of the primary tumor and affected lymph nodes. Accurate tumor delineation is vital as it directly impacts treatment outcomes. Figure 1.8 displays a patient's CT with key organs, Organ-at-Risk (OAR) and Planning Target Volume (PTV) clearly delineated. Errors in delineation can lead to underdosing or overdosing, affecting tumor control and toxicity risks (Cui et al. 2015). Tumor segmentation often exhibits interobserver variation, even among expert radiation oncologists, influencing treatment plan quality and patient survival likelihood (Wuthrick et al. 2015). Current segmentation tools that rely on reference images, such as segmentation atlases. Segmentation atlases are libraries of already segmented medical images that are registered onto each new image with deformable registration (Walker et al. 2014). Yet, atlases can be unreliable, expensive, and still depend on user input (Y. Kim et al. 2016; Delpon et al. 2016; Johnstone et al. 2018).

AI offers a promising solution to enhance the efficiency, quality and standardization of

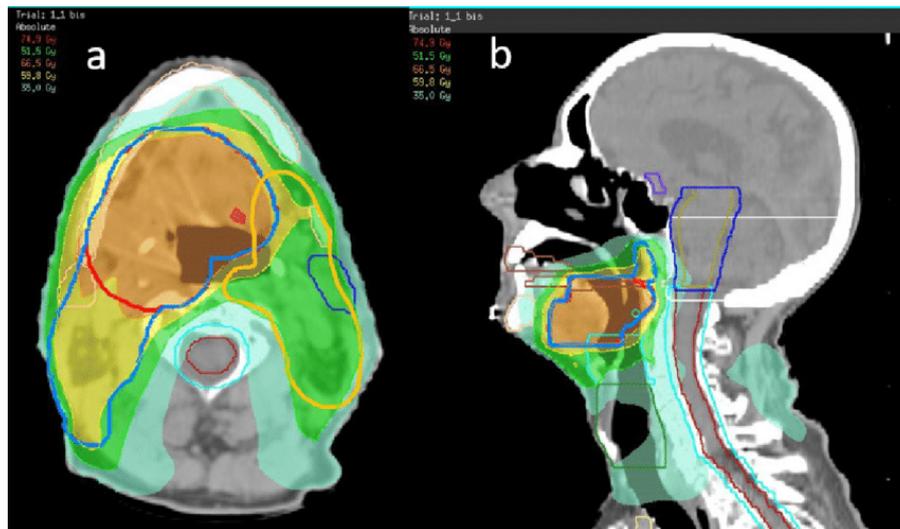


Figure 1.8: CT scan with overlapped delineation of organs involved in radiation therapy (Martini et al. 2020). This figure displays segmentation of key organs prior to treatment planning. The precision of the delineation is crucial for the next steps of the treatment workflow to ensure precise and safe dose delivery.

radiation treatment planning by enabling almost fully automated segmentation, demonstrating high accuracy compared to human experts. However, further prospective studies are needed to assess AI's efficiency, accuracy, and reproducibility within the clinical workflow. Adjacent to tumor segmentation, organs need to be delineated to calculate radiation doses and ensure they stay within safe limits. AI tools have shown potential in segmenting various organs in the body: head and neck (Ibragimov et al. 2017; Nikolov et al. 2021; Zhu et al. 2019), thorax (Lustberg et al. 2018; T. Wang et al. 2020; Ullah et al. 2023; Zhong et al. 2019; Khalil et al. 2022), abdominal organs (Bongratz et al. 2023; Ding et al. 2022; Yu et al. 2022; Aparna et al. 2023) and brain (Pflüger et al. 2022; R. Li et al. 2023). Yet, they are limited by the small number of samples available to integrate in the deep learning model's training set.

**Treatment planning:** Once provided with medical images, tumor and organ segmentations, and the dose prescription, medical dosimetrists aim to create the optimal treatment plan. The process of radiotherapy plan optimization involves finding the most effective solution to meet conflicting objectives: delivering a high radiation dose to the target while minimizing radiation exposure to nearby critical OARs. Achieving this balance requires a series of iterative adjustments of parameters that dictate radiation dose deposition, through an agreed upon set of rules or optimization using statistical methods. This iterative process can be time-consuming, and there's no guarantee that the clinically acceptable plan is the most optimal one. Moreover,

these methods are often limited in accounting for variations in plan complexity and patient-specific trade-offs, and the quality of radiation treatment plans is heavily influenced by human factors, such as decisions regarding the angles of radiation beams and optimization parameters. These decisions can result in substantial variations in the quality of treatment plans, both within and between different medical institutions (Berry et al. 2016).

To achieve high-quality automated treatment plans, these AI-based algorithms must incorporate complex decision-making processes, alike the strategies used in AI applications such as playing games. With the rise of Reinforcement Learning (RL) (Mnih et al. 2015) to perform human-level decisions in complex games such as Atari or Go (Schrittwieser et al. 2020; J. X. Chen 2016), retrospective studies have adopted a paradigm shift towards gamification of the optimization of LINAC’s parameters. In W. Hrinivich et al. 2023 and William Thomas Hrinivich et al. 2020, , the authors adopted RL to automatically define the parameters of the LINAC, such that the positions of the leaves of the MLC or movement of the upper jaws. In the same idea, this game strategy also enabled to build on more classical methods (Boutillier et al. 2015) as seen in Shen et al. 2019, where the authors use reinforcement learning to optimize the weights required to establish an optimal objective function for treatment planning.

**Dose Prediction:** Dose simulation is a crucial step of treatment planning. Several works investigate the feasibility of calculating the dose distribution from the patient’s imaging (CT or MRI scans) (Campbell et al. 2017). Computing the dose distribution from the LINAC’s parameters, such as beam angles, number of gantry arcs, or MLC positions, can be a tedious task. Several classical methods like collapse cone convolution and pencil beam, exist to solve this task and are commercialized in treatment planning systems (TPS). Nevertheless, these methods suffer from uncertainty on the effective dose deposited in the target organs and OARs. Moreover, computing the dose requires intensive calculations and computational power that are mitigated with approximations on the physics underpinning dose deposition.

To overcome these limitations, deep learning offers promising avenues towards precise and fast dose prediction (Xing et al. 2020). As exemplified by Figure 1.9, research focus on harnessing convolutional neural networks to directly predict the dose from information regarding the patient’s anatomy and contours delineating the OARs and PTVs (Kearney et al. 2018; Fan et al. 2019). To further enhance the capabilities of deep learning some works investigate introducing priors related to the LINAC’s beam specificities to add contextual information to the models (Vandewinckele et al. 2022). The use of generative adversarial networks (Goodfellow et al. 2014) (GANs), self-supervised learning and adding constraints in the loss functions, is

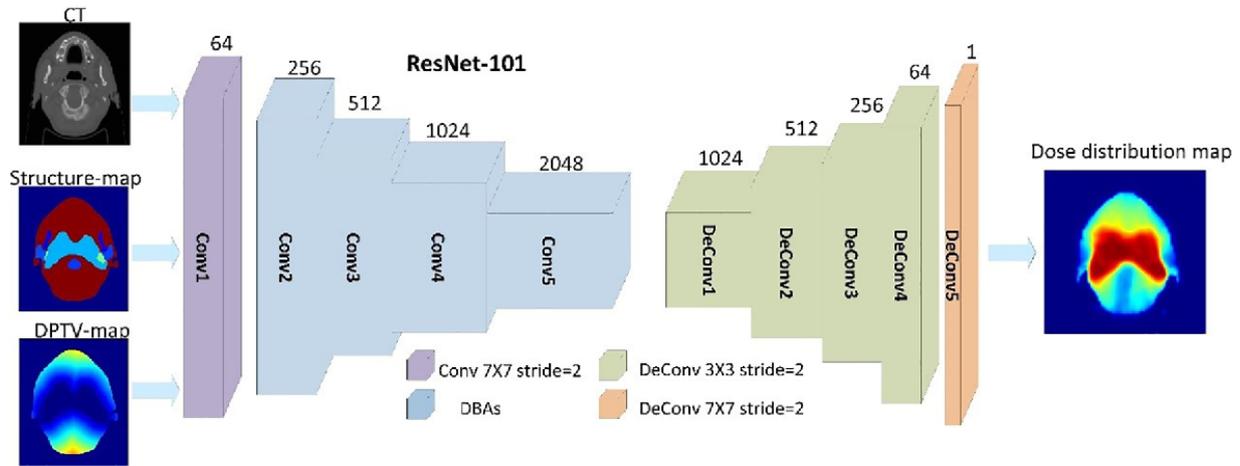


Figure 1.9: **Example of a DL-based dose prediction pipeline (X. Chen et al. 2023).** In this example, the UNet-like model takes as input a patient’s CT, delineation of the involved organs and a map indicating the distance of each organs to the planning target volume. The model is then trained to output a corresponding optimal dose distribution map.

also investigated to better extract and exploit multi-scale features (Zhan et al. 2022).

One of the main challenges faced by deep learning solutions for dose prediction from patient’s scans and segmentation masks resides in the lack of available training data. Discussed research systematically displays a limited number of training samples, rarely above hundred. This lack of data jeopardizes the coveted generalization power (Qilin et al. 2022) of neural networks. To circumvent the data scarcity issue, recent work tries to leverage transfer learning (Kandalan et al. 2020; L. Wen et al. 2023).

### 1.5.2 Deep Learning and Monte-Carlo dose computation

As pointed out in section 1.3.2, the MC method remains the most precise simulation tool for dose computation, but its clinical adoption is hindered by its heavy computational requirements. Our objective centers upon harnessing the potential of deep learning as a strategic accelerator for MC simulations. This entails leveraging the innate computational efficiency of MC simulations conducted with a constrained number of samples. Such simulations, while less accurate, can serve as computationally affordable building blocks that a well-structured deep learning model can assimilate and synthesize to generate dose distributions alike high precision MC computations. The driving force behind this approach is predicated on the symbiotic relationship between the computational efficiency of low precision MC simulations and the remarkable capacity of deep learning models to discern intricate patterns within datasets. By orchestrating

this synergy, our intention is to effectively emulate the outcomes of MC simulations that would otherwise be infeasible within the time constraints of clinical decision-making.

Thus, the current overarching question that needs to be answered is: **How and to what extent can neural networks yield an accelerated approximation of MC simulations?**

Several works have been proposing the use of deep learning models to generate dose maps approximating MC precision from the patient’s CT scans (Götz et al. 2020; Lee et al. 2019; Voss et al. 2023). For example, Pastor-Serrano et al. 2022 train their models on patches of CT scans from 30 patients and leverage the ability of Transformers (Vaswani et al. 2017) to process sequences to achieve milliseconds dose predictions meant approximate MC simulations. Yet, data scarcity still prevents from generalizing to other patient geometries and anatomies, like the presence of metal implants. Additionally, this technique still relies on the type of machine that performed the CT, thus necessitating a substantial amount supplementary data from other machines to ascertain generalization.

Another approach consists in directly accelerating the computation of MC dose simulations by denoising low precision and cheap to compute MC simulations (Bai et al. 2020). Moreover, Deep Learning engines for denoising anatomy specific MC simulations have been proposed in the literature. In Peng et al. 2019 the authors proposed an encoder-decoder architecture to predict high precision simulations from low precision ones in rectal cancer patients treated with IMRT. Neph et al. 2019a used combined UNets (Ronneberger et al. 2015a) coupled with additional CT scans as input to solve the same problem in MR-guided beamlet dose for head and neck patients. Vasudevan et al. 2020 investigated GANs to denoise dose simulations in water phantoms reporting promising results. More recently, Dijk et al. 2022 proposed a deep learning-based denoising approach and alleviated the need for an extensive training set by training the model using synthetic data.

## 1.6 Thesis overview

Our envisioned strategy entails a paradigm shift, where deep learning models are harnessed to emulate and approximate precise MC simulations by exploiting the foundational insights gleaned from less computationally intensive iterations. This integration holds the promise of not only expediting the adoption of MC methodologies within clinical contexts but also elevating the precision and scope of their applications, thereby ushering in a new era of computational efficiency and accuracy in medical decision-making processes. Unlocking the power of fast MC simulation would mean more precise treatments and safer delivery. It will also lead the way

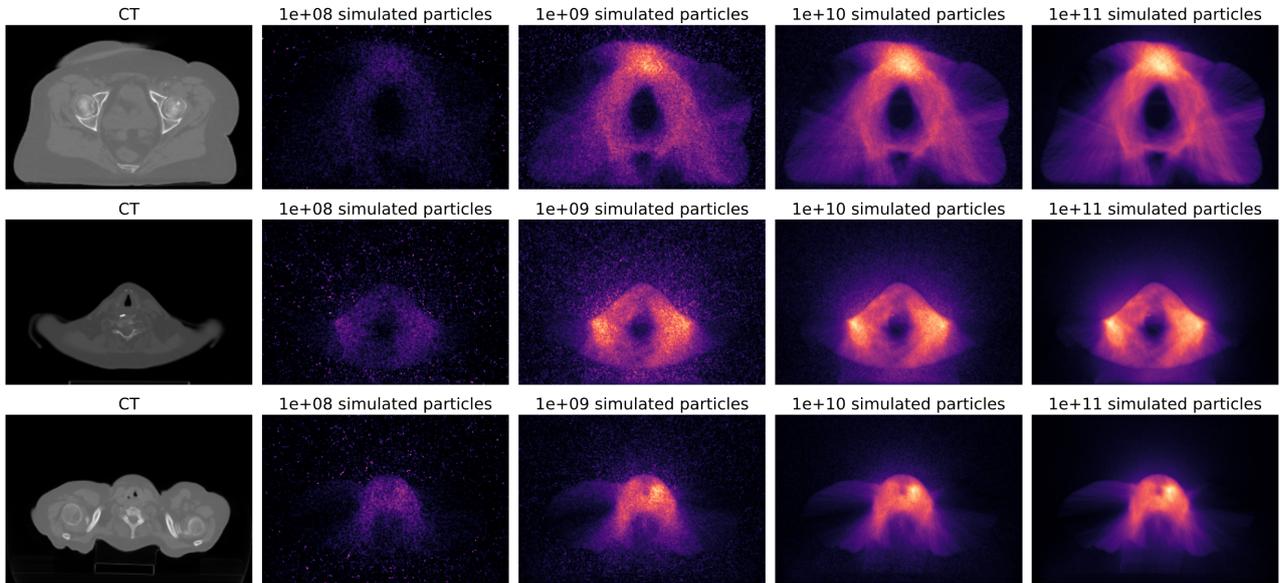


Figure 1.10: **Patients' CT and corresponding normalized dose distributions computed with Monte-Carlo at various numbers of simulated particles.** The figures highlight that the MC method converges and the simulated dose becomes cleaner as the number of simulated particles increases.

towards real-time adaptation methods during delivery of the treatment.

The central objective of this thesis is to develop novel methodologies that seamlessly integrate Deep Learning architectures with MC simulations, thereby enhancing the efficiency and expediency of radiotherapy dose calculations. By capitalizing on the computational strengths of Deep Learning, this research aims to address the challenges posed by the time-consuming nature of MC simulations, making strides towards achieving clinically practical dose calculation times. In the following sections we detail our contributions.

To sum up, we propose to utilize sequences of low precision MC simulations to infer high precision dose distributions with AI. In order to conduct our experiments and optimize the chosen AI models, we generated a unique in the world dataset comprising dose distribution simulated with MC from real-world radiotherapy patients. With this dataset, we investigate several deep learning frameworks to evaluate our hypothesis. Finally, we present a new family of loss functions that integrate a clinical criterium directly into the optimization pipeline to enhance further our training recipe.

### 1.6.1 New perspective on Monte-Carlo simulations

Unlike typical methods explored for speeding up MC simulations using Deep Learning, we took a different approach. Instead of directly using complex mathematical computations, we presented the deep learning model with a series of simulations that became progressively clearer over time. This unique strategy allowed us to provide the model with valuable insights into how the radiation dose accumulates inside the patient’s body. In simple terms, a MC simulation for radiotherapy dose is like trying to estimate how radiation interacts with the patient’s body. It involves simulating the behavior of billions of tiny particles (like photons) as they travel through the body and deposit energy. These simulations generate a lot of data, and we can think of them as a sequence, like frames in a movie, where each frame provides a more detailed picture of the radiation dose.

Our challenge was to teach the model to predict the next frame in this ‘movie’ based on what it had seen so far. This predictive ability is crucial for accurately estimating the final radiation dose. To do this, we needed to consider not just the current frame but also the ones that came before it. This allowed us to capture how radiation accumulates over time and how it spreads out within the patient’s body. We believe that by training the model in this way, it can learn to recognize patterns in how radiation spreads, even when there is some uncertainty due to the limited number of particles used in the simulation. This innovative approach aims to improve the accuracy of radiotherapy treatments by helping us predict radiation doses more effectively. We adopt this perspective to train networks to derive correlation and causation between the levels of noisiness induced by the different numbers of simulated particles. Indeed, from a human eye and as displayed by Figure 1.10, some spatial correlation prevails and might be an indicator of how the energy is deposited in patient tissue depending on the number of generated particles. It is our hypothesis that a model could learn from observing this relationship between uncertainty and spatial distribution of the dose conditioned by the number of simulated particles.

### 1.6.2 Monte-Carlo dataset of dose distributions

To conduct our deep learning experiments to investigate our approach developed in subsection 1.6.1, we generated our own, first and unique in the world dataset. This dataset comprises dose distributions simulated with MC and delivered using the VMAT strategy. To simulate these dose distributions, we had to model the lower part of a LINAC’s head using a software specialized for modelling particle transport and carried out the MC simulations using this model

and the CT scans of 50 patients.

The unique nature of our dataset and still limited resources constricted the simulations to 50 patients. Open-source datasets for our specific application are inexistent. Thus, this first and unique dataset, while quite small, remains our only reference for all subsequent experiments, de facto introducing an inherent limitation to the work presented in this thesis.

### 1.6.3 Designing a learning loss function for Monte-Carlo

Based on our research and experimental observations, we present a novel loss function to train neural networks on the specific task of dose denoising. This loss function approximates the clinical evaluation metric used for dose quality assessment in a differentiable manner. Thus our work integrates a clinical criterion directly into the optimization of the model’s weights.

The primary objective behind these advancements is to circumvent the pitfalls associated with proxy problems induced by training models on ill-suited optimization targets. By tailoring the training process to align closely with the intricacies of radiotherapy dose distribution modeling, we seek to attain a higher level of precision and reliability in our simulations. These developments represent a significant step forward in our ongoing efforts to accelerate MC simulations for the field of radiotherapy and to yield clinically trustworthy and explainable dose predicting models. Therefore, this loss function opens a plethora of possibilities for devising more nuanced and effective training strategies for the task of modeling radiotherapy dose distributions and dose generation.

### 1.6.4 Overview of the contributions

To summarize, the contributions presented in this manuscript are threefold:

1. The first, unique and open-access dataset of MC dose simulations in the case of VMAT patients, generated using a supercomputer.
2. 3D fully convolutional recurrent framework to learn denoising of MC simulations from sequences. We benchmarked our method against several other models. Implementation and trained models are open-source.
3. A new family of loss function based on a gold standard clinical metric used to compare and evaluate dose distributions. This work yielded a pending patent.

Our open source code is available at <https://git.io/JcbxD> and <https://tinyurl.com/ycytz9z2>.

These contributions led to various scientific publications and awards.

## Conference Papers

- "High-particle simulation of Monte-Carlo dose distribution with 3D ConvLSTMs", Sonia Martinot, Norbert Bus, Maria Vakalopoulou, Charlotte Robert, Eric Deutsch and Nikos Paragios, *Medical Image Computing and Computer Assisted Intervention (MICCAI)* - Paper and Oral, 2021.
- "Differentiable Gamma Index-Based Loss Functions: Accelerating Monte-Carlo Radiotherapy Dose Simulation", Sonia Martinot, Nikos Komodakis, Maria Vakalopoulou, Norbert Bus, Charlotte Robert, Eric Deutsch and Nikos Paragios, *Image Processing in Medical Imaging (IPMI)* - **Paper and Best Poster Award**, Bariloche, Argentina, 2023.

## Conference Abstracts

- "Weakly supervised 3D ConvLSTMs for Monte-Carlo radiotherapy dose simulations", Sonia Martinot, Norbert Bus, Maria Vakalopoulou, Charlotte Robert, Eric Deutsch and Nikos Paragios, *Medical Imaging for Deep Learning (MIDL)* - Abstract and Oral, 2021.
- "Fast Monte-Carlo dose simulation with recurrent deep learning", Sonia Martinot, Norbert Bus, Maria Vakalopoulou, Charlotte Robert, Eric Deutsch and Nikos Paragios, *European Society for Therapeutic Radiology and Oncology (ESTRO)* - Abstract and Oral, Madrid, Spain, 2021.
- "Deep Particles Embedding for Monte-Carlo simulations", Sonia Martinot, Nikos Komodakis, Maria Vakalopoulou, Norbert Bus, Charlotte Robert, Eric Deutsch and Nikos Paragios, *European Society for Therapeutic Radiology and Oncology (ESTRO)* - Abstract and Poster, Vienna, Austria, 2023.

## Awards and other activities

- Best Poster Award of IPMI conference, 2023.
- Winner of the challenge *Ma thèse en 6 minutes* of Université Paris Saclay.

- Teaching assistant for Deep Learning in Medical Imaging course of the MVA Masters at ENS.
- External lecturer for the Python course at ESSEC.



# Chapter 2

## Monte-Carlo Dataset

### Contents

---

<b>2.1</b>	<b>Patient cohort</b>	<b>44</b>
<b>2.2</b>	<b>Simulation of dose distributions</b>	<b>45</b>
2.2.1	Geant4 and OpenGate	45
2.2.2	Modelization of the LINAC	46
2.2.3	Organization of the MC simulations	50
2.2.4	Uncertainty on the dose of the simulations	51
<b>2.3</b>	<b>Experimental setup</b>	<b>53</b>
2.3.1	Training data generation pipeline	53
2.3.2	Quantitative Image Quality Assessment	55
2.3.3	Quantitative Clinical Dose Evaluation	56
<b>2.4</b>	<b>Conclusion</b>	<b>59</b>

---

To conduct our research into the potential of deep learning for expediting MC simulations, our initial focus was on creating a comprehensive MC dataset encompassing radiotherapy dose distributions. In this chapter, we outline the methodology employed for dataset generation, justify the tools chosen to achieve our goal, detail the patient data utilized in this process, and provide essential dataset statistics that serve as quality assurance measures for the simulations.

## 2.1 Patient cohort

In the context of the collaboration contract framed by the Framework Agreement Noesia (l'intelligence artificielle au service de la radiothérapie et oncologie), the Gustave Roussy Institute provided the clinical data of 50 radiotherapy patients undergoing VMAT treatment. Each patient in the dataset received treatment following a VMAT plan comprising two gantry rotations around the patient. Each plan was designed for a Varian LINAC TrueBeam (Glide-Hurst et al. 2013) and a photon beam of energy 6 MeV.

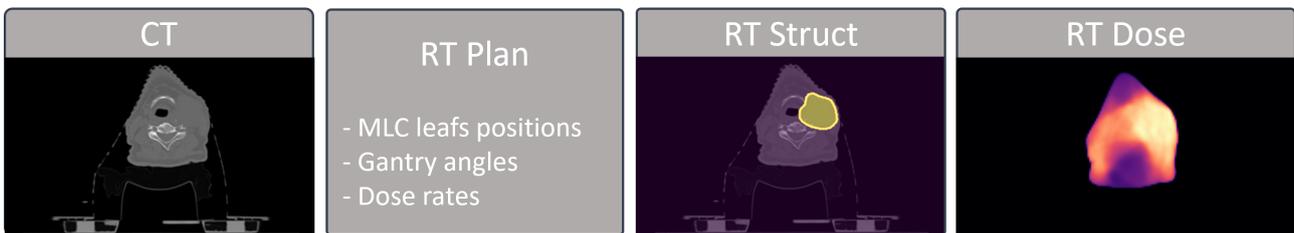


Figure 2.1: **Available data for each patient in the cohort.** The data for each patient comprise the CT scan, the Radiation Therapy Plan (RT Plan) defining all the machine parameters and delivery type, the Radiation Therapy Structure Set (RT Struct) indicating the segmentation masks of the organs involved, and the Radiation Therapy Dose (RT Dose) which is the corresponding delivered dose distribution.

Depicted in Figure 2.1, the dataset comprises the patients' CT scans and files characterizing the radiotherapy plans. Each radiotherapy plan includes several files:

- the CT scan of the patient around the area comprising the target cancerous cells. Each CT has a spatial resolution of  $1mm^3$ .
- the Radiation Therapy Structure Set (RT Struct) which refers to a structure set. It is a type of DICOM object that defines the contours or regions of interest within a patient's anatomy. These contours represent various structures such as the target volumes (tumors), critical structures (OARs) and normal tissues within the patient's body. It specifies the shape, size, and location of these structures, aiding in treatment planning and delivery.
- the Radiation Therapy Dose (RT Dose). This file describes the radiation dose distribution within the patient's body. It provides a 3D map of the delivered radiation dose, allowing for visualization of how radiation is distributed within the target and surrounding tissues.
- the Radiation Therapy Plan (RT Plan), which outlines the details of the treatment plan, including beam configurations, treatment techniques, prescribed dose, and dose

constraints. This file provides information about the individual treatment beams, their parameters, energy levels, angles of delivery, placements of the LINAC's source, jaws and MLC leaves. It defines how the treatment will be administered and serves as a blueprint for the LINAC.

The patients in the dataset underwent radiotherapy to target cancerous cells present either in the brain, Oto-Rhino-Laryngology (ORL) region, pelvis or thorax regions.

## 2.2 Simulation of dose distributions

Deep learning methods are data hungry, requiring usually an extensive amount of data to train the models. One of our major contributions is a dataset comprising the MC simulations of the dose distributions corresponding to the real-world radiotherapy patients introduced in previous section 2.1. To that end, we modeled the head of a LINAC to simulate the whole VMAT treatment plan and the interactions of the simulated radiative particles within the patients' body.

### 2.2.1 Geant4 and OpenGate

**Geant4:** Geant4 (Agostinelli et al. 2003), or Geometry and Tracking, version 4, is a versatile particle tracking toolkit used in high-energy physics, medical physics, astrophysics, space science, and radiation protection. It allows precise modeling of complex geometries, simulates particle trajectories, offers a wide range of physics models, supports scoring and visualization, and provides user-friendly C++ and Python interfaces. Developed and maintained by the CERN (European Organization for Nuclear Research), Geant4 has an active user community that offers support, documentation, and continuous development. It is used for simulating particle behavior in various applications (Archambault et al. 2003):

- High-energy physics experiments: Geant4 is used in the design and optimization of particle detectors for experiments at accelerators like the Large Hadron Collider.
- In Medical Physics: Geant4 plays a crucial role in simulating the behavior of radiation in human tissues, aiding in treatment planning for radiation therapy and studying the effects of radiation on biological tissues.
- In Space Science: Geant4 helps understand how cosmic rays interact with spacecraft and astronauts in space missions.

- In Radiological and Nuclear Physics: It is used for simulating radiation interactions and evaluating the impact of radiation exposure.

Several works already used Geant4 to model and study LINACs (Sardari et al. 2010; Didi et al. 2022) and simulate the radiation dose of radiotherapy plans (Gasteuil et al. 2019; Aitkenhead et al. 2020). However these works are not all open-source and do not match our specific requirements of simulating VMAT plans and computational resources.

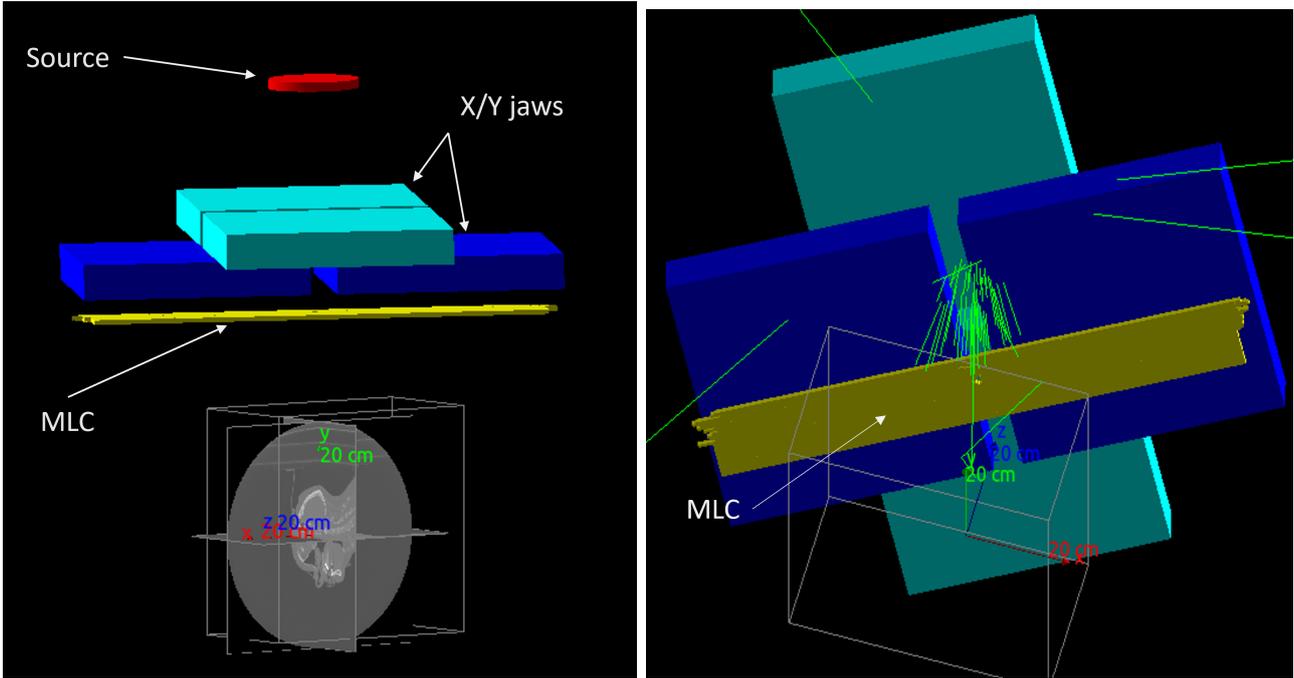
**OpenGATE:** In order to systematically explore the influence of the number of simulated particles, which directly relates to the uncertainty within the MC simulations, we opted to utilize the OpenGATE software. This choice afforded us the flexibility to terminate the simulations at distinct particle counts, enabling a meticulous examination of the simulation’s progressive refinement. OpenGATE (Sarrut et al. 2022) is a project that builds on Geant4, extending its capabilities and providing a more user-friendly framework for specific applications in medical physics and radiation therapy.

OpenGATE is a versatile and open-source software platform primarily used for simulating the transport of particles in matter. It’s particularly well-known and widely used in the field of medical physics and nuclear medicine (Grevillot et al. 2011; Teixeira et al. 2019; Krim et al. 2020). It incorporates advanced models for simulating the interactions of radiation with biological tissues. This is crucial for applications in radiation therapy and nuclear medicine. Users can integrate medical imaging data (e.g., CT, MRI) into simulations, making it valuable for treatment planning and research in medical physics. OpenGATE allows researchers and professionals to simulate the behavior of various types of radiation, including gamma rays, X-rays, and positrons, as they interact with different materials such as tissues or detectors.

### 2.2.2 Modelization of the LINAC

To match the specificities of the patient cohort described in section 2.1, we modeled a Varian Truebeam LINAC with 6 MeV photon beam. Modeling a LINAC can require to model the source of the radiation and all the beam-shaping objects (jaws, MLC, flattening filter). In our case, we used OpenGATE to model only the lower part of the head of a LINAC, i.e. we restricted the modelization to the upper jaws and the MLC. Figure 2.2a displays a transversal view of our modelization.

**Beam shapers:** We modeled the jaws and MLC with rectangular volumes without simulating wedge parts. We set the dimensions of the upper jaws, also called the X, Y jaws, and MLC to



(a) Transversal view of the modeled LINAC and perspective with a patient's CT. (b) View from below the modeled LINAC head. Green lines are the trajectories of the simulated photons.

Figure 2.2: **Visualizations of the OpenGATE modelization of the lower part of a LINAC's head.** The MLC is depicted in yellow, the upper jaws in blue and cyan, and the location of the source is indicated by the red disk. The CT of the patient is introduced below the LINAC's head, awaiting radiation. Green lines indicate the trajectories of simulated ionizing particles.

match the schematics of a 6 MeV photon beam of a Varian Truebeam LINAC. We extracted the gantry angles and corresponding dose rates as well as positions of the source, X/Y jaws and MLC from the original clinical RT plans and incorporated them in OpenGATE to simulate the VMAT plans. For clarity, we present a visual representation of this modelization in Figure 2.2a.

We introduced a simplification in the modeling process to reduce computational time. Specifically, we employed a technique referred to as "KillActors" in the treatment of the MLC jaws and leaves. KillActors effectively eliminate any incident particle upon collision with the MLC components. Each MLC leaf/jaw was modeled by a rectangular volume respecting dimensions from schematics of the LINAC. In Figure 2.2b, green lines represent the trajectories of simulated photons generated from the phase plane. The phase plane is located just above the X-Y jaws. Most photons are stopped by the X-Y jaws and MLC KillActors. Consequently, the modeling framework does not account for the penumbra effect. The rationale underlying this choice pertains to the significant computational overhead associated with modeling the intricate interactions between particles and the individual metal leaves and jaws, which would

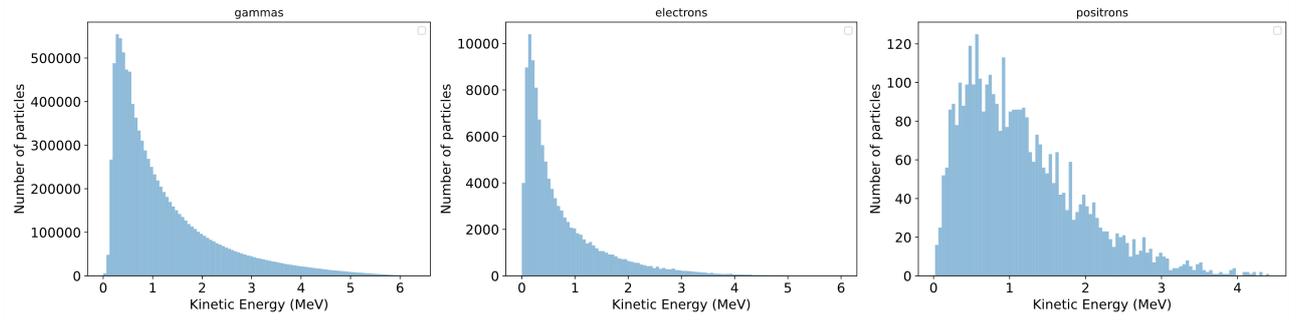
impose a substantial increase in computational time.

**Source and Phase space:** We represented the location of the radiation source as a cylindrical volume, visible in red in Figure 2.2a. The source was located 1 *m* from the patient’s isocenter. An isocenter, short for "isocenter of rotation", is a crucial point in radiation therapy and medical imaging, particularly in treatments that involve radiation or focused beams of energy. It is a specific location within the patient’s body where multiple radiation beams or imaging lines converge or intersect. More precisely, LINACs or diagnostic imaging equipment like CT scanners, are designed to rotate around the patient. The isocenter is the point around which the machine or imaging system revolves.

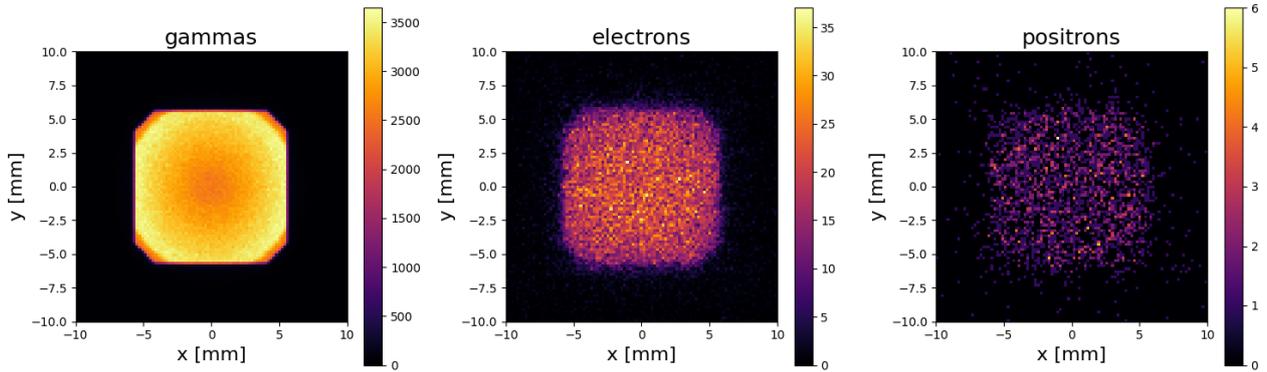
Rather than simulating the entire trajectory from source to the target isocenter, we leveraged a publicly accessible phase space dataset made available by the International Atomic Energy Agency (IAEA). The phase space of a LINAC is a multidimensional space that characterizes the distribution of particles such as electrons or photons, both in terms of position and momentum. It plays a crucial role in understanding the behavior of the particles beam generated by the LINAC. The phase space of a LINAC includes the following key parameters:

- **Position in Space:** This component of the phase space describes the spatial coordinates ( $x, y, z$ ) of the particles or photons within the beam. It defines where each particle or photon is located within the beam at a specific point in time.
- **Momentum:** Momentum refers to the velocity of the particles or photons multiplied by their mass. In LINAC physics, momentum is a crucial parameter, and it characterizes the kinetic energy and direction of each particle or photon in the beam.
- **Energy:** Energy is closely related to momentum and describes the total kinetic energy of each particle or photon in the beam. It is particularly important in radiation therapy, where the LINAC beam must deliver a precise dose of energy to the treatment target.
- **Direction:** The direction component of the phase space specifies the trajectory or path that each particle or photon follows as it moves through the LINAC. This direction can be defined in terms of angles (e.g., polar and azimuthal angles) or unit vectors.
- **Phase Space Volume:** The entire phase space volume encompasses all possible combinations of position and momentum for the particles or photons within the LINAC beam.

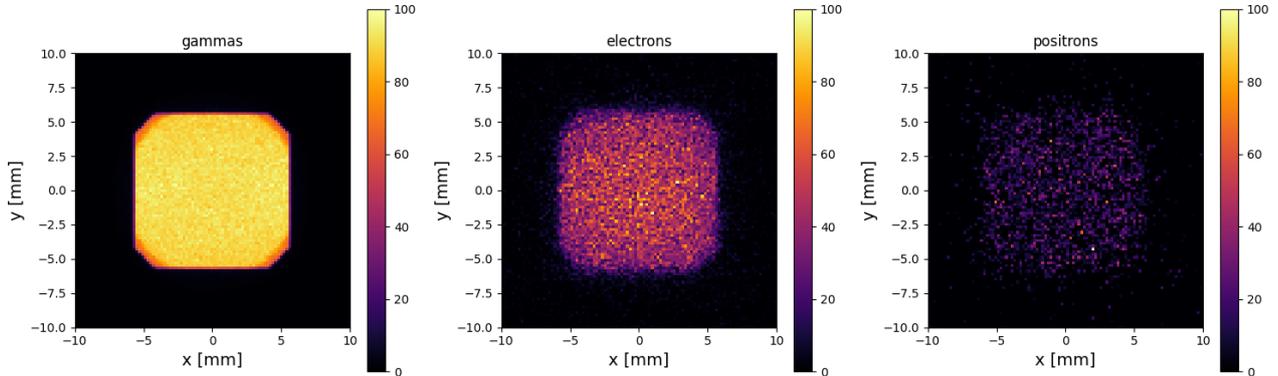
Our chosen phase space (Constantin et al. 2010), comprises the information of 52477334 particles, including 51865724 photons, 591763 electrons and 19847 positrons. Each of these



(a) Energy distribution of particles generated by the phase space.



(b) 2D histograms showing the spatial distribution of the amount of each type of particles.



(c) Normalized 2D histograms showing the spatial distribution of the energy of the particles.

Figure 2.3: **Spatial and energy distributions of the particles comprised in the phase space.** Several ionizing particles are included in the chosen phase space: electrons, photons (gamma rays) and positrons. The phase space comprises a majority of photons of high energy.

particles has associated momentum, initial position and energy. Figure 2.3a displays the distributions of each type of particles accounted for in the phase space, depending on their kinetic energy. Gamma rays, i.e. photons, have much higher energy than the electron and positrons. Figure 2.3b shows how the particles are distributed spatially, while Figure 2.3c displays the

normalized distribution of energy depending on the considered type of particles. During MC simulations, these particles are randomly selected to simulate their trajectories. In MC simulations requiring the simulation of a greater number of particles, particles are reused for simulation with a different random starting point to allow for a different stochastic behavior of the particles and therefore of their interactions with various mediums.

Once the modelization created, we place the CT scan of the patient such that the distance between the isocenter and the source equals 1 *m*. The initial setting of the simulations can look like Figure 2.2a. The CT provides the electron density information that characterizes the tissues of the organs within the patient’s body. Also, the resolution of the provided CT defines the resolution of the output dose distribution. During the simulation, a user-defined number of particles, mixing electrons and photons, is sampled from the provided phase space. The behavior and interactions of these particles with the patient’s tissues are computed using the MC method to yield the deposited energy.

### 2.2.3 Organization of the MC simulations

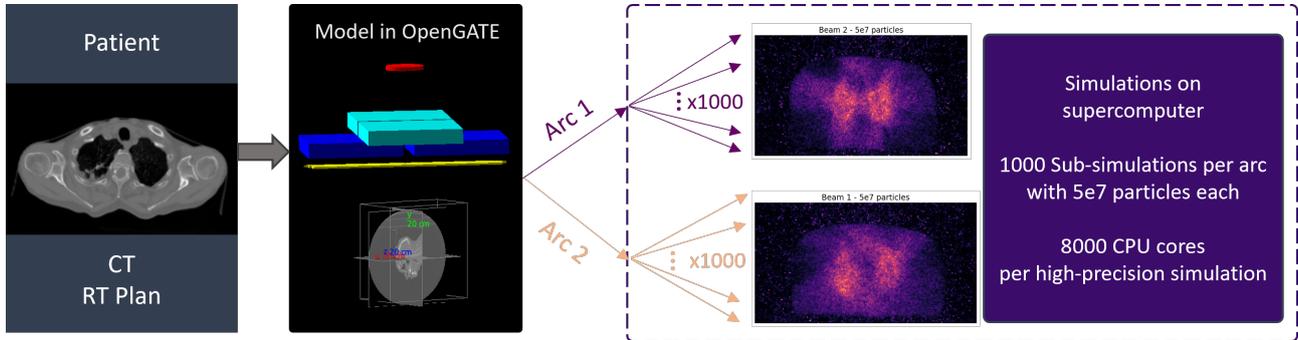


Figure 2.4: **Supercomputer computation strategy of MC simulations.** Once the treatment plan modelled in OpenGATE, the high-precision simulation is split into many sub-simulations that require fewer particles and therefore a reduced computation time. These sub-simulations are heavily parallelized on the supercomputer. Each complete simulation necessitated 8000 CPUs.

We conducted experimental tests to determine the required amount of simulated particles to achieve acceptable dose simulations. We decided to simulate the behavior of  $10^{11}$  particles for each patient in the dataset. For this number of particles, the computed maximum uncertainty defined by equation 1.1 in section 1.4.2, remains under 3% in regions where dose values are within 20% of the maximum dose.

To simulate the behavior of  $10^{11}$  particles for each patient case in a reasonable amount of

time, we parallelized each simulation of  $10^{11}$  particles into a thousand sub-simulations with random initializations. One sub-simulation required computing the behaviour of  $10^8$  particles. Each plan comprises 2 gantry rotations, i.e. 2 arcs. To further parallelize computations, we separated the simulations of each arc. Thus, the dose simulation involving  $10^{11}$  particles was parallelized into 2000 sub-simulations - 1000 sub-simulations per arc, each requiring the simulation of the behavior of  $5 \times 10^7$  particles. To further reduce the computational overhead and as depicted in Figure 2.4, we halved the resolution of the provided CT to  $2 \text{ mm}^3$ . Finally, thanks to the unbiased property of the MC estimator, complete sub-simulations (2 arcs) were obtained by summing one sub-simulation of each arc, yielding a complete sub-simulation involving  $10^8$  particles. Subsequently, the high precision simulation with  $10^{11}$  particles was simply obtained by summing the thousand corresponding sub-simulations with  $10^8$  particles.

To compute the thousands of sub-simulations for our dataset, we used a supercomputer on which we had one million hours of dedicated computation time. The Très Gros Centre de Calcul (TGCC) is one of the biggest computers in the world, with a computational power of petaflops per second and around 30 petabytes available for storage. On this supercomputer, we carried out the simulations by using heavy multiprocessing and thousands of Central Processing Unit (CPU). More precisely, each per-arc sub-simulation required around 2 hours of computation on 4 cores. The duration of the simulation could vary up to 3 hours depending on the complexity of the patient’s anatomy. The storage space required for the output volumes of a single patient’s simulation amounted to 1TB. Therefore, the high precision simulation for one patient treated with a VMAT plan comprising 2 arcs necessitated a total accumulated computation time of approximately 16k hours - almost 2 years - spanned across 8000 CPU cores. Generating the dataset astoundingly showcased the overwhelming temporal limitations of the MC simulation engine.

## 2.2.4 Uncertainty on the dose of the simulations

Figure 2.5 shows the MC simulation of a selected patient from the cohort described in section 2.1. The uncertainty maps highlight the relationship revealed by formula 1.1 in section 1.4.2. As the number of simulated particles increases, the MC uncertainty decreases. Uncertainty dose maps show that there are some discrepancies regarding the areas that are less targeted by the radiation beam. Indeed, when comparing the uncertainty maps with the dose distributions, we note that the noise level remains higher in areas that have not been irradiated. Nevertheless, the outcomes of the MC simulations that modelled the behaviour of  $10^{11}$  particles yield uncertainty values that remain under 10%, as further confirmed by the violin plots of Figure 2.6. As we did

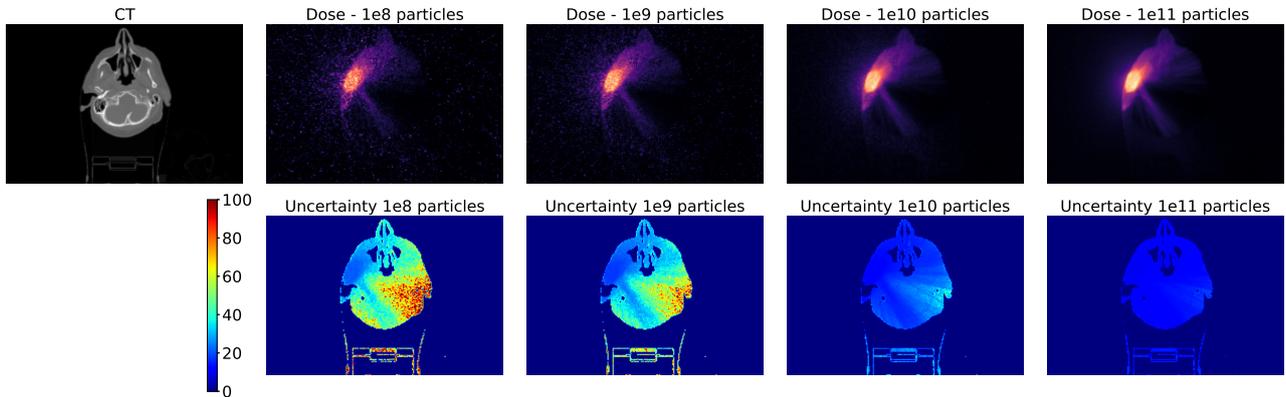


Figure 2.5: **Example of a patient’s CT, corresponding MC simulations at various noise levels and associated uncertainty maps.** This figure highlights that, as the number of simulated particles increases during MC computations, the uncertainty surrounding the deposited dose decreases at a diminishing rate. The uncertainty maps corresponds to the dose maps in the sense that areas at which many particles are targeted display a reduced uncertainty when compared to other areas.

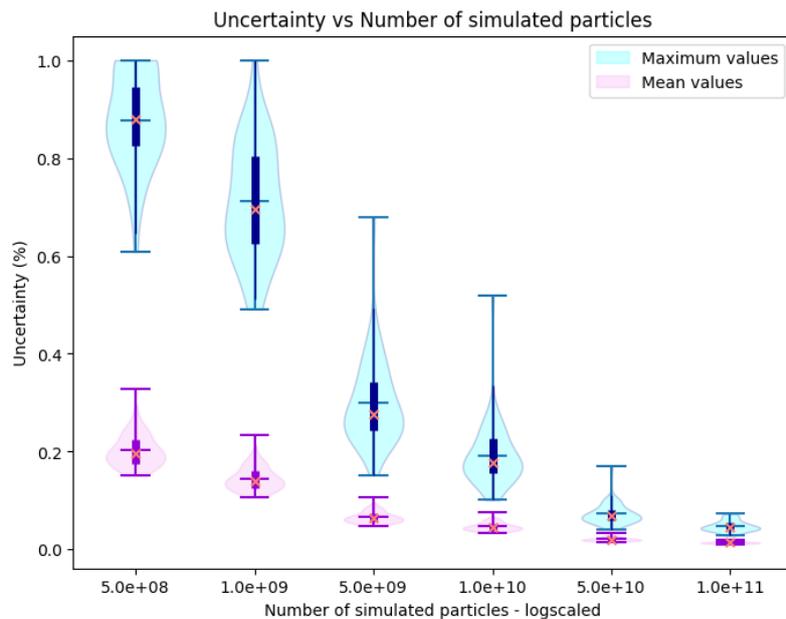


Figure 2.6: **Uncertainty values accross dataset vs number of simulated particles.** The figure shows violin plots of the mean and maximum uncertainty values across the complete dataset, depending on the number of simulated particles. Values stem from voxels in areas with dose values superior to 20% of the maximum dose in the dose distribution. The figure confirms that we reach an acceptable uncertainty level in the high-precision simulations computed with  $10^{11}$  particles.

not have enough computational resources to improve these statistics and since the qualitative evaluation of the obtained dose distributions was satisfactory, we stopped the simulations after

$10^{11}$  particles.

Figure 2.6 displays violin plots of mean and maximum uncertainty values of the MC simulations relative to the number of simulated particles and computed over the entire patient cohort. Uncertainty values were computed in areas within 20% - 100% of the dose maximum. The plots reveal that while all plans were simulated with the same number of particles, there is some variability in the uncertainty values. In other words, the uncertainty on the energy deposited in the patient varies because discrepancies in anatomy complexity.

## 2.3 Experimental setup

In this section, we first develop our strategy to generate training samples quickly from the dataset presented in section 2.2. Subsequently, we present several similarity measures that we used in all experiments to evaluate the quality of the dose distributions generated by our deep learning models, compared to the reference dose distributions.

### 2.3.1 Training data generation pipeline

As detailed in section 2.2, we computed 1000 sub-simulations with  $10^8$  simulated particles each, for each patient comprised in the cohort. Since the MC method provides unbiased estimators, summing all the sub-simulations yields a simulation with a precision corresponding to simulating  $10^{11}$  particles, and we consider the latter simulations as our ground-truth reference data in all presented experiments.

However, as part of our experiments with sequential MC simulations, we need to generate intermediate simulations with fewer particles (for example  $10^9$  or  $10^{10}$  simulated particles) to create the input sequences that are to be processed by the deep learning models during training. To create such simulations, we need to sum the corresponding number of sub-simulations. However, creating and storing these intermediate dose simulations would require an extensive amount of storage space. To circumvent this potential bottleneck, we create these intermediate dose simulations on-the-fly i.e., directly during training of deep learning models.

During training, we create batches of simulations of varying noise levels on-the-fly using the sub-simulations as shown in Figure 2.7. We sum subsets of the sub-simulations in order to generate noisy simulations corresponding to previously decided numbers of simulated particles. As we only have 50 patients represented in the dataset, we add some variability at least in the MC noise by aggregating sub-simulations randomly to the desired number of simulated

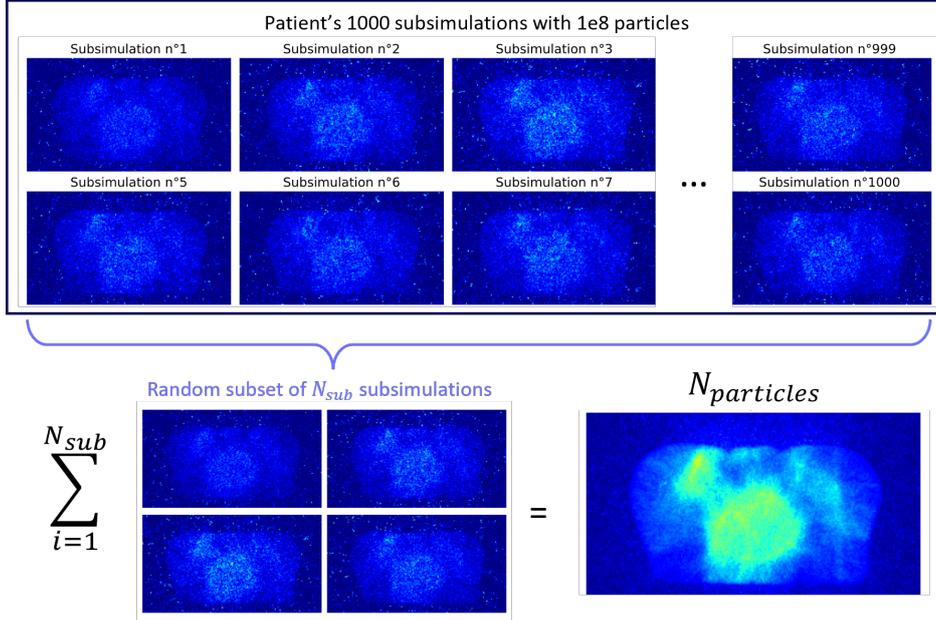


Figure 2.7: **Creation of intermediate noisy simulations from the computed sub-simulations.** 1000 sub-simulations per patient were created on the supercomputer. As they are independent, simply adding them yields intermediate simulations with various numbers of simulated particles, i.e. noise levels, that depend on the number of summed sub-simulations.

particles.

Moreover, let  $N_{particles}$  be the decided number of particles corresponding to the desired level of noise in a simulation. Then, under the satisfied condition that  $N_{particles} \geq 10^8$ , the required number of sub-simulations to sum in order to create the desired simulations is found by:

$$N_{sub} = \frac{N_{particles}}{10^8}$$

As there are 1000 sub-simulations available, all independently and identically distributed, then the number of available dose volumes  $N_{comb}$  per patient with that number of simulated particles corresponds to the combinations of the sub-simulations defined as:

$$N_{comb} = \binom{1000}{N_{sub}}$$

It is important to note that as  $N_{sub}$  increases, the resulting combinations share more information in terms of dose deposition and noise, and are decreasingly independent from each other. We also note  $N_{ind}$  the number of independent dose volumes created to match  $N_{particles}$

simulated particles defined by:

$$N_{ind} = \frac{1000}{N_{sub}}$$

We define  $N_{ind}^{tot}$  the total number of independent dose volumes across the whole MC dataset. As the dataset comprises 50 patients:

$$N_{ind}^{tot} = 50 \times N_{ind}$$

$N_{particles}$	$N_{sub}$	$N_{ind}$	$N_{comb}$	$N_{ind}^{tot}$
5e8	5	200	> 8e12	10000
1e9	10	100	> 2e23	5000
5e9	50	50	> 9e84	2500
1e10	100	10	> 6e139	500
1e11	1000	1	1	50

Table 2.1: Number of available dose volumes per patient in the dataset, depending on the number of simulated particles.

Table 2.1 displays in order of magnitude the number of combinations of sub-simulations that are possible to create less noisy simulations from them. As the input of the model will be a sequence of decreasingly noisy simulations, we ensure that the least noisy ones comprise the same sub-simulations as the noisiest ones belonging to the same input sequence. This ensures that the simulations correlate at least in terms of deposited energy.

For each of the presented experiments, the processing pipeline of the dataset varies prior to training. Specifically, variations encompass patch-based or whole image trainings, the exploration of various normalization techniques, and the integration of data augmentation strategies. These variations are further presented and discussed in the next chapters.

### 2.3.2 Quantitative Image Quality Assessment

In order to assess the fidelity of the generated dose images, we employed well-established metrics that have been widely adopted in image quality evaluation. These metrics include the Structural Similarity Index Measure, the Mean Squared Error, and the Mean Absolute Error. These metrics collectively serve as robust quantitative measures to gauge the likeness and precision of the generated dose images in comparison to the ground truth, facilitating a comprehensive evaluation of the model’s performance. In the following presentations of said metrics, we use the given notations:

- $x$  and  $y$  are the compared images.
- $N$  and  $M$  represent the dimensions of the images.

**Structural Similarity Index Measure (SSIM):** The SSIM (Zhou Wang et al. 2004) quantifies the structural similarity between the generated dose image and the ground truth. It takes into account luminance, contrast, and structure, offering a comprehensive assessment of image similarity. With the conventions that  $\mu_x$  and  $\mu_y$  represent the means of  $x$  and  $y$ ,  $\sigma_x^2$  and  $\sigma_y^2$  are the variances of  $x$  and  $y$ ,  $\sigma_{xy}$  is the covariance between  $x$  and  $y$ ,  $c_1$  and  $c_2$  are constants to stabilize the division, the SSIM is defined as:

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2)}{(\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_2)} \quad (2.1)$$

**Mean Squared Error (MSE):** The MSE (Chai et al. 2014) calculates the average squared differences between pixel values in the generated image and the ground truth. It provides a measure of the overall discrepancy between the two images. Mathematically, the MSE is defined as:

$$\text{MSE}(x, y) = \frac{\sum(x - y)^2}{N \cdot M} \quad (2.2)$$

**Mean Absolute Error (MAE):** The MAE (Chai et al. 2014) computes the mean of the absolute pixel-wise differences between the generated and ground truth images. It offers insights into the average magnitude of the errors between the two images. Mathematically, the MAE is defined as:

$$\text{MAE}(x, y) = \frac{\sum|x - y|}{N \cdot M} \quad (2.3)$$

### 2.3.3 Quantitative Clinical Dose Evaluation

The presented metrics are useful to compare two dose distributions from a clinical standpoint and widely used in practice.

**Gamma index Passing Rate (GPR):** The GPR (Low et al. 1998) is a widely used metric in the field of medical dosimetry and radiation therapy to assess the quality of dose distributions in radiotherapy treatment planning. It quantifies the agreement between a reference

dose distribution (typically the desired or planned dose) and the delivered or calculated dose distribution, taking into account both dose differences and distance-to-agreement criteria. Let us define the relevant terms and parameters:

$D_{\text{ref}}(x, y, z)$ : Dose at a specific point  $(x, y, z)$  in the reference or planned dose distribution.

$D_{\text{calc}}(x, y, z)$ : Dose at the same point  $(x, y, z)$  in the calculated or delivered dose distribution.

$D_{\text{max}}$ : Predefined dose threshold, often expressed as a percentage of the maximum dose in the reference distribution (e.g., 3% of the maximum dose).

$\Delta d$ : The distance-to-agreement (DTA) measures how close the two points  $(x, y, z)$  are in space between the reference and calculated distribution.

$\delta$ : A predefined DTA criterion, often expressed in millimeter.

The GPR is computed based on the values of the **gamma index matrix**  $\gamma$ . The gamma index matrix has the same dimensions as the evaluated dose distribution. Let  $\Delta D = D_{\text{ref}} - D_{\text{calc}}$  be the dose difference between the reference and calculated doses at a specific point, then the gamma index value  $\gamma_{(x,y,z)}$  for a specific point  $(x, y, z)$  is calculated as follows:

$$\gamma_{(x,y,z)} = \sqrt{\left(\frac{\Delta D}{D_{\text{max}}}\right)^2 + \left(\frac{\Delta d}{\delta}\right)^2}$$

The GPR is calculated by assessing how many points in the dose distribution meet predefined criteria. Typically, a passing point is defined as:

$$\gamma_{(x,y,z)} \leq 1$$

In other words, a point  $(x, y, z)$  is considered as satisfying the gamma index test if  $\gamma_{(x,y,z)}$  is less than or equal to 1. The GPR is then expressed as a percentage of the total number of points evaluated:

$$\text{GPR (\%)} = \frac{\text{Number of Passing Points}}{\text{Total Number of Points}} \times 100\%$$

A higher GPR indicates better agreement between the reference and calculated dose distributions and is often used as a measure of plan quality in radiotherapy. Typically, a GPR of 95% or higher is considered acceptable for clinical treatment plans, but the specific criteria can vary depending on clinical protocols and guidelines.

**Dose Volume Histograms (DVH):** A DVH (Drzymala et al. 1991) is a graphical representation and quantitative analysis tool employed in radiotherapy to characterize and visualize the distribution of radiation doses within a patient's anatomical structure. It provides valuable insights into how radiation is delivered to specific regions of interest, including tumors and nearby healthy organs. Key components of a DVH include:

- **X-Axis (Dose):** The horizontal axis of a DVH represents the range of radiation doses delivered to the tissues or structures within the patient. It is typically measured in units of Gray (Gy), which quantifies the amount of radiation energy absorbed by the tissue.
- **Y-Axis (Volume):** The vertical axis represents the volume of the tissue or structure receiving a particular radiation dose. It is often expressed as a percentage of the total volume. For instance, 50% on the Y-axis indicates that half of the volume of the structure received the corresponding dose on the X-axis.

A DVH is constructed based on data from the radiotherapy treatment plan, which includes information about the prescribed dose, the distribution of radiation beams, and the geometry of the patient's anatomy. The DVH curve represents the cumulative percentage volume of the anatomical structure receiving equal to or less than a specific dose. A steeper slope on the DVH curve indicates that a larger volume of the structure receives higher doses of radiation, which can be indicative of effective tumor targeting. Conversely, a shallower slope implies better sparing of healthy surrounding tissues. DVHs serve several critical purposes in radiotherapy:

- **Treatment Planning:** DVHs help radiation oncologists and medical physicists design treatment plans that optimize the balance between delivering a sufficient radiation dose to the tumor while minimizing exposure to adjacent healthy organs.
- **Plan Evaluation:** DVHs aid in the evaluation of treatment plans by quantifying how well the intended dose distribution aligns with clinical goals and constraints.
- **Patient-Specific Assessment:** DVHs allow for patient-specific assessment of radiation dose distributions, enabling adjustments to treatment plans as needed.

In summary, a Dose-Volume Histogram (DVH) is a fundamental tool in radiotherapy that offers a comprehensive and quantitative assessment of how radiation doses are distributed within a patient's anatomy. It plays a pivotal role in optimizing treatment plans, evaluating their effectiveness, and ensuring the safe and targeted delivery of radiation therapy.

## 2.4 Conclusion

In this chapter, we presented the cohort of VMAT patients we utilized to simulate MC simulations of dose distributions. In the OpenGATE software, we modeled the lower part of the head of a LINAC in order to better reproduce the patient’s original radiation treatment plan. We give a detailed and justified description of the various modelization and simplification choices we made. MC computations were made using a supercomputer. The dataset is publically available at <https://hosting.therapanacea.eu/data/miccai2021/>. Because of the amount of MC simulations generated for this dataset, we described our on-the-fly data generation strategy to yield training data when optimizing deep learning models. Finally, we presented the computer vision metrics and clinical criteria we selected in order to evaluate the quality of the dose distributions predicted by deep learning models in comparison with the reference MC dose distributions computed with  $10^{11}$  simulated particles.

In the Annexe 6.2.3, we provide visualizations of several patients’ MC dose simulations and the corresponding uncertainty maps.



# Chapter 3

## Denoising Monte-Carlo simulations

### Contents

---

<b>3.1</b>	<b>Introduction</b>	<b>62</b>
3.1.1	New perspective on Monte-Carlo	62
3.1.2	Problem statement	64
<b>3.2</b>	<b>ConvLSTMs for 2D denoising</b>	<b>65</b>
3.2.1	Definitions	65
3.2.2	Methods	67
3.2.3	Results	68
3.2.4	Discussion	70
<b>3.3</b>	<b>3D-ConvLSTM</b>	<b>70</b>
3.3.1	Introduction	70
3.3.2	Methods	70
3.3.3	Methods	73
3.3.4	Results	73
3.3.5	Discussion	76
<b>3.4</b>	<b>Ablation studies</b>	<b>76</b>
3.4.1	Alterations in the input data	76
3.4.2	Towards lighter 3D ConvLSTM	77
3.4.3	Evaluation per anatomy area	78

3.4.4	Discussion . . . . .	78
<b>3.5</b>	<b>Weakly-Supervised ConvLSTMs . . . . .</b>	<b>79</b>
3.5.1	Introduction . . . . .	79
3.5.2	Methods . . . . .	80
3.5.3	Results . . . . .	80
3.5.4	Discussion . . . . .	81
<b>3.6</b>	<b>Conclusion . . . . .</b>	<b>82</b>

---

## 3.1 Introduction

### 3.1.1 New perspective on Monte-Carlo

**Modelization:** In contrast with usual approaches studied to accelerate Monte-Carlo simulations using Deep Learning, we choose to give a sequence of increasingly clean simulations as input to the model. This has the advantage of giving additional information to the model about how the dose increases within the patient depending on the number of particles generated.

A MC simulation of radiotherapy dose requires inferring the dose deposited by billions of photons in the human body. The method consists in drawing independent random samples from an unknown distribution by means of sequentially sampling empirical measures that describe the dose deposition of individual photons. Let us denote by  $M_{N_i} \in \mathbb{R}_+^{d_1 \times d_2 \times d_3}$  the 3D dose volume result of a simulation performed with  $N_i$  photons. Since several MC dose simulations for the same patient are independent from each other, the following equation holds:

$$M_{N_i} + M_{N_j} = M_{N_i+N_j}$$

Repeating this addition multiple times allows us to achieve simulations with a high number of samples. We can then assess this cumulative process as a temporal one, where the indices  $N_i$  correspond to consecutive time steps. In that case, a dose simulation can be represented by a stochastic variable  $X_{N_i}$  that we observe over time, as the number of simulated photons grows. Then, considering a sequence  $(X_{N_1}, \dots, X_{N_T})$  with  $T$  observations of that variable, our denoising problem amounts to predicting the most likely observation  $X_{N_{T+1}}$  based on the given sequence:

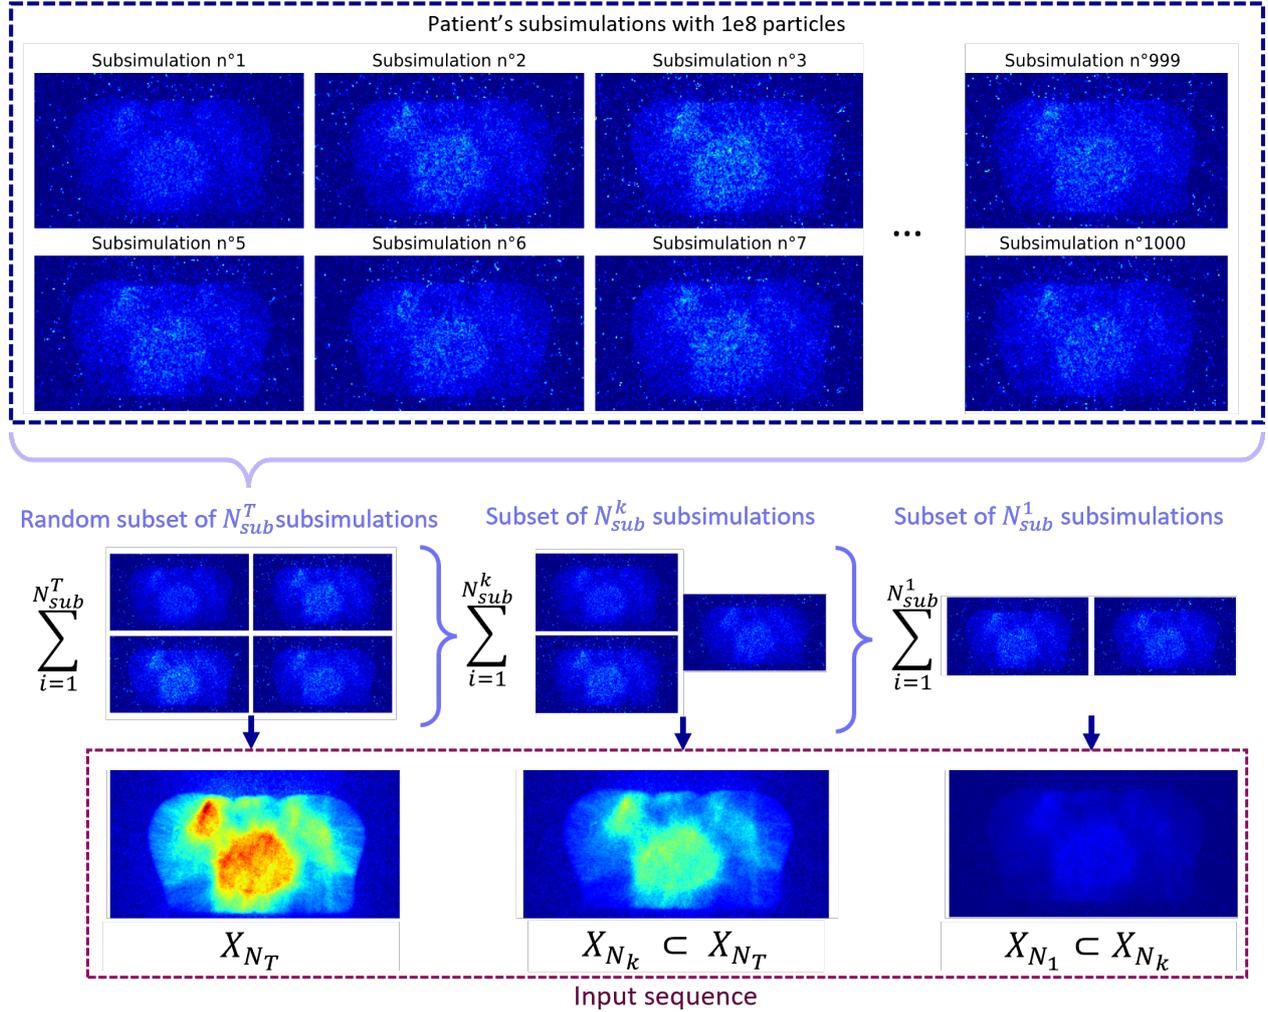


Figure 3.1: **Aggregation procedure of the sub-simulations to create input dose sequence.** To create the input sequence of MC dose simulations, a subset of sub-simulations is selected depending on the highest number of simulated particles in the sequence. These sub-simulations are then summed to yield the corresponding simulations. To create the noisier MC simulations of the sequence, subsets of the original subset are selected for aggregation. As a result, each noisier simulation is included in the simulation that has a higher level of precision.

$$X_{N_{T+1}} = \underset{X_{N_{T+1}}}{\operatorname{argmax}} p(X_{N_{T+1}} | X_{N_1}, \dots, X_{N_T})$$

where  $p$  denotes an unknown probability. For our denoising task, an observation of  $X_{N_i}$  is the radiotherapy dose delivered to a patient at each time step, i.e.  $M_{N_i}$ . Hence, there is a need to exploit both spatial and temporal information of the given sequence before inferring the highly sampled dose. This formulation allows us to exploit the temporal and spatial coherence

in the process of simulation.

**Dataset training aggregation pipeline:** During training, we compute the noisy sequence for each batch by aggregating sub-simulations on-the-fly to match the required noise levels, according to the pipeline presented in Figure 3.1. Let us consider an input sequence of length  $T$ . Each distribution  $X_{k|k=1..T}$  of the input sequence of dose distributions corresponds to  $N_{k|k=1..T}$  simulated particles for a given patient case. The aggregation of the sub-simulations takes place such that:

$$X_{N_1} \subset \dots \subset X_{N_k} \subset \dots \subset X_{N_T}$$

Therefore, the sub-simulations summed to yield  $X_{N_k}$  are also used to create all subsequent dose distributions of reduced noise levels  $X_{N_{t>k}}$ . This strategy ensures the continuous aggregation of sub-simulations and preserves information across noise levels.

### 3.1.2 Problem statement

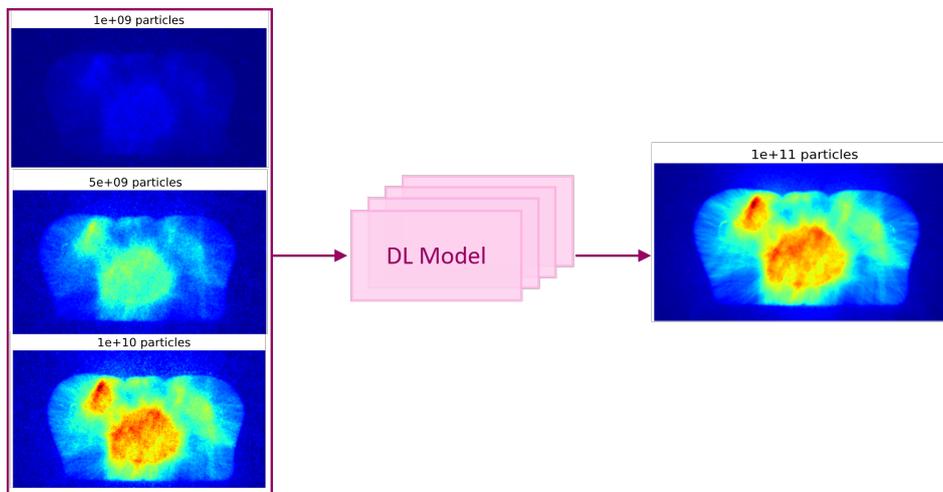


Figure 3.2: **Global studied learning pipeline.** The model learns to generate a precise dose distribution from a sequence of noisy MC simulations.

Provided with the new paradigm of section 3.1.1, the underlying question guiding the presented experiments in this chapter is: **Does training deep learning models on sequences of Monte-Carlo dose simulations lead to better and more accurate dose distributions?**

The global framework of study is displayed in Figure 3.2. Under the prism of the afore-

mentioned approach, we study several deep learning architectures in hope to find a model that yields optimal denoising performances under the studied hypothesis. In particular, we study a recurrent architecture called Convolutional Long-Short Term Memory (Shi et al. 2015) in both two dimensional and three-dimensional settings. We analyse performances of trained models in contrast to benchmarked models and discuss the limitations of the proposed frameworks through extensive ablation studies. Finally, we propose and investigate another architecture to disentangle noisy representations using feature embeddings that overcomes previous limitations.

## 3.2 ConvLSTMs for 2D denoising

### 3.2.1 Definitions

**Long Short-Term Memory (LSTM)** (Memory 2010) networks are a class of recurrent neural networks (RNN) (Schuster et al. 1997) designed to mitigate the vanishing gradient problem encountered by traditional RNNs. LSTMs excel in modeling sequential data with long-range dependencies.

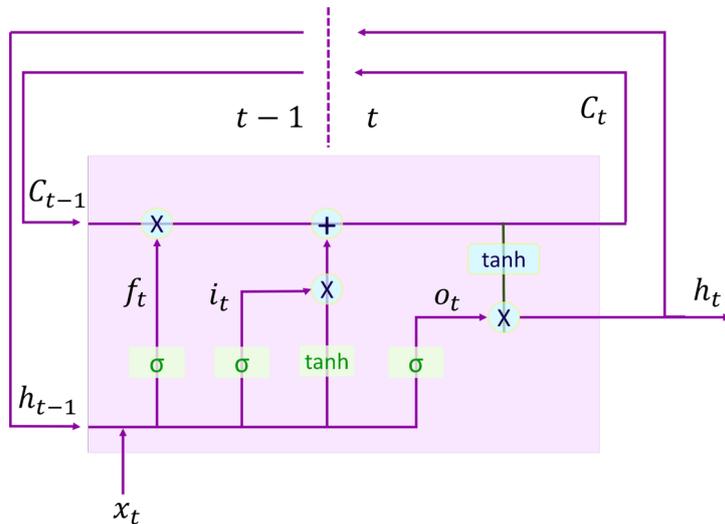


Figure 3.3: **Architecture of an LSTM cell.** An LSTM cell comprises several gates that fulfill a specific function each like forgetting information or keeping in memory meaningful features.

Let  $h_t$  signify the hidden state (output),  $W_f$ ,  $W_i$ ,  $W_c$ , and  $W_o$  weight matrices, while  $b_f$ ,  $b_i$ ,  $b_c$ , and  $b_o$  are bias terms.  $\sigma$  represents the sigmoid activation function, and  $\tanh$  corresponds to the hyperbolic tangent activation function. Then LSTM cells are characterized by a cell structure displayed in Figure 3.3, that comprises three pivotal gates:

1. **Forget Gate:** The input of the forget gate  $f_t$  is a concatenation of the previous hidden state  $h_{t-1}$  and the current input  $X_t$ . This combination allows the forget gate to consider both the historical context from  $h_{t-1}$  and the new information from  $X_t$ , when deciding what to forget.

The forget gate is responsible for determining what information from the previous cell state  $C_{t-1}$  should be retained or discarded. The cell state serves as a dynamic storage unit that can hold information relevant for the current time step while letting go of outdated or irrelevant information. The forget gate computes a value between 0 (forgetting) and 1 (retaining) for each element in the cell state.

$$f_t = \sigma(W_f \cdot [H_{t-1}, X_t] + b_f)$$

2. **Input Gate:** The activation of the input gate  $i_t$  governs the integration of new information into the cell state. It generates a candidate cell state  $\tilde{C}_t$  by blending the current input  $x_t$  and the previous cell state  $C_{t-1}$ , governed by the following equations:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$

$$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c)$$

3. **Output Gate:** It operates on the current input  $x_t$ , the previous cell state ( $C_{t-1}$ ), and the candidate cell state  $\tilde{C}_t$ , ultimately yielding the new cell state  $C_t$  and the current output  $h_t$ . Noted as  $o_t$ , the output gate controls the selection of information to be exposed as the output:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$$

$$C_t = f_t \cdot C_{t-1} + i_t \cdot \tilde{C}_t$$

$$h_t = o_t \cdot \tanh(C_t)$$

**Convolutional Long Short-Term Memory (ConvLSTM)** extends the LSTM architecture by integrating convolutional layers within the LSTM cells. In a ConvLSTM cell, the forget, input, and output gates operate akin to those in standard LSTM. However, the incorporation of spatial information is achieved through convolutional operations. The mathematical expressions remain similar, with convolutional operations denoted by  $*$  in the gate calculations.

This augmentation is particularly valuable in scenarios involving spatiotemporal data anal-

ysis, such as video processing, where both spatial and temporal dependencies play crucial roles. Applications using this architecture range from next frame prediction in videos to weather forecasting. For the first time, we conduct preliminary experiments with this architecture on the radiotherapy dose Monte-Carlo simulation denoising task.

### 3.2.2 Methods

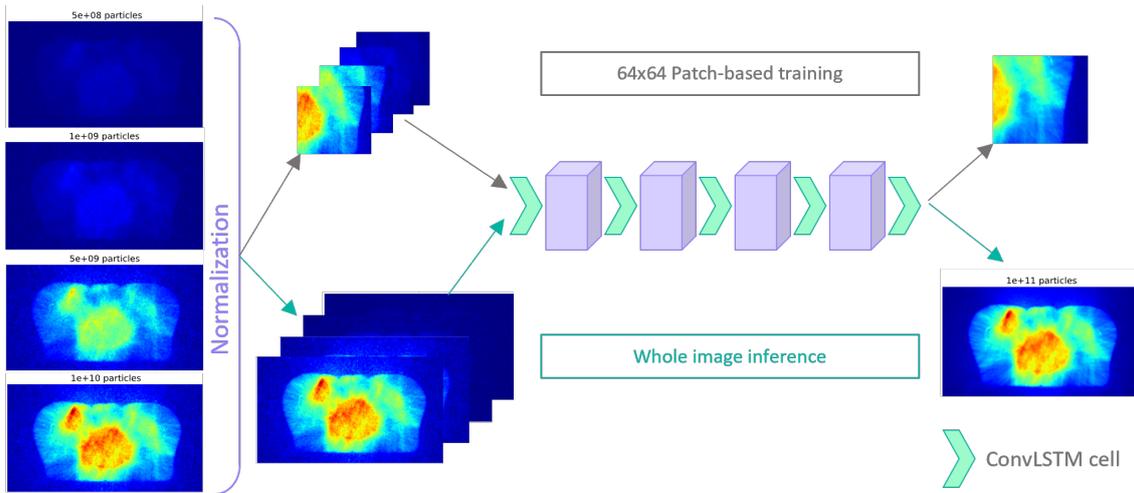


Figure 3.4: **Training and inference pipelines of the 2D ConvLSTM-based model.** The input sequence of MC simulations is first normalized prior to patch extraction for training. As the architecture is fully convolutional, whole images can directly be processed at inference time.

**Model:** Our model comprises 5 ConvLSTM cells stacked on top of each other. Nonetheless, the spectral dimensions has a varying number of channels across the architecture. No spatial downsampling is introduced and thus remain equal to that of the input through all propagation. All convolutional layers have  $3 \times 3$  filters. Each convolutional layer is followed by a Swish activation function (Ramachandran et al. 2017).

**Dataset preprocessing:** The input sequence of the model consists of 4 slices of a patient’s dose volume simulated with respectively 5e8, 1e9, 5e9 and 1e10 particles using MC. The denoised slice comes from the corresponding dose simulated with 1e11 particles. This slice constitute the optimal ground-truth of the input sample. Thus, the model is trained to infer highly sampled dose simulations from lower precision simulations. The training set comprises axial slices from 40 patients while the validation and test sets each draw slices from 5 patients. We select the slices in areas within 30%-100% of the maximum dose in the denoised simulations. The

normalization applied consists in dividing each input simulation by the maximum dose value computed over all ground-truth simulations from the training set. The same normalization is applied to samples from the validation and test sets. This normalization has the advantage that it keeps the physical fact that more simulated particles generate more deposited energy.

**Training scheme:** The training strategy is patch-based to encourage robustness of the model’s generalization power. We set the patch size to  $64 \times 64$  pixels i.e.  $12.8 \times 12.8 \text{ cm}^2$  for each slice. We use random horizontal and vertical flipping as sole augmentation techniques. We use AdamW (Loshchilov et al. 2017a) as optimizer and the linear combination of the MSE and the SSIM as loss function. The initial learning rate was set to  $10^{-3}$  and reduced by a factor of 10% each time the validation loss plateaued for more than 100 iterations. Early stopping was performed to stop the training.

### 3.2.3 Results

We compare the model’s performance to the Non Local Means (NLM) algorithm. To ensure a fair comparison, we set NLM to denoise simulations with  $1e10$  particles, i.e. the last slice of the input sequence of our model. Nonetheless, we additionally normalize the input simulation to further help NLM. We stress that without this normalization step, NLM fails to denoise the simulation in the sense that NLM does not account for the increase of deposited dose caused by the additional simulated particles.

Method	MSE	SSIM(%)	GPR 3%/3mm (%)	GPR 2%/2mm (%)
Raw $1e10$ particles simulation	3.13e-4	87.5	81.6±7.04	31.3±8.73
NLM on $1e10$ particles simulation	2.04e-4	94.4	80.5±11.2	35.9±13.2
<b>Stacked ConvLSTM</b>	<b>9.73e-5</b>	<b>97.3</b>	<b>83.2±12.9</b>	<b>42.7±15.2</b>

Table 3.1: Quantitative evaluation of the dose generated by our model and the NLM method.

We evaluate the model on the test set according to the GPR, MSE and SSIM which were presented in the previous chapter. The GPR is computed using a dose threshold of 10% and a dose-to-agreement / dose tolerance ratio of  $3mm/3\%$ . Table 3.1 displays the quantitative comparison.

Figure 3.5 displays the distributions for one sample case. The NLM denoised dose distribution still comprises noise. Figure 3.6 shows the gamma index matrices of the simulation

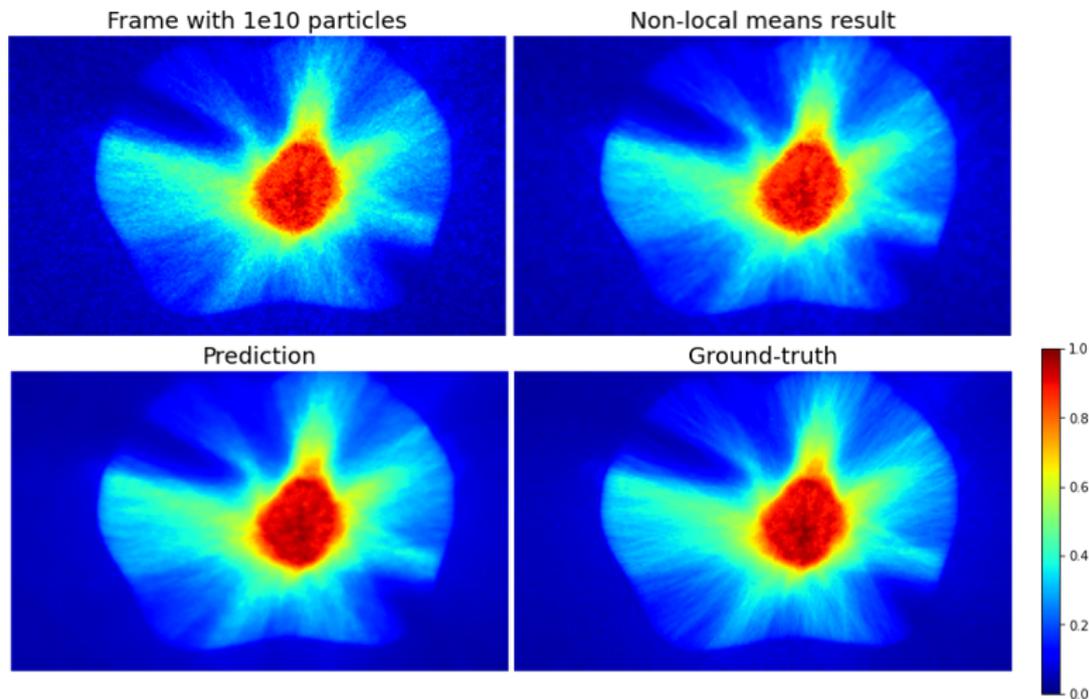


Figure 3.5: **Dose distributions generated by the model and NLM.** This figure displays slices of the less noisy dose simulation in the input sequence, NLM denoised distribution, model's prediction and reference dose. This visualization confirms the enhanced prediction of the ConvLSTM-based model compared to the classical NLM method, which distribution remains noisy.

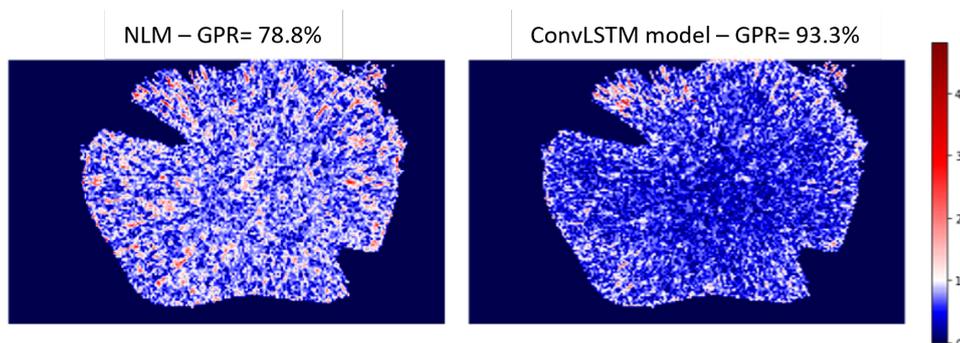


Figure 3.6: **Visualization of gamma index matrix for the NLM and our model's denoised dose distributions.** The prediction of the ConvLSTM model scores a better GPR than NLM. Blue voxels correspond to an accurately predicted dose while red voxels fail to satisfy the passing criterion of the GPR.

denoised by NLM compared with our model's. These matrices showcase the discrepancies in terms of GPR and highlights that our model outperforms NLM. Moreover, our model further decreases the inherent noise further in high dose regions.

### 3.2.4 Discussion

The results of this preliminary experiment showcase that our approach outperforms the traditional non local means algorithm. Moreover, our model not only denoise quite successfully the input simulations but also accounts for the increase in dose that occur conjointly with the increase in number of simulated particles. This means that our framework, designed to handle sequential information, succeeds in generating the extra missing dose corresponding to the ground-truth simulations. However, 2D frameworks suffer from a lack of continuity in the prediction. When considering a whole dose volume, each axial slice of the volume is processed by the 2D model. The reconstructed predicted volume lacks smoothness along the axial view, indicating a bad continuity of predictions impeding the generation of accurate 3D dose distributions. Therefore, these first results encouraged us to directly pursue experiments with our recurrent framework and generalize it to handle dose volumes.

## 3.3 3D-ConvLSTM

### 3.3.1 Introduction

Due to the limitations of the 2D approach, we generalize the ConvLSTM framework to cope with three-dimensional inputs for denoising of whole 3D dose distributions. As 3D trainings usually require an increased amount of GPU resources, we also investigate incorporating 3D ConvLSTM cells into a 3D UNet (Çiçek et al. 2016b), hoping that the spatial downsampling that occurs can alleviate the computational bottleneck.

This study aims to demonstrate that this novel recursive framework harnesses its strength from the sequential nature of its input and its ability to derive correlation between the levels of noisiness induced by the different number of particles simulated in the 3D space.

### 3.3.2 Methods

**3DConvLSTM cells:** As we are considering data that present spatial information in three dimensions, we extended the ConvLSTM framework to deal with temporal sequences of 3D volumes. This can be achieved by using 3D convolutional operators indicated by  $*$ . In that structure,  $W_z$  and  $b_z$  in the equations below denote the parameters (filters and bias) of the considered convolutional layers.  $\odot$  stands for the Hadamard product and  $\sigma$  for the sigmoid function. The following equations (3.1-3.5) describe how gates are activated and states modified:

$$i_t = \sigma(W_{xi} * X_t + W_{hi} * H_{t-1} + W_{ci} \odot C_{t-1} + b_i) \quad (3.1)$$

$$f_t = \sigma(W_{xf} * X_t + W_{hf} * H_{t-1} + W_{cf} \odot C_{t-1} + b_f) \quad (3.2)$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tanh(W_{xc} * X_t + W_{hc} * H_{t-1} + b_c) \quad (3.3)$$

$$o_t = \sigma(W_{xo} * X_t + W_{co} \odot C_t + b_o) \quad (3.4)$$

$$H_t = o_t \odot \tanh(C_t) \quad (3.5)$$

All the states of the 3D ConvLSTM cell are initialized with zeros which corresponds to ignorance of the future states. To ensure that dimensions match between the input and the various states inside the cell, padding is applied before convolutions. Therefore, the output of the 3DConvLSTM has the same spatial dimensions as the input.

This extension of ConvLSTM allows processing of medical volumetric sequential data in a fully convolutional manner. Each voxel's future state can be seamlessly predicted using contextual information brought by both temporal and spatial features from its own and its neighbours' past states in all dimensions. In the following subsections we present the two different setups that we used to integrate the 3D ConvLSTM cells.

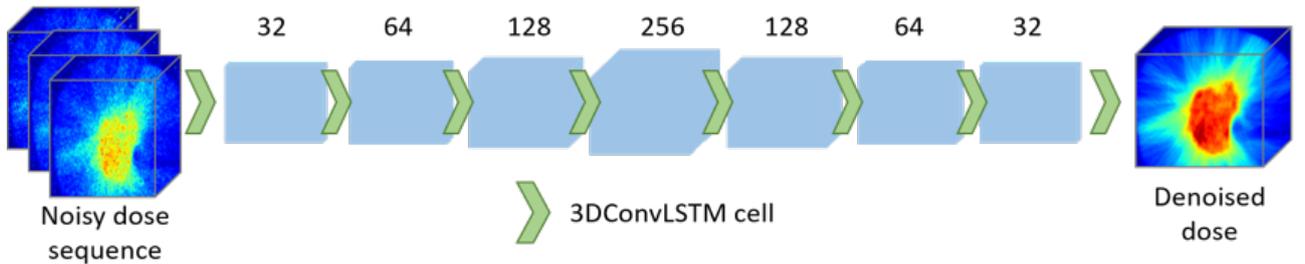


Figure 3.7: **Proposed architecture with stacked 3D ConvLSTM cells.** The number of output channels after each block appears above or below the layers' output volumes.

**Proposed model with stacked 3DConvLSTM cells:** The model consists of 7 cells 3D ConvLSTM stacked on top of each other, without introducing any spatial downsampling. All

convolutional layers in the 3D ConvLSTM cells contain  $3 \times 3 \times 3$  filters. Figure 3.7 shows the architecture of the Stacked ConvLSTM model. The spatial dimensions remain equal to that of the input through all propagation.

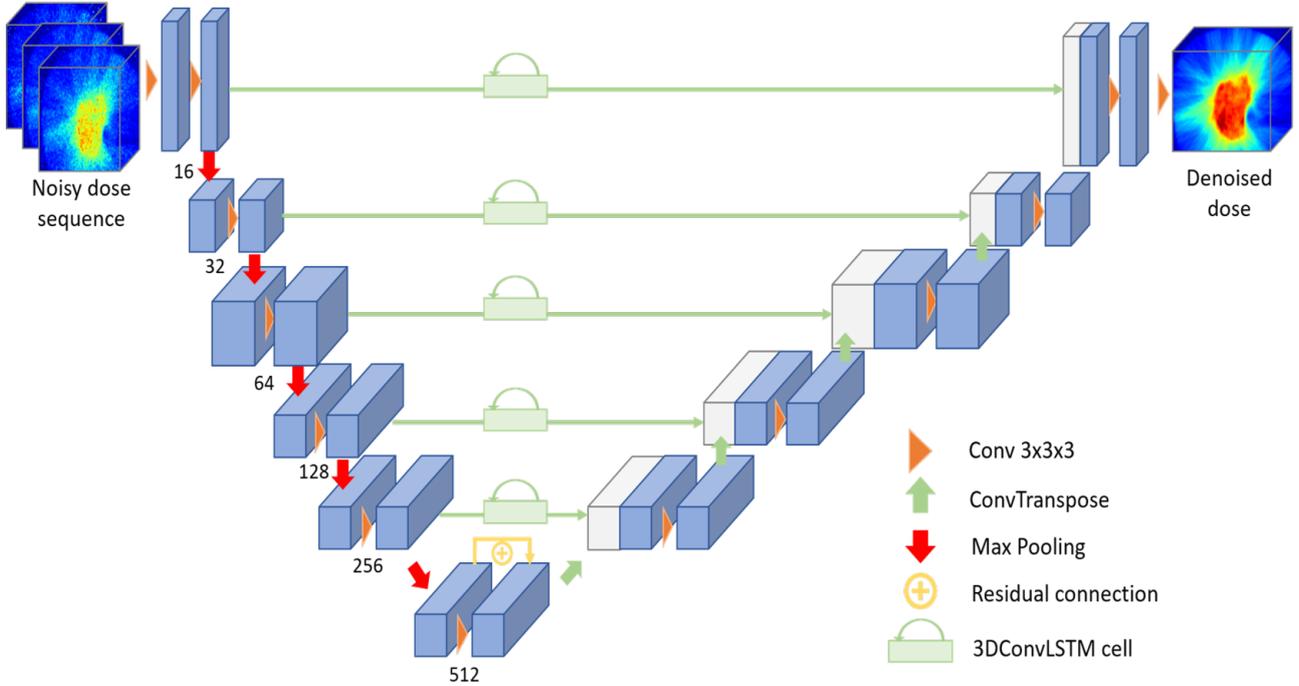


Figure 3.8: **3D Unet enhanced with 3D ConvLSTM cells in the skip connections.** The introduction of 3D ConvLSTM cells in each skip connection of the UNet enables the model to process sequences of volumes.

**Proposed UNet-based models with 3DConvLSTM cells in skip connections:** We also introduce a model based on the 3D UNet [Çiçek et al. 2016a](#) architecture and enhanced with 3D ConvLSTM cells in the skip connections to further extract features at each of the 5 down-sampling steps. Down-sampling is performed using max pooling with kernel size and stride of 2 and comprises two identical convolutional layers. All convolutions have  $3 \times 3 \times 3$  filters. The bottleneck has two identical convolutions with a residual connection to further exploit the deep features. The up-sampling blocks have a transpose convolution for up-sampling and a regular convolutional layer for further processing. Each convolutional layer is ended with a LeakyReLU [Maas 2013](#) activation function, and batch normalization [Ioffe et al. 2015](#) is used for faster convergence. Details regarding the number of channels at each stage are shown in Figure 3.8. This model is trained in the same setting as the proposed model with stacked 3D

ConvLSTM cells.

### 3.3.3 Methods

**Dataset** : The dataset was split to 40, 5 and 5 patients for training, validation and test respectively. Anatomies were distributed as evenly as possible between these sets. The input sequence comprises 3 decreasingly noisy dose volumes simulated with  $5 \times 10^8$ ,  $10^9$  and  $5 \times 10^9$  particles of the same patient case. We use random horizontal and vertical flipping as sole augmentation techniques. The ground-truth was the corresponding highly sampled simulation with  $10^{11}$  particles. Each sample was selected and fed to the model along the axial view.

**Loss function:** To optimize the weights of the considered models, we considered a hybrid loss function that consists in a weighted sum of the SSIM and the L1 loss (3.6):

$$\mathcal{L} = \sum_{i=0}^{N_{samples}} \left( \left\| X_{N_{T+1}}^{(i,estimated)} - X_{N_{T+1}}^i \right\|_1 + SSIM \left( X_{N_{T+1}}^{(i,estimated)}, X_{N_{T+1}}^i \right) \right) \quad (3.6)$$

where  $X_{N_{T+1}}^{(i,estimated)}$  is the model’s estimation of the  $i$ -th denoised dose volume sample.

**Optimization:** A patch-based training was implemented by randomly selecting sub-volumes from the 3D input sequences - ground-truth pairs, in areas within 30% - 100% of the dose maximum. The patch size was 64  $mm^3$  subvolumes, i.e. 12.8  $cm^3$ . We used Adam optimizer with learning rate, weight decay, beta1, beta2 and epsilon parameters set to  $10^{-5}$ ,  $10^{-4}$ , 0.9, 0.999 and  $10^{-8}$  respectively. The learning rate was reduced by half when the validation loss stagnated, i.e. when difference in loss was inferior to  $1e^{-2}$  for more than 200 iterations. The batch size was set to 8. All models were trained for  $3 \cdot 10^5$  iterations. The final model we kept was the one that performed best on the validation set.

### 3.3.4 Results

**Benchmarked models:** We compare our method with other commonly used learning based denoising methods in the literature. Our first benchmarking model is a 3DUNet [Çiçek et al. 2016a](#) with 5 down-sampling blocks. The second one is Pix2Pix [Isola et al. 2017](#), a generative adversarial framework. Pix2Pix has been adapted to a 3D setting. Moreover, since we are considering smaller data in terms of height and width, we remove one down-sampling block

and the corresponding decoding block from the generative model. The adapted generator thus consists of 5 down-sampling blocks, giving a fair comparison with the proposed 3DUNet architecture. Since these models don't handle sequential data, the input is the last volume of the sequence fed to the recurrent architecture, i.e. the least noisy simulation of the sequence. Finally, we also compared with the recently proposed BiONet architecture [Xiang et al. 2020](#) after adapting it to 3D data and also limiting the number of down-sampling blocks.

Method	SSIM (%)	GPR (%)	L1	#
Inputs 5e9 particles	$58.1 \pm 0.1$	$59.1 \pm 2.1$	$0.149 \pm 0.050$	
3DUNet <a href="#">Ronneberger et al. 2015a</a>	$80.0 \pm 2.4$	$61.2 \pm 2.8$	$0.088 \pm 0.007$	10 M
Pix2Pix 3D <a href="#">Isola et al. 2017</a>	$55.4 \pm 8.6$	$66.6 \pm 14.4$	$0.102 \pm 0.009$	120 M
3D BiONet <a href="#">Xiang et al. 2020</a>	$93.0 \pm 0.2$	$90.6 \pm 1.2$	$0.080 \pm 0.001$	178 M
3DUNet ConvLSTM	$64.5 \pm 6.1$	$79.1 \pm 1.2$	$0.037 \pm 0.004$	36 M
Stacked 2D ConvLSTM	$81.6 \pm 3.2$	$74.1 \pm 3.1$	$0.021 \pm 0.003$	1.5 M
Stacked 3D ConvLSTM	<b><math>97.9 \pm 0.9</math></b>	<b><math>94.1 \pm 1.2</math></b>	<b><math>0.004 \pm 0.001</math></b>	5 M

Table 3.2: **Evaluation metrics for the performance of the models on the test set.** The 3D ConvLSTM-based models are benchmarked against several deep learning architectures and also against our 2D ConvLSTM model.

**Quantitative results:** Extensive quantitative comparison using the L1 error, SSIM and gamma passing rate (GPR) for each model on the test set are presented in Table 3.2. We evaluate the GPR criteria with a dose to agreement and tolerance on dose values of 3%/3mm within 30% - 100% of maximum dose. Results show that Stacked 3DConvLSTM outperforms all benchmark models in all metrics while having the lowest number of trainable parameters. We also trained the original ConvLSTMs, on slices of dose volumes. Results in Table 3.2 reveal that the 2D version still performs better than 3DUNet and Pix2Pix3D with regards to all metrics with only 1.5 million parameters indicating the need of sequential data for this task. Nevertheless, it does not outperform its 3D counterpart nor 3D BiONet. This fact highlights that our 3D model as well as 3D BiONet extract volumetric features that greatly improve the quality of the predictions. Moreover, Stacked 3D ConvLSTM achieves the lowest L1 value and displays GPR scores with standard deviations of 1.2. Although BiONet also shows comparative robustness in its predictions, the quality of the denoised dose volumes remain inferior to that of our proposed models. Another remark stemming from these results is that, despite having a higher SSIM than Pix2Pix3D, 3DUNet's GPR is lower. This might indicate that 3DUNet is able to infer structural coherence in the dose volumes but lacks in precision at a voxel level. In contrast, the proposed recurrent 3DUNet outperforms 3DUNet on the GPR by 18% even

though its SSIM score fails to match the 3DUNet.

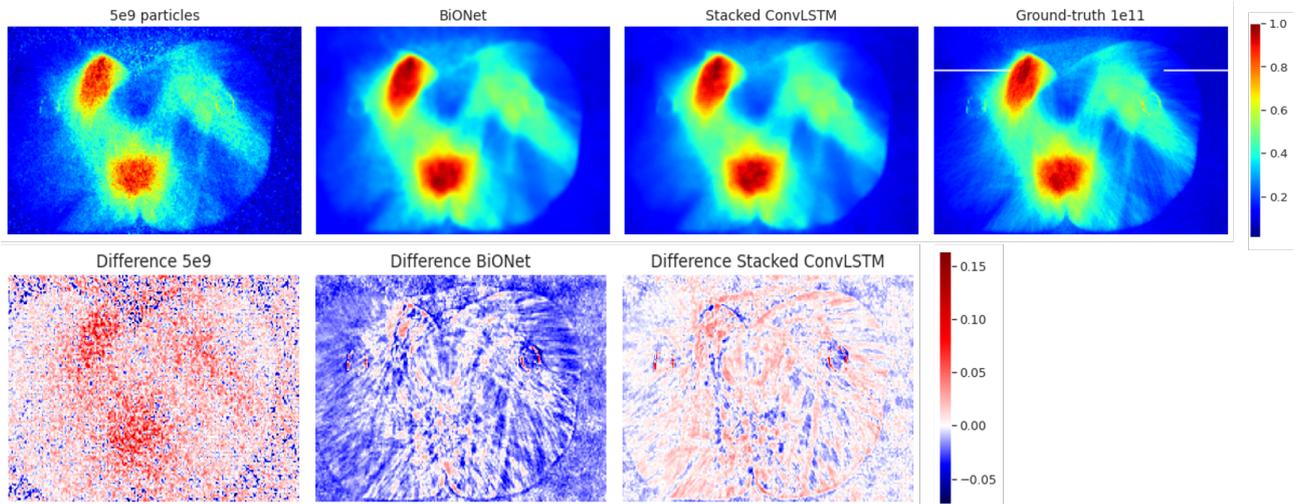


Figure 3.9: **Comparison of dose distributions generated by 3D ConvLSTMs and benchmarked methods.** On the first row from left to right: a single slice of the  $5 \cdot 10^9$  dose volume, BiONet’s, Stacked 3D ConvLSTM’s predictions and ground-truth  $1 \cdot 10^{11}$  dose volume. On the second row from right to left error maps for the three different representations.

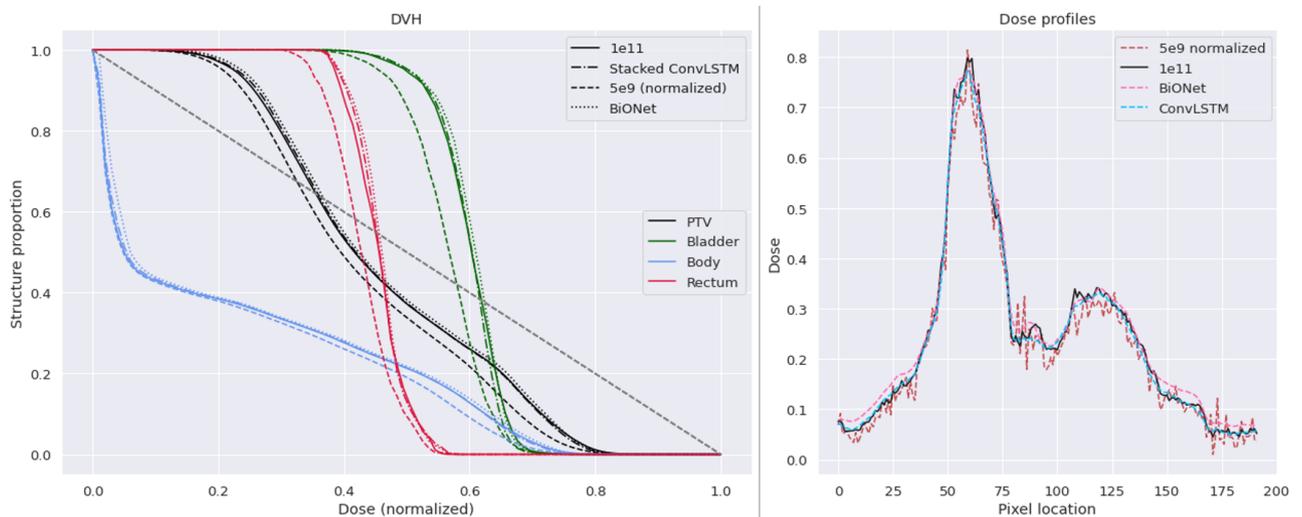


Figure 3.10: **Comparison of dose profiles and DVHs generated by 3D ConvLSTMs and benchmarked methods.** Left: DVH curves, showing that the dose gradients are reproduced faithfully. Right: Dose profile along the line indicated on the ground truth image of Figure 3.9.

**Qualitative results:** Figure 3.9 shows the predictions of the best performing models, namely BiONet and the Stacked 3DConvLSTM, on a test case. Both models reproduce high dose regions well. To further assess the denoising ability of the models, dose profiles are provided in Figure 3.10. Both models succeed in smoothing the noise of the low-simulation input. Error maps between predictions and ground-truth dose associated with BiONet point out that BiONet globally overestimates the dose in low dose gradient regions. Stacked 3DConvLSTM performs better in those regions but underestimates dose in high dose gradient regions where denoising is expected to be more challenging.

However, we can notice that both models unfortunately tend to smooth fine details of the Monte-Carlo ground-truth simulation. Figure 3.10 plots the dose volume histogram (DVH) corresponding to the patient studied in Figure 3.9. Both models substantially improve the DVHs towards the ground-truth DVHs. Nevertheless, the DVHs of BiONet indicate that the model still slightly overestimates the dose in voxels, in contrary to the Stacked 3DConvLSTMs.

### 3.3.5 Discussion

Independently of GPU-accelerated computation, MC simulation time can be further decreased using deep learning-based frameworks. The goal of this work is to highlight how considering the MC simulation task as a spatiotemporal problem can be an asset to reach accurate and fast computation of dose. Extensive experiments and comparisons with other state of the art methods highlight the potential of our method. However, the fact that our model does not perform any spatial down-sampling implies that the required GPU memory usage could still be reduced. Achieving high GPR scores while decreasing the computational load could enable real-time Monte-Carlo dose simulation. Future work aims to reduce the number of simulated particles, or in other words increase the level of noise of the input dose volumes.

## 3.4 Ablation studies

Following these results, we performed ablation studies aiming at quantifying the impact of changes in input data and architecture.

### 3.4.1 Alterations in the input data

**Length of input sequence:** In prior works, the input sequence comprises 3 decreasingly noisy Monte-Carlo simulation of the radiotherapy dose of a single patient. Each simulation

corresponds to a number of simulated particles.

Input length	Test GPR	Test SSIM	Test L1 ( $\times 10^{-3}$ )
1	$91.98 \pm 8.43$	$97.20 \pm 0.64$	$5.70 \pm 1.01$
2	$90.62 \pm 1.31$	$96.04 \pm 1.01$	$7.45 \pm 1.64$
3	$95.40 \pm 1.37$	$97.34 \pm 0.63$	$5.46 \pm 1.08$

Table 3.3: Impact of the sequence length on quantitative performances of 3D ConvLSTMs results.

In this experiment, we train the model with the input comprising only of the first and noisiest simulation or the two noisiest simulations. Results are displayed in table 3.3 and showcase that the performance of the model deteriorates when the noise level increases. Additionally, with a single input, the model displays instability in the GPR values, underlining a reduced robustness of the model towards outliers and a loss of generalization power.

**Impact of the CT:** We added the CT as a second channel in the input of the model while training. Table 3.4 compares both trainings and shows that adding the CT as extra information in the input does not improve nor jeopardize the model’s performance significantly.

CT	Test GPR	Test SSIM	Test L1 ( $\times 10^{-3}$ )
No	$95.31 \pm 1.55$	$97.26 \pm 0.65$	$5.58 \pm 1.10$
Yes	$95.27 \pm 1.66$	$97.27 \pm 0.63$	$5.57 \pm 1.08$

Table 3.4: Lightweight 3DConvLSTM: Impact of CT addition to the input of the model.

### 3.4.2 Towards lighter 3D ConvLSTM

By reducing the number of output channels of each ConvLSTM layer in the model, we lowered the number of parameters almost by a factor 10. Performances are reported in 3.5 and point out that the decrease in complexity does not impact the performances negatively. This means that fewer parameters allow efficient denoising of the Mone-Carlo dataset.

N parameters	Test GPR	Test SSIM	Test L1 ( $\times 10^{-3}$ )
$3 \times 10^6$	$95.89 \pm 1.35$	$97.50 \pm 0.73$	$5.25 \pm 1.25$
$5 \times 10^5$	$95.31 \pm 1.55$	$97.26 \pm 0.65$	$5.58 \pm 1.10$

Table 3.5: Impact of complexity reduction of 3D ConvLSTMs on quantitative metrics.

### 3.4.3 Evaluation per anatomy area

Metric	Pelvis + Thorax	ORL
GPR (%)	$96.32 \pm 1.14$	$94.93 \pm 0.84$
SSIM (%)	$97.30 \pm 0.31$	$97.62 \pm 0.92$
L1 ( $\times 10^{-3}$ )	$5.53 \pm 0.58$	$4.98 \pm 1.55$

Table 3.6: Quantitative results per anatomy group on the test set for the lightweight 3D ConvLSTM model.

Table 3.6 displays quantitative results depending on the type of anatomy considered. Performances displayed highlight a satisfactory generalization to various anatomies. A slight discrepancy can be noted when comparing the GPR values of ORL patients versus pelvis/thorax patients. Nonetheless, ORL structures - PTV or OAR - are usually smaller than structures encountered in pelvis/thorax patients. This difference of size can explain smaller GPR values in the case of ORL patients.

### 3.4.4 Discussion

We carried out ablation studies to determine the sensitivity of our framework to several parameters: the length of the input sequence comprising noisy MC simulations, considering the CT as physical prior and input to the model, and reducing the number of trainable parameters in the model. Results show that a lower number of parameters deteriorates performances of the model in terms of GPR and robustness to outliers. Adding the CT as input to the model did not lead to better results, nor did it worsen the performances. We hypothesize that no additional information can be found in the CT that is not already present in the input sequence of MC simulations. Reducing the number of trainable parameters in the model by a factor of 10 did not have any impact on the performances. These results suggest that the task of denoising 3D MC simulations is either too simple or that the dataset does not comprise enough variability in 3D. Based on these ablation studies, the following experiments either focus on using a weakly supervised training setting to overcome the small dataset size in 3D, or are conducted in 2D with a supervised training to yield significant results.

## 3.5 Weakly-Supervised ConvLSTMs

### 3.5.1 Introduction

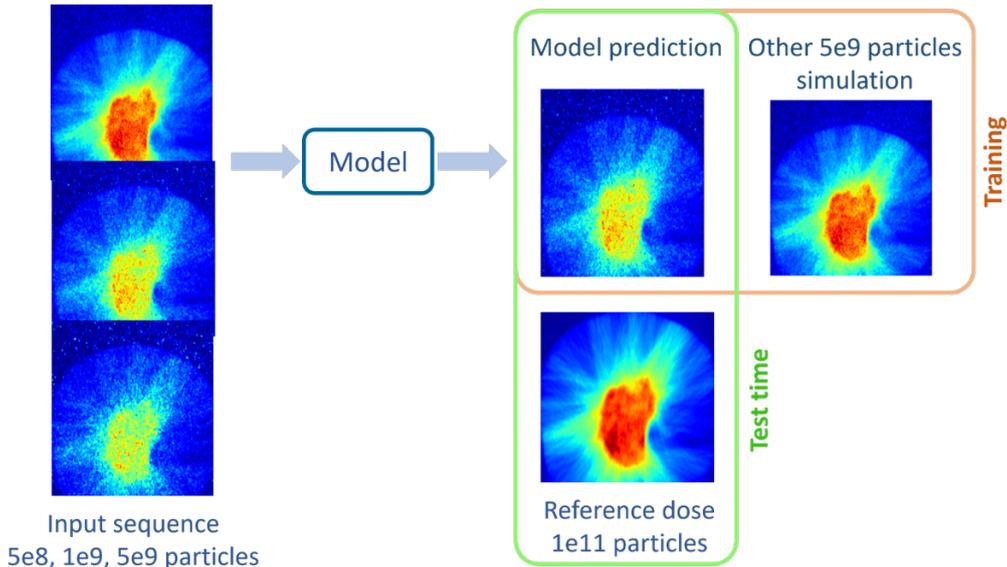


Figure 3.11: **Noise2Noise training pipeline applied to our sequence denoising task.** The optimization of the model’s parameters during training relies on the comparison of the model’s prediction with another noisy simulation for the same patient.

The ablation studies in the previous section pointed to the small dataset size. Indeed, considering a traditional supervised training strategy calls for a substantial dataset, which in our case implies high-precision and therefore expensive-to-compute ground-truth dose simulations. To circumvent this impeding requirement, we leverage the weakly supervised Noise2Noise (N2N) method [Lehtinen et al. 2018](#).

The N2N framework is a deep learning paradigm aimed at denoising images or data without requiring clean reference samples during training. It leverages pairs of noisy observations,  $x_1$  and  $x_2$ , of the same underlying clean data  $y$  to train a neural network  $F$  to predict the clean data from the noisy inputs, thus minimizing a suitable loss function  $L$  between the network’s output,  $F(x_1)$  or  $F(x_2)$  and one of the noisy inputs, i.e. minimizing  $L(F(x_1), x_1)$  or  $L(F(x_2), x_2)$ . N2N has found application in various image restoration tasks, including denoising, deblurring, super-resolution, and inpainting, as it offers a robust solution when access to clean reference data is challenging or non-existent. Using this weakly-supervised framework we take full advantage of

the abundance of cheap-to-compute noisy dose simulations at our disposal to train our model. Using the N2N framework, the objective of this section is to confirm that 3D ConvLSTMs outperform other models as suggested by the results of sections 3.2 and 3.3.

### 3.5.2 Methods

**Dataset:** The cohort described in section 2.1 was split to 40, 5 and 5 patients for training, validation and test respectively, while anatomies were distributed as evenly as possible across sets. As precised in table 2.1 of section 2.1, to each patient correspond 20 independent noisy MC dose distributions computed with  $5e9$  simulated particles. The input-output pairs during training are comprised of a noisy sequence for the input, and the corrupted reference is another distinct simulation with  $5e9$  particles. The training and testing pipeline is displayed in Figure 3.11. Following the N2N framework, the training loss is computed between the model’s prediction and the corrupted reference, i.e. the distinct noisy MC simulation during training. At inference we performed direct denoising of entire dose volume, enabled by the fully convolutional character of the ConvLSTM architecture.

**Optimization:** For all considered architectures, training was patch-based to promote better robustness of the models. The patch size was  $32 \times 32 \times 32$  i.e.  $6.4 \text{ mm}^3$  and the batch size was 8, to match our hardware limitations in terms of available GPU RAM. The training lasted until the difference in validation loss was inferior to  $10e^{-2}$  for more than 1000 iterations using AdamW optimizer (learning rate=0.001, betas=(0.9, 0.999), weight decay=0.01). The loss function combines the SSIM and the L1 loss as follows:  $Loss = 20 \times L1 + SSIM$ . The indicated weighing coefficient presented here was found after thorough hyperparameter search and yielded the best performing model on the test set.

### 3.5.3 Results

The results of our experiments in section 3.3 promoted 3D BiONet as sole competitor to our recurrent framework. Thus, in this section we benchmarked our model against 3D BiONet trained with the N2N scheme. Furthermore, we carried out two trainings for BiONet, while keeping the same number of trainable parameters:

- The full input sequence was given as input to BiONet, as it was to the 3D ConvLSTMs.
- Only the noisy MC dose distribution computed by simulating  $5 \times 10^9$  particles was given as input.

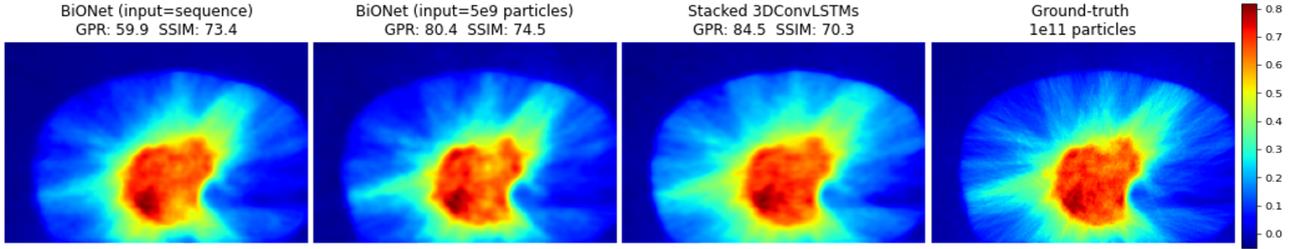


Figure 3.12: Visual comparison of denoised volumes by each model and the ground truth.

Figure 3.12 displays slices of denoised volumes by each model for further qualitative comparison. Unfortunately, we observe that both models give visually smoothed predictions compared to the fine-grained ground-truth.

Model	SSIM (%)	L1	GPR (%)
3D BiONet (input: noisy input sequence)	<b>91.1±2.2</b>	$1.02e^{-1} \pm 2.2e^{-2}$	50.5±4.8
3D BiONet (input: noisy $5e^9$ particles)	90.4 ± 3.4	<b><math>7.22e^{-2} \pm 1.1e^{-2}</math></b>	71.9±5.1
Stacked 3D ConvLSTMs	86.4 ± 5.2	$1.04e^{-1} \pm 2.2e^{-2}$	<b>83.8±3.4</b>

Table 3.7: Quantitative results for the different evaluated models.

Table 3.7 shows the SSIM, the L1 loss and the GPR computed between the models’ denoised volume and the corresponding ground-truth dose volume (simulated with  $1e^{11}$  particles) over the test set. The GPR is computed with a dose to agreement and tolerance on dose values of 3%/3mm within 30% - 100% of maximum dose.

Regarding the input given to BiONet during training, results suggest that BiONet does not extract meaningful information from the given input sequence. Indeed, its performance with respect to the GPR are drastically improved by 20% when using the  $5 \times 10^9$  simulation alone. We note that the SSIM scores are equivalent between the two training and outperform 3D ConvLSTMs. Yet, despite having the lowest number of parameters, the 3DConvLSTMs remain unrivaled with respect to the GPR values.

### 3.5.4 Discussion

In this experiment, we leveraged the N2N framework to overcome potential limitations linked to the small dataset size when considering full 3D distributions. This weakly supervised setting allow us to fully exploit the abundance of low precision MC simulations within our dataset to train models. We benchmarked our model against BiONet. Performances of both remained competitive in terms of SSIM. However, 3D ConvLSTMs outperformed BiONet by 20% with respect to GPR values. A side conclusion from the experiments stems from comparing the

impact of the input simulation on the performances of BiONet. When training BiONet using full input sequences as input while preserving the number of trainable parameters, BiONet’s performances deteriorates drastically. This suggests that BiONet’s architecture is ill-suited to handle sequences, despite having a distinctly higher number of trainable parameters than our model comprising 3D ConvLSTMs.

## 3.6 Conclusion

In this chapter, we conducted several experiments to determine the validity of our approach to accelerate MC simulations of dose distributions. Our approach consisted in predicting a high precision dose using an input sequence of low precision simulations as input to the model, therefore taking advantage of the information regarding the deposition of the dose and spatial correlations comprised in the sequence. To verify this approach, we investigated a recurrent yet fully convolutional architecture, ConvLSTMs, that we generalized to handle 3D volumes. We benchmarked our ConvLSTM-based models against state-of-the-art and classical methods and evaluated the quality of the generated dose maps using the metrics presented in section 2.3. Results indicate that 3D ConvLSTMs outperform other methods, especially with respect to the clinical metric (GPR).

We performed ablation studies to further test the sensitivity of our architecture to changes in the input sequence, physical prior and architectural complexity. The main limitation arising from our experiments is the small dataset size when considering 3D trainings. We therefore explored a weakly supervised setting to bypass this challenge and validated the results observed under a supervised training. 3D ConvLSTMs display drastically improved GPR values than other models at test time. Moreover, ConvLSTMs seem to take full advantage of the input sequence of MC simulations while its contenders fail to extract meaningful features and perform best when considering a single input simulation, i.e. in a one-to-one prediction setting. Thus, ConvLSTMs show promising performances and seem to confirm the advantages of our many-to-one approach to denoise and accelerate MC simulation. Nevertheless, this architecture necessitates computing the input sequence, even at inference time, to generate a high precision distribution. This requirement boils down to an additional computational bottleneck problem that we try to address in the following chapter.

# Chapter 4

## Accelerating Monte-Carlo with intermediate representations

### Contents

---

<b>4.1</b>	<b>Introduction</b>	<b>83</b>
<b>4.2</b>	<b>Cascaded Denoising</b>	<b>84</b>
4.2.1	Introduction	84
4.2.2	Methods	84
4.2.3	Results	86
4.2.4	Discussion	87
<b>4.3</b>	<b>Deep Particle Embeddings</b>	<b>87</b>
4.3.1	Introduction	87
4.3.2	Methods	87
4.3.3	Results	90
<b>4.4</b>	<b>Conclusion</b>	<b>91</b>

---

### 4.1 Introduction

In the previous chapter, we studied ConvLSTM-based approaches to generate precise dose distribution from a sequence of Monte-Carlo simulations with higher noise levels. Nevertheless, a significant drawback of this approach lies in the substantial computational resources needed

to simulate input sequences for both model training and inference. In this chapter, we explore frameworks that circumvent this impediment at inference time. These architectures not only enhance the speed of dose prediction but also make effective use of our sequential perspective throughout the training process. We precise that we conducted the following experiments in 2D to emancipate trainings from the data scarcity issue and encourage significant results and evaluation of the models' generalization power.

## 4.2 Cascaded Denoising

### 4.2.1 Introduction

Cascaded U-Nets (Christ et al. 2016) represent a powerful architectural innovation in the field of image processing, particularly in medical imaging and computer vision applications like medical image segmentation (Aswathy et al. 2022; Lachinov et al. 2018; H. Liu et al. 2019). These networks are an extension of the popular U-Net architecture, known for its exceptional performance in tasks like image segmentation and object detection. The concept behind Cascaded U-Nets is to stack multiple U-Net models sequentially, creating a cascading structure that refines and enhances the accuracy of the output at each stage. The primary use of Cascaded U-Nets is to tackle complex image processing challenges where high levels of precision and detail are required. In this work we try to leverage the refining power of cascaded networks to denoise MC simulations while alleviating the computational bottleneck of simulating the input dose distributions at inference time.

### 4.2.2 Methods

**Pipeline:** Figure 4.1 displays the training pipeline to achieve our goal. First, we train a 2D UNet to denoise pairs of incrementally noisy simulations. More precisely, this first training focuses on generating from  $10^8$  particles to  $10^9$ , from  $10^9$  to  $10^{10}$  and from  $10^{10}$  to  $10^{11}$ . Secondly, we duplicate this pre-trained model into 3 identical models. We finetune each of them on one single transition -  $10^8$  to  $10^9$ ,  $10^9$  to  $10^{10}$ , or  $10^{10}$  to  $10^{11}$ . After convergence of these finetuned models, we cascade them, thus yielding a global model that directly process dose distributions simulated with  $10^8$  particles. The cascaded model finally generates a dose map similar to the high precision distribution simulated with  $10^{11}$  particles.

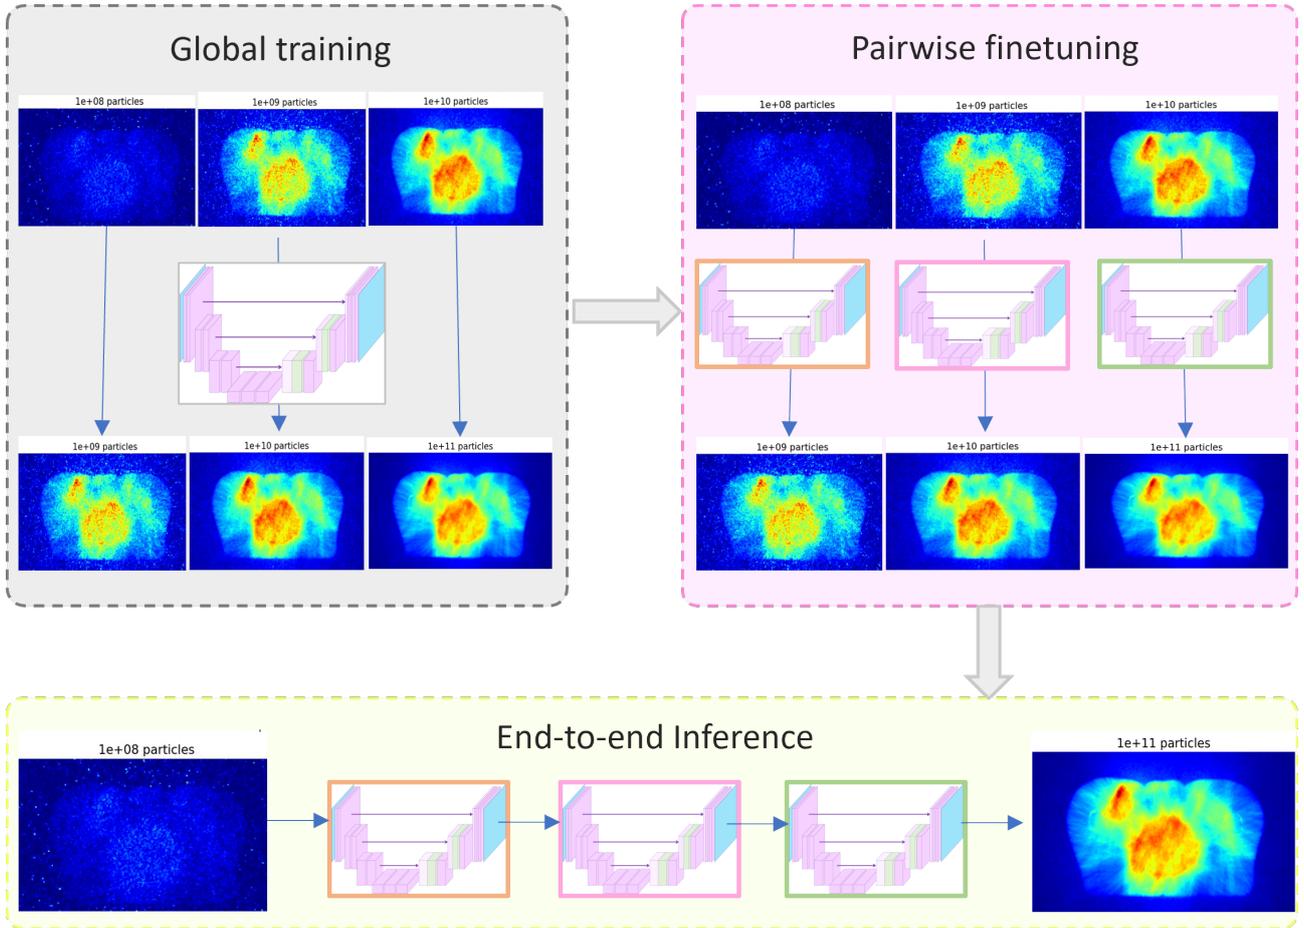


Figure 4.1: **Pipeline for end-to-end dose denoising with cascaded UNets.** The first step of the pipeline consists in training a 2D UNet to denoise pairs of MC simulation with incremental noise levels. Then, this model is duplicated into three separated versions that are each finetuned to map selected noise levels. Finally, these finetuned copies are cascaded to yield a denoising pipeline.

**Optimization:** The dataset detailed in section 2.1 was split in 35, 5 and 10 patients for respectively the train, validation and test sets with anatomies dispatched as evenly as possible across the partition. In contrast with previous experiments from Chapter 3, we added 5 patients to the test from the train set to help better assess the behavior of the model at test time. Each sample consisted of a slice of noisy MC simulation and its corresponding reference slice in the ground-truth dose volume. Slices were selected in areas with dose values superior to 10% of the maximum dose. We trained our model on whole images. To that end, we enabled batch training by padding each dose map in the datasets to the size  $256 \times 256$ .

Our 2D UNet model comprised 3 downsampling-upsampling stages. Convolutional blocks in the encoder comprised 2 convolutional layers with Swish activation functions (Hendrycks et al.

2016) and residual connections combined with group normalization, following the suggestions in Kolesnikov et al. 2020. The bottleneck comprised 4 convolutional layers with residual connections. Overall the architecture is very similar to the UNet model modified and presented in Ho et al. 2020. However, our model does not perform positional encoding, nor does it include any kind of attention module.

We trained the models using the AdamW optimizer with default parameter values and an initial learning rate of  $10^{-3}$ . We performed early stopping to decide when to stop the training. The learning rate was decayed by a factor of 10% each time the validation loss plateaued for more than 200 iterations. Batch size was set to 16. We used a hybrid loss function combining the SSIM and the L1 loss similarly to experiments in chapter 3.

### 4.2.3 Results

Denoising level	Global model	Finetuned models
1e8 to 1e9	$49.39 \pm 6.94$	$48.42 \pm 3.99$
1e9 to 1e10	$73.12 \pm 1.97$	$80.33 \pm 3.18$
1e10 to 1e11	$70.90 \pm 9.77$	$74.03 \pm 9.17$
End-to-end models		
1e8 to 1e10	$51.24 \pm 6.47$	$64.32 \pm 7.56$
1e9 to 1e11	0	$69.39 \pm 8.08$
1e8 to 1e11	0	$60.40 \pm 7.13$

Table 4.1: GPR (3mm/3%) values computed on the test set in areas within 20% of the maximum dose.

Quantitative results are displayed in table 4.1. The global model refers to the UNet trained on all pairs of MC simulations. The finetuned models are the global model duplicated and finetuned on each pair of MC simulations. Results show that cascading the global model with itself without further finetuning yields abismal GPR scores. Overall, finetuning improves on simply training a UNet to perform pairwise denoising. Yet, when cascading the finetuned models, results remain modest and the high standard deviation across all combinations suggests that the framework is not robust. As noise and error propagation seem to prevail when cascading the models, we argue that the cascaded framework may be ill-suited for our denoising task.

#### 4.2.4 Discussion

We tried to alleviate the computational requirements to generate the input data at inference with a cascaded UNet framework. Results remain modest, suggesting that the cascaded framework propagates error at inference. We use this experiment as baseline to benchmark against the next approach.

### 4.3 Deep Particle Embeddings

#### 4.3.1 Introduction

In previous experiments, we used ConvLSTMs to generate dose distributions with a quality rivaling Monte-Carlo precision. The input of this architecture was a sequence of low precision MC simulations. Through our experiments, we gave proof that our ConvLSTM models take advantage of the extra information enclosed within the sequence and outperform other benchmarked models. Nevertheless, one limitation remains: the proposed framework still suffered from the impeding need to simulate the input sequence, resulting in a computation bottleneck prior to inference. The previous experiment with cascaded UNets in section 4.2 failed to provide a competitive alternative, most probably because of error propagation.

In this section, we present a novel framework that uses feature positional embeddings to circumvent both the computational overhead at inference time and bypass the issue of error propagation induced by separate model trainings. Additionally, the proposed framework promises to infer high precision dose distributions from any level of noise that was present in the training sequences of MC simulations.

#### 4.3.2 Methods

**Sinusoidal Positional Embeddings:** Sinusoidal positional embeddings (Vaswani et al. 2017), i.e. positional encodings, are a technique commonly used in natural language processing and computer vision tasks to provide deep learning models with information about the positions of elements in a sequence or grid-like data. These embeddings are designed to capture the positional relationships between elements without relying solely on sequential order. They are often used in conjunction with other embeddings, such as word embeddings in NLP or image embeddings in computer vision.

The idea behind sinusoidal positional embeddings is to encode positional information using

a set of sine and cosine functions with different frequencies and phases. These functions create distinct patterns that help the model differentiate between positions. Let  $pos$  represent the position of the element in the sequence or grid,  $i$  refer to the dimension or feature index within the embedding, and  $d_{model}$  be the dimensionality of the embeddings. Then sinusoidal positional embeddings  $PE$  are described with the following formula:

$$\begin{aligned} PE_{(pos,2i)} &= \sin\left(\frac{pos}{10000^{(2i/d_{model})}}\right) \\ PE_{(pos,2i+1)} &= \cos\left(\frac{pos}{10000^{(2i/d_{model})}}\right) \end{aligned}$$

The sinusoidal positional embeddings are computed for each position in the sequence (from 1 to the maximum sequence length) and for each dimension within the embedding. The use of both sine and cosine functions with varying frequencies and phases ensures that the embeddings encode different positional information for different dimensions. The choice of 10,000 in the formula is a commonly used constant but can be adjusted based on the problem and sequence length.

These sinusoidal positional embeddings are then added element-wise to the existing word or feature embeddings to incorporate positional information into the input representation fed into the neural network. This allows the model to learn positional dependencies and relationships, which is especially important in tasks like machine translation, text summarization, and image segmentation where the order or spatial arrangement of elements matters. By using sinusoidal positional embeddings, models can capture both the content and the position of elements in the input data, leading to improved performance on a wide range of sequence-based tasks.

More recently, positional embeddings have been used within diffusion models (Ho et al. 2020) to encode the time step, i.e. noise level the model is operating. Inspired by this usage of positional encoding, we employ this architecture to embed the noise level of considered MC simulations, i.e. the number of simulated particles.

**Deep Particle Embedding (DPE):** The model displayed in Figure 4.2 has the same architecture than the UNet used for the experiments of section 4.2. The main difference resides in the addition of the computation of sinusoidal positional embeddings of the number of particles simulated to compute the corresponding input dose distribution. Indeed, much like the assumptions in the works with diffusion models (Ho et al. 2020), embedding the number of particles boils down to encoding the corresponding MC noise level.

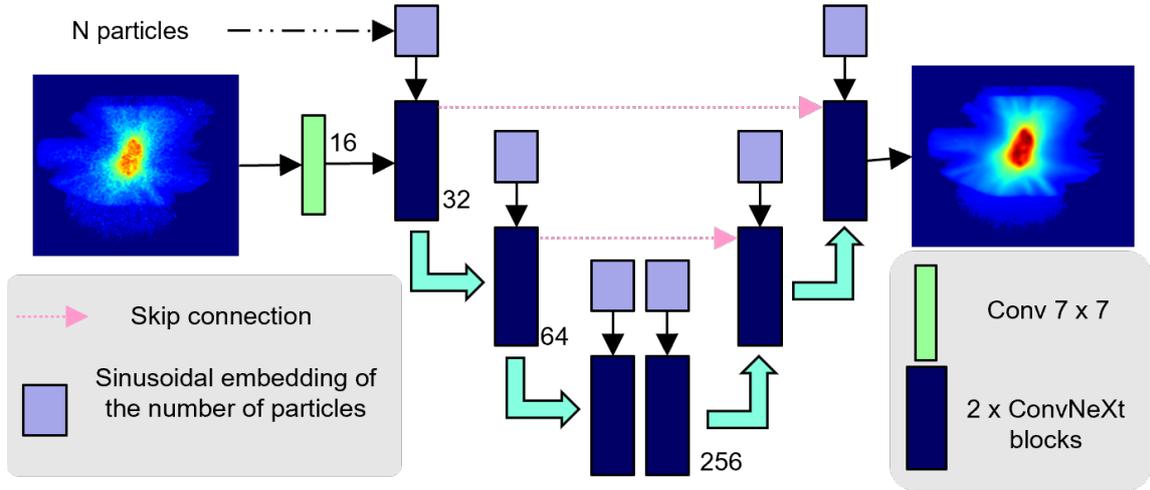


Figure 4.2: **Architecture of the Deep Particle Embedding Model and denoising pipeline.** The model is a UNet with integrated sinusoidal positional embeddings of the number of simulated particles associated with the input MC dose simulation.

Instead of considering pairs of decreasingly noisy simulations as in the approach of section 4.2, we train the model to directly predict the high-precision dose from any noisy MC simulation, i.e. any level of noise. We note that the main difference between this model and the standard UNet used as benchmark against ConvLSTMs in section 3.3 stems from the computation of the particle embeddings. Similarly, the model has 10 million trainable parameters.

During training, the model learns to infer the high-precision dose from any of the 3 associated low-precision distributions and associated positional embeddings of the corresponding numbers of simulated particles,  $10^8$ ,  $10^9$ ,  $10^{10}$ . At convergence of the training, we obtain a model that takes as input a single noisy MC simulation and its number of simulated particles to yield a high-precision dose map. Thus, at inference, the model can infer from any of these low-precision distributions independently. With this framework, we take full advantage of both the information residing in the input sequence and the reduced computational bottleneck implied by the one-to-one prediction.

**Optimization:** The training, validation and test sets draw slices of dose volumes from respectively 35, 5 and 10 non overlapping patients as described in section 4.2. We train the model on complete slices extracted from the patients' dose volumes. We enabled batch training by padding all slices to match the size  $256 \times 256$ . We apply random horizontal and vertical flipping as data augmentation during training. We use AdamW optimizer and the loss function is the

sum of the MSE and the SSIM. function. We stopped the training when necessary according to early stopping and decayed the learning rate by 10% when the validation loss plateaued for more than 200 iterations. At inference, no padding is required as the architecture is fully convolutional.

### 4.3.3 Results

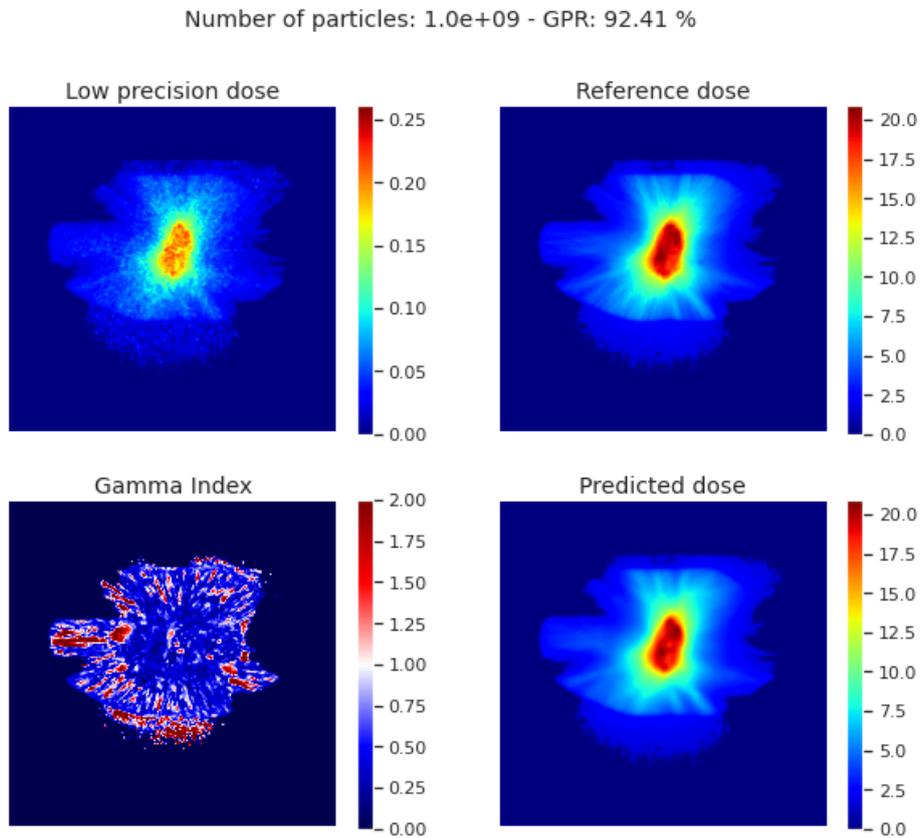


Figure 4.3: Denoising results - dose prediction and gamma index matrix - of a 1e9 simulation with the deep particle embedding model.

For benchmark, we trained the same model using only one low-precision simulation as input. The chosen simulation was the one computed with  $10^9$  particles as it is the intermediate MC simulation of the input sequence. We also compare with our 2D ConvLSTM model presented in section 3.2 (Martinot et al. 2021) and with our cascaded UNets from section 4.2. We evaluate the model using the GPR, the SSIM and the MAE. Table 4.2 show the results on the test set.

Input simulation	GPR (%)	SSIM	MAE
Cascaded UNets - 1e9 to 1e11	69.39 ± 8.08	N/A	N/A
DPE trained to denoise 1e9 alone	77.1 ± 17.9	90.1 ± 6.1	8.6e-1 ± 3.0e-2
DPE - 1e8	61.4 ± 6.6	74.3 ± 7.2	1.1e-1 ± 4.1e-2
DPE - 1e9	<b>83.2 ± 10.2</b>	91.1 ± 3.1	8.3e-2 ± 2.3e-2
DPE - 1e10	<b>95.7 ± 3.6</b>	94.5 ± 2.1	6.8e-2 ± 2.1e-2
<i>Martinot et al. 2021</i> (1e10)	83.2 ± 12.9	<b>97.3 ± 2.7</b>	N/A

Table 4.2: Evaluation metrics for the performance of the models on the test set.

First we notice that training the model on the single MC simulation with  $10^9$  particles yields GPR scores that are consistent with the performances of the finetuned and cascaded UNets. Secondly, denoising simulations computed with  $10^9$  particles with our deep particle embedding model trained to denoise any MC noise level, yields a relatively better GPR score, MAE value and even SSIM score, than the model trained on the sole task of denoising the simulations computed with  $10^9$  particles. This observation suggests that training the model to denoise various levels of noise leads to superior denoising capabilities. Finally, our deep particle embedding model outperforms our previous work with 2D ConvLSTMs (*Martinot et al. 2021*) on the task of denoising simulations computed with  $10^{10}$  particles by 10% on the GPR scores, setting a new record.

Figure 4.3 displays a low-precision dose map computed with  $10^9$  particles, the associated ground-truth simulation, the dose predicted by our deep particle embedding model and the gamma index matrix computed between the predicted dose map and the reference distribution. Qualitatively, the generated dose map seems to smooth details in dose deposition in both low dose and high dose areas. According to the look of the gamma index matrix, failing voxels appear significantly in low dose areas. Some voxels fail to satisfy the gamma index criterion ( $\gamma \leq 1$ ) in the high dose areas that correspond to regions with high gradients, i.e. in regions where the dose seems to vary a lot from one voxel to another.

## 4.4 Conclusion

In chapter 3, we had presented the capabilities of deep learning architectures based on the recurrent ConvLSTM modules on the task of accelerating MC simulations. However, the approach still experienced one critical flaw stemming from the computational overhead induced by the computation of the complete input sequence. In this chapter, we offered to improve on

the recurrent framework and to overcome this limitation. The goal was to bypass the computational requirement of computing all the simulations in the input sequence, at least at inference time.

First, we proposed to investigate cascaded UNets to progressively denoise each noisy MC simulation. Each UNet was trained on pairs of decreasingly noisy MC simulations. Cascading these models then resulted in a single model capable of denoising simulations computed with  $10^8$  particles into a high-precision dose distribution. Nonetheless, the quantitative evaluation of this model revealed that it suffered from error propagation, resulting in subpar performances.

Subsequently, we investigate another deep learning architecture that enables to discard any problem related to error propagation but still exploit the information contained in the input sequence while reducing the computational overhead at inference. We proposed to use positional embeddings to encode the number of simulated particles associated to each input MC simulation, i.e. to encode the noise level and give it as extra information to the model. At inference, any MC simulation is required to predict the high-precision dose map. The outcomes of training this model highlight that using both positional embeddings to encode the number of simulated particles, and training the model to denoise from any MC noise level allowed the model to extract more meaningful features and to outperform our previous frameworks. The results showcase a new efficient way to accelerate MC computations using neural networks while reducing the computational overhead.

# Chapter 5

## Deep Gamma index-based Loss functions

### Contents

---

<b>5.1</b>	<b>Introduction</b>	<b>94</b>
<b>5.2</b>	<b>Related work</b>	<b>95</b>
5.2.1	Loss functions	95
5.2.2	Gamma index	95
<b>5.3</b>	<b>Definitions</b>	<b>96</b>
5.3.1	The Gamma index	96
5.3.2	The Gamma index Passing Rate	96
<b>5.4</b>	<b>Methods</b>	<b>97</b>
5.4.1	Minimization problem	97
5.4.2	Soft counting with Sigmoid-GPR	97
5.4.3	Accelerating gamma index computations	99
<b>5.5</b>	<b>Experimental design</b>	<b>99</b>
<b>5.6</b>	<b>Results</b>	<b>102</b>
5.6.1	Training with GPR-based loss functions	102
5.6.2	Speed-up of GPR-acceleration approach	104
<b>5.7</b>	<b>Conclusion</b>	<b>106</b>

---

## 5.1 Introduction

In experiments presented in the previous chapter, we focused on proving the efficiency of adopting a many-to-one strategy instead of one-to-one to predict dose distributions with a quality approximating Monte-Carlo precision. Nevertheless, we trained all models involved to optimize loss functions based on weighted combinations of common computer vision metrics. Even if such methods provide a good trade-off between time and performance, training on such loss functions amounts to solving a proxy problem, with no strict assurance to conjointly optimize the clinical validation of the generated dose, which is performed using the GPR.

The GPR is one of the most essential and commonly used clinical evaluation metric for verification of complex radiotherapy dose delivery such as Intensity Modulated Radiation Therapy or Volumetric Modulated Arc Therapy (VMAT) [Quan et al. 2012](#). As such, the GPR provides a clinical criterion to assess the quality of the model’s predictions. Therefore, training directly with the GPR as primary objective would yield more accurate training from a clinical standpoint. However, the GPR has two main limitations that deter from using it as loss function. First, training neural networks in a supervised setting requires a differentiable loss function to allow backpropagation and to update the model’s weights. Yet, the GPR is non-differentiable thus jeopardizing gradient descent. Secondly, despite efforts to bridge the gap, current Gamma index and GPR computations remain time-consuming, especially when comparing high dimensional dose distributions.

We circumvent these challenges by incorporating the GPR as an optimization criterion during training of neural networks. We created a tensorized and differentiable approximation of the GPR that asymptotically tends towards the true GPR values. According to our knowledge, this is the first study to create a new class of loss functions based on the GPR and to bring the speed of gamma index computations down to milliseconds, both for 2D and 3D dose distributions. We provide a proof-of-concept showcasing deep learning acceleration of MC dose simulations with models trained to optimize the presented GPR-based loss functions. Finally, we study the behavior of the GPR-based loss functions and benchmark them against the SSIM, the MAE, the MSE and combinations of these.

## 5.2 Related work

### 5.2.1 Loss functions

When training a neural network on a task, the choice of the loss function is crucial. Loss functions such as the Dice Loss [Sudre et al. 2017](#), the Focal Loss [T.-Y. Lin et al. 2017](#), or the Structural Similarity Index Measure (SSIM) [Zhou Wang et al. 2004](#) have revolutionized, respectively, segmentation, object detection, and image processing tasks. Moreover, all loss functions do not yield the same impact on the training and inference, as explained in the study introducing the Multiscale-SSIM [H. Zhao et al. 2016](#).

This problem becomes even more evident in the medical field, in which models need to ensure reliable performance. For this reason, integrating mathematical objectives that train the models to optimize clinically relevant properties is of utmost importance for their integration into clinical practice.

In light of these considerations, we overcome the mathematical challenge of the GPR and turn this clinical metric into a viable loss function for our task of accelerating the simulation of MC radiotherapy dose distributions. We provide a family of GPR-based criteria that are therefore in adequacy with clinical requirements.

### 5.2.2 Gamma index

The main challenges of computing the gamma index matrix reside in the pixel-wise computation of gamma index values that can be time-consuming proportionally to the dimensionality of the evaluated dose distribution. Prohibitive calculation time hinders the potential of the GPR as loss function. Many works propose ways to decrease the computation complexity, either by changing the mathematical formalism or accelerating the calculations. In [Gu et al. 2011](#), Gu et al. use a geometric method with a GPU-accelerated radial pre-sorting technique to speed up calculations. [Chen et al. 2009](#) consider reducing the search distance by using a fast Euclidean distance transform.

In this paper, we present an acceleration approach adequate for deep learning frameworks that significantly reduces the calculation speed and enables fast training with our GPR-based loss functions.

## 5.3 Definitions

### 5.3.1 The Gamma index

Let  $D_r$  and  $D_e$  be two dose distributions ( $\mathbb{R}^k \rightarrow \mathbb{R}$ ), respectively the reference and the evaluated. In our case, the evaluated dose distribution is the model's prediction. To each of them corresponds a grid of points in which each point,  $P_r$  of  $D_r$ , and  $P_e$  of  $D_e$  has a coordinate vector, respectively  $\vec{d}(P_r)$  and  $\vec{d}(P_e)$ , and a dose value,  $D_r(P_r)$  and  $D_e(P_e)$ .

Let us consider a point  $P_r$  in  $D_r$  and the points  $P_e$  in a vicinity  $V(P_r)$  around  $P_r$ . Then the gamma index  $\Gamma$  is defined as a function of real values such that for all  $P_r \in D_r$ ,  $\Gamma(P_r)$  writes as follows:

$$\Gamma(P_r) = \min_{P_e \in V(P_r)} \sqrt{\frac{\|\vec{d}(P_e) - \vec{d}(P_r)\|^2}{DTA^2} + \frac{(D_e(P_e) - D_r(P_r))^2}{\Delta^2}} \quad (5.1)$$

where  $DTA$  is the tolerance on the Distance-To-Agreement (DTA), commonly in  $mm$ , and  $\Delta$  is the tolerance on the relative dose difference expressed as a percentage of the reference dose value  $D_r(P_r)$ . This definition entails that each point  $P_r$  has its own gamma index value in  $\Gamma$ , which indicates whether there exists an evaluated point  $P_e$  in its neighbourhood that is close enough to its dose value.

### 5.3.2 The Gamma index Passing Rate

Let us introduce a dose threshold  $\delta$  and consider a point  $P_r$  of the reference distribution such that  $D_r(P_r) \geq \delta$ . Then, given a DTA and dose tolerance  $\Delta$ , the evaluated distribution matches the reference at  $P_r$ , if the passing criterion is satisfied, i.e. if:

$$\text{Passing criterion: } \Gamma(P_r) \leq 1 \quad (5.2)$$

The Gamma index Passing Rate (GPR) is defined as the percentage of points  $P_r$  that satisfy the condition in eq. 5.2 while  $D_r(P_r) \geq \delta$ .

Let  $\mathbb{1}_{D_r \geq \delta}$  and  $\mathbb{1}_{\Gamma \leq 1}$  be the indicator functions defined such that:

$$\mathbb{1}_{D_r \geq \delta}(P_r) = \begin{cases} 1 & \text{if } D_r(P_r) \geq \delta. \\ 0 & \text{otherwise.} \end{cases} \quad \mathbb{1}_{\Gamma \leq 1}(P_r) = \begin{cases} 1 & \text{if } \Gamma(P_r) \leq 1. \\ 0 & \text{otherwise.} \end{cases} \quad (5.3)$$

Then we can write the GPR as follows:

$$GPR(D_r, D_e) = \frac{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r) \cdot \mathbb{1}_{\Gamma \leq 1}(P_r)}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \quad (5.4)$$

## 5.4 Methods

### 5.4.1 Minimization problem

With the GPR formulation in eq. 5.4, maximizing the GPR amounts to minimizing the corresponding loss function  $L_{GPR}^\delta$  which draws values in  $[0, 1]$ :

$$L_{GPR}^\delta(D_r, D_e) = 1 - \frac{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r) \cdot \mathbb{1}_{\Gamma \leq 1}(P_r)}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} = 1 - GPR \quad (5.5)$$

Due to the fact that the indicator function  $\mathbb{1}_{D_r \geq \delta}$  does not depend on  $\Gamma$ , the gradient of  $L_{GPR}^\delta$  (with respect to the trainable parameters) can be written as follows:

$$\frac{\partial L_{GPR}^\delta}{\partial w} = \frac{1}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \cdot \sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r) \frac{\partial \mathbb{1}_{\Gamma \leq 1}(P_r)}{\partial \Gamma} \frac{\partial \Gamma(P_r)}{\partial w}, \quad (5.6)$$

where  $w$  represents any of the trainable parameters of the neural network.

The problem with the above definition of  $L_{GPR}^\delta$  is that it generates zero gradients, which is a direct consequence of the fact that the indicator function  $\mathbb{1}_{\Gamma \leq 1}(\cdot)$  is stepwise constant with respect to  $\Gamma$ , therefore preventing training with gradient descent and parameter updates.

To address this issue, in the following we propose the use of a soft approximation of the objective function  $L_{GPR}^\delta$  with non-zero gradients.

### 5.4.2 Soft counting with Sigmoid-GPR

To avoid the propagation of null gradients, we propose to use the sigmoid function,  $\sigma(x) = (1 + \exp^{-\beta x})^{-1}$ , to approximate counting passing voxels. The slope of the sigmoid depends on the value of its sharpness  $\beta$  that we consider as a hyperparameter.

Moreover, we note that for all  $P_r \in D_r$ , it stands that:

$$\lim_{\beta \rightarrow +\infty} \sigma(\beta \cdot (1 - \Gamma(P_r))) = \mathbb{1}_{\Gamma \leq 1}(P_r) \quad (5.7)$$

Hence, the asymptotic behaviour of the sigmoid function combined with shifting the gamma

index values can provide an estimate of the count of passing voxels by summation over all points  $P_r$ . The accuracy of the estimation then depends on the value of  $\beta$ : the bigger the  $\beta$ , the more precise the estimation will be.

Thus, we approximate the loss  $L_{GPR}^\delta$  in eq. 5.5 with  $L_{\sigma-GPR}^\delta$  defined using the sigmoid function:

$$L_{\sigma-GPR}^\delta = 1 - \frac{\sum_{P_r \in D_r} \sigma(\beta \cdot (1 - \Gamma(P_r))) \mathbb{1}_{D_r \geq \delta}(P_r)}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \quad (5.8)$$

$L_{\sigma-GPR}^\delta$  is differentiable everywhere and, provided the sharpness  $\beta$  is not too high, gradients are non-zero and allow gradient descent to update the model's weights during backpropagation.

Given eq. 5.7, we remark that  $L_{\sigma-GPR}^\delta$  accurately approximates the true GPR loss function, i.e.,  $L_{\sigma-GPR}^\delta \rightarrow L_{GPR}^\delta$  as  $\beta \rightarrow +\infty$ .

#### 5.4.2.1 Annealing schedule of $\beta$ :

In light of the equations above, we propose to consider  $\beta$  as a hyperparameter. At the beginning of training, the model usually predicts poorly and the majority of voxels fail to satisfy the gamma index passing criterion. This implies that the corresponding loss computed with  $L_{\sigma-GPR}^\delta$  generates zero gradients everywhere if the value of  $\beta$  is set too high. To avoid this behaviour, we propose an annealing schedule for  $\beta$  that starts with low initial values and progressively increases  $\beta$  over the training. Moreover, when  $\beta \sim 0^+$ , the Taylor series expansion of the sigmoid function yields:

$$\sigma(\beta \cdot (1 - \Gamma(P_r))) \sim \beta \cdot (1 - \Gamma(P_r)) \quad (5.9)$$

Given eq. 5.9, we can write the Taylor expansion of  $L_{\sigma-GPR}^\delta$  when  $\beta \sim 0^+$ :

$$L_{\sigma-GPR}^\delta \sim 1 - \beta + \frac{\sum_{P_r \in D_r} \Gamma(P_r) \mathbb{1}_{D_r \geq \delta}(P_r)}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \quad (5.10)$$

Consequently, we introduce the loss function  $L_\Gamma^\delta(D_r, D_e)$  to model the linear behaviour of  $L_{\sigma-GPR}^\delta$  at the start of the annealing schedule:

$$L_\Gamma^\delta(D_r, D_e) = \frac{\sum_{P_r \in D_r} \Gamma(P_r) \cdot \mathbb{1}_{D_r \geq \delta}(P_r)}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \quad (5.11)$$

As the training continues and the loss decreases, the annealing scheme proceeds in progressively increasing  $\beta$  in order to improve the approximation of the GPR loss  $L_{GPR}^\delta$  defined in Eq. 5.5. As  $\beta$  increases and is acquiring larger values, minimizing the loss amounts to getting

failing voxels (voxels with  $\Gamma > 1$ ) to satisfy the passing criterion.

To model the behaviour of  $L_{\sigma-GPR}^{\delta}$  at this stage (i.e. as  $\beta$  is acquiring larger values), we modify  $L_{\Gamma}^{\delta}$  in eq. 5.11 by introducing the loss function  $L_{\Gamma>1}^{\delta}$ . To prevent backpropagation of zero gradients with respect to  $\mathbb{1}_{\Gamma>1}$ , we use the *stopgrad* operation:

$$L_{\Gamma>1}^{\delta}(D_r, D_e) = \frac{\sum_{P_r \in D_r} \Gamma(P_r) \cdot \mathbb{1}_{D_r \geq \delta}(P_r) \cdot \text{stopgrad}(\mathbb{1}_{\Gamma>1}(P_r))}{\sum_{P_r \in D_r} \mathbb{1}_{D_r \geq \delta}(P_r)} \quad (5.12)$$

For the sake of characterizing the behaviour of  $L_{\sigma-GPR}^{\delta}$  in eq. 5.8, we also study trainings that involve the use of loss functions  $L_{\Gamma}^{\delta}$  and  $L_{\Gamma>1}^{\delta}$  in the following experiments.

### 5.4.3 Accelerating gamma index computations

Having a differentiable GPR loss does not make it directly applicable for neural network training since, by definition, it requires iterating over all voxels in the given distributions, therefore leading to prohibitive computation time when considering high-resolution distributions. To deal with this issue, we propose an accelerated version of GPR for faster calculations.

To avoid physical incoherence when computing gamma index values, we sample the evaluated and reference distributions to the resolution  $1 \text{ mm}^3$  with bilinear interpolation. By the definition in eq. 5.1, one can observe that the evaluated voxels located farther than  $DTA \text{ mm}$  from  $P_r$  automatically yield a gamma index superior to 1. Thus we limit the search within an invariant vicinity defined by the chosen  $DTA$ . More precisely, the gamma index value of a reference point then stems from comparing gamma values computed with voxels in a cube comprising  $(2 \times DTA + 1)^k$  voxels, in the case of  $k$  dimensional dose distributions. We then use unfolding to extract sliding local blocks of the evaluated distribution generated by the model. This operation creates one channel per voxel in the vicinity defined by the  $DTA$ . We then apply the minimum operation over the channel dimension to get the minimal gamma index value.

The approach enables fast computation of the gamma index distribution  $\Gamma$ , which is necessary for the calculations of  $L_{\sigma-GPR}^{\delta}$ ,  $L_{\Gamma}^{\delta}$  and  $L_{\Gamma>1}^{\delta}$  presented in eq. 5.8, 5.11, 5.12. Computation times are discussed in Section 5.6.

## 5.5 Experimental design

**Dataset & Preprocessing:** We carried out the experiments on the publicly available dataset presented in [Sonia Martinot et al. 2021](#) comprising 50 patients treated with VMAT plans. Each patient has a reference dose distribution computed from  $1 \times 10^{11}$  particles and a low precision

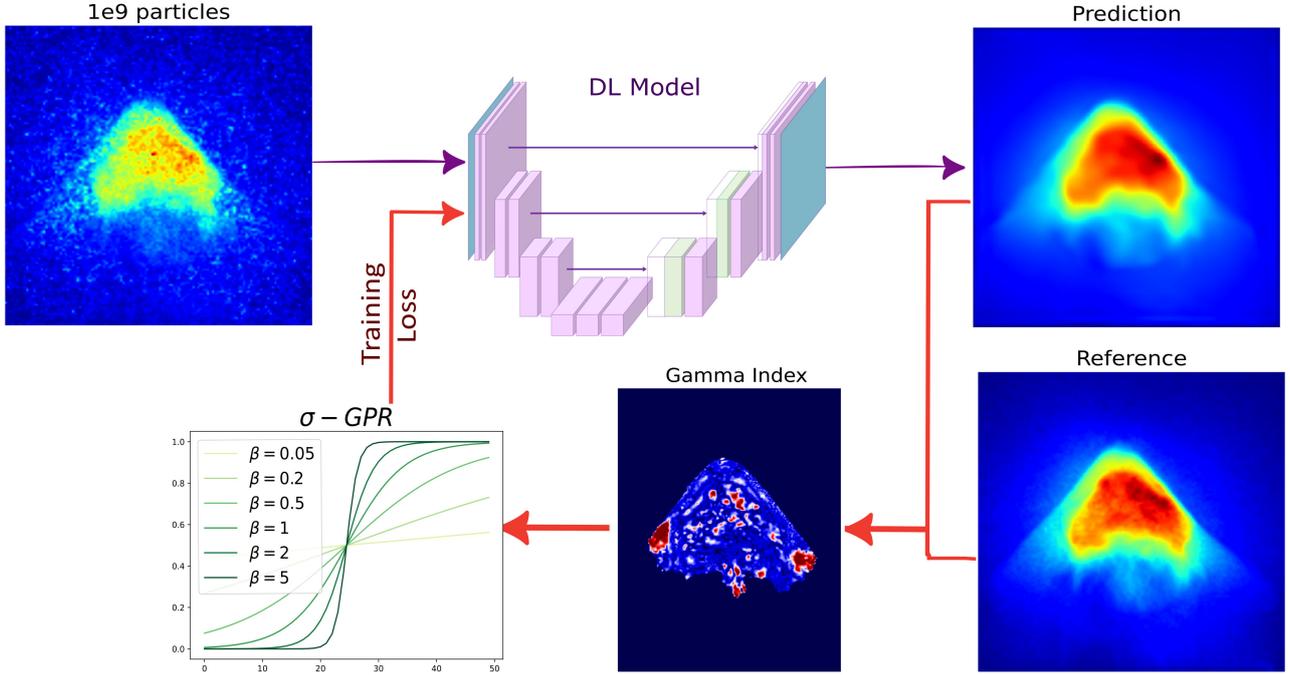


Figure 5.1: **Pipeline of training using sigmoid-GPR-based loss functions.** The Gamm-index based loss function require to compute the gamma index matrix between the dose predicted by the model and the reference distribution. Then, a soft GPR is computed using a sigmoid function. The sharpness parameter  $\beta$  of this sigmoid functions plays as a hyperparameter that can be tuned to take various values depending on the advancement of the training.

simulation computed from  $1 \times 10^9$  particles. The main goal of the methods benchmarked on this task is to generate the high precision simulation of the dose from the available low precision one. More details about the dataset can be found in the original publication. For our experiments, we split patients to 35-5-10 for respectively, the train, validation and test sets. The cases in the dataset correspond to various anatomies and therefore, we split them as equally as possible between sets to avoid biases.

Even though our approach enables training on 3D dose distributions, the dataset comprises a small number of samples. Thus, we decided to carry out the experiments in 2D to favour significant experiments and a relevant benchmark. In this setting, a training sample corresponds to an axial slice of a patient’s dose volume. The 2D training dataset therefore comprises around 11k training samples, where a sample is a pair of corresponding slices of low precision and high precision dose simulation.

We normalized both low precision and reference distributions using the average dose maximum computed over the reference dose volumes from the training set. We then applied the

same normalization on the validation and test sets. To enable batch training, we padded each training sample with zeros in order to match a fixed size of  $256 \times 256$ . To further help the model generate accurate dose predictions, we added the corresponding CT slice as second input channel to incorporate the corresponding anatomy. We applied minimax normalization to CT volumes so voxel values remain in  $[0, 1]$  range.

**Model:** In all experiments, the model is a standard UNet architecture [Ronneberger et al. 2015b](#) with skip connections between the encoder and the decoder. The encoder part of the model performs downsampling twice with convolutional layers using  $4 \times 4$  filters and a stride of 2. Symmetrically, transposed convolutions upsample feature maps in the decoder. Each stage of the UNet comprises two convolution blocks before downsampling or upsampling. Much like the convolutional block presented by Liu et al. in [Z. Liu et al. 2022](#), a convolution block first applies a convolution with  $7 \times 7$  filters and  $3 \times 3$  padding, and then two convolutions with  $3 \times 3$  filters to further process the features maps. Each convolution is followed by Gaussian Error Linear Units (GELU) activation units. The block ends with a residual connection to keep high frequency details from the block’s input. Overall, the model has around 10 million trainable parameters.

**Optimization set-up:** In all trainings, we trained the model using AdamW optimizer [Loshchilov et al. 2017b](#). We set the initial learning rate to  $3e^{-4}$  and decreased it progressively during the training when the validation loss stagnated. Weight decay was set to  $5e^{-4}$  and batch size to 16. We trained for 20k iterations on a NVidia GeForce RTX 3090 GPU. The trainings were stopped when overfitting appeared by adopting the early stopping strategy. With this training scheme, early stopping occurred after around 15k iterations, when the validation loss fails to decrease 2% after 500 iterations.

**Loss functions:** To train with the GPR-based loss using sigmoid count  $L_{\sigma-GPR}^{\delta}$  presented in eq. 5.8, we designed the following annealing schedule for the sharpness parameter  $\beta$ . We set the initial value of  $\beta$  to  $2 \times 10^{-2}$  for the first 150 iterations. Then,  $\beta$  increased by a factor of 5% every 50 iterations until it reached an intermediate value of 3 where updates slowed down to 5% every 100 iterations. Increasing updates stopped when  $\beta$  reached a chosen ceiling value of  $\beta_{max} = 5$ . Setting  $\beta_{max}$  prevented the slope of the sigmoid from getting too sharp and the loss from encountering a vanishing gradient problem, which would stop the updates of gradient descent. Additional benchmarks with the approximating loss functions  $L_{\Gamma}^{\delta}$  and  $L_{\Gamma>1}^{\delta}$  have been also conducted, in order to better characterize the behaviour of  $L_{\sigma-GPR}^{\delta}$  for small values of  $\beta$ .

For all GPR-based functions, we set the dose threshold  $\delta$  to 20%. This means that, while loss functions compute the gamma index distribution by considering all voxels, the computed approximated GPR value takes into account only voxels  $P_r$  for which the dose value is superior to 20% of the maximum dose of the reference distribution, i.e.  $D_r(P_r) \geq 20\% \cdot \max_{P_r \in D_r} D_r(P_r)$ .

To benchmark against our proposed GPR-based loss functions, we considered several other loss functions commonly used in computer vision. The benchmark includes the MAE and the MSE for a comparison with pixel-wise errors. Finally, we considered the combination of pixel-wise errors with the SSIM. More precisely, the benchmark includes SSIM-MAE and SSIM-MSE, which are the equally weighted sum of respectively the SSIM and MAE, and the SSIM and MSE. For each training on the loss functions considered above, we used the exact same model architecture and optimization strategy, in order to promote the reliability and fairness of the comparison.

## 5.6 Results

### 5.6.1 Training with GPR-based loss functions

Extensive quantitative comparison on the test set for each training, using the MAE, MSE, SSIM and GPR with various values of DTA and dose tolerance  $\delta$  are summarised in Table 5.1. As the test set comprises 10 patients, we computed the metrics over each slice of each patient’s volume, and then average over the test set per considered metric. Results point out that models trained with GPR-based loss functions tend to outperform others with respect to the GPR, the MAE and MSE. In contrast, models trained with SSIM-MAE and SSIM-MSE show the highest SSIM scores. With a closer look however, one can observe that they report among the lowest performance for the rest of the metrics. This result indicates that the SSIM may not be a well-suited metric to evaluate the quality of dose distributions, since it seems to be biased.

To assert statistical significance of the results, we take an in-depth look at each patient in the test set to explain the high standard deviation values observed in Table 5.1. Boxplots a), b) and d) in Figure 5.2 point out the presence of an outlier patient case on which models tend to fail with respect to the GPR, SSIM and MSE. In contrast with SSIM, MSE and MAE-trained models, we observe that models trained with GPR-based loss functions not only display robustness to this outlier, but also show smaller standard deviation over the whole test set.

Figure 5.2 also allows to compare discrepancies within the family of GPR-based loss func-

Loss function	GPR 3%/3mm	SSIM (%)	MAE	MSE
MAE	66.5 ± 10.0	83.3 ± 9.2	0.39 ± 0.10	0.51 ± 0.28
MSE	69.5 ± 10.3	83.9 ± 11.5	0.34 ± 0.09	0.43 ± 0.29
SSIM+MAE	64.6 ± 11.9	<b>94.0 ± 8.8</b>	0.35 ± 0.09	0.74 ± 0.36
SSIM+MSE	67.5 ± 13.2	93.3 ± 3.1	0.30 ± 0.07	0.30 ± 0.15
$L_{\Gamma}^{\delta}$	79.2 ± 4.3	87.5 ± 7.6	0.28 ± 0.07	0.30 ± 0.23
$L_{\Gamma>1}^{\delta}$	79.0 ± 4.4	86.7 ± 8.1	0.27 ± 0.07	0.25 ± 0.16
$L_{\sigma-GPR}^{\delta}$	<b>81.4 ± 4.0</b>	88.2 ± 7.5	<b>0.24 ± 0.06</b>	<b>0.22 ± 0.16</b>

Table 5.1: Evaluation metrics over the dose distributions comprised in the test set. Different benchmarks over the considered loss functions for different metrics are highlighted with their mean and standard deviation. With bold we indicate the best performing methods per metric.

Metric	3D dose Time(ms)	2D dose Time(ms)
MSE	0.14 ± 0.02	0.13 ± 0.03
MAE	0.23 ± 0.06	0.19 ± 0.04
SSIM	27.49 ± 7.73	5.06 ± 3.32
$L_{\sigma-GPR}^{\delta}, L_{\Gamma>1}^{\delta}, L_{\Gamma}^{\delta}$	<b>30.54 ± 0.01</b>	<b>4.51 ± 0.00</b>
Exhaustive	985 ± 515	8.01 ± 2.28
PyMedPhys	> 1 second	> 100 ms

Table 5.2: Speed comparison of metrics computed over 2D and 3D dose distributions.

tions. While all of them produce better performing models with respect to all evaluation metrics except the SSIM, the loss function with sigmoid counting  $L_{\sigma-GPR}^{\delta}$  outperforms  $L_{\Gamma}^{\delta}$  and  $L_{\Gamma>1}^{\delta}$ . We explain this behaviour by the fact that both  $L_{\Gamma}^{\delta}$  and  $L_{\Gamma>1}^{\delta}$  focus only on minimizing gamma index values, and not directly maximizing the number of voxels satisfying the passing criterion. We conclude that  $L_{\sigma-GPR}^{\delta}$  yields better maximization of the GPR and is therefore the better approximation of the true GPR loss function  $L_{GPR}^{\delta}$ .

The MSE-trained model outperforms other models trained with non GPR-based loss functions with respect to the GPR, so we chose to display its dose prediction conjointly with the dose generated by the  $L_{\sigma-GPR}^{\delta}$  trained model in Figure 5.3. Although both trainings achieved convergence, the prediction of the MSE-trained model manifests important artefacts at the bottom of the generated dose. Additionally, the dose itself seems to be smoother than the dose predicted with the  $L_{\sigma-GPR}^{\delta}$  training. Finally, the MSE-trained model appears to overestimate the dose in low-dose regions to a greater extent than the  $L_{\sigma-GPR}^{\delta}$ -trained model.

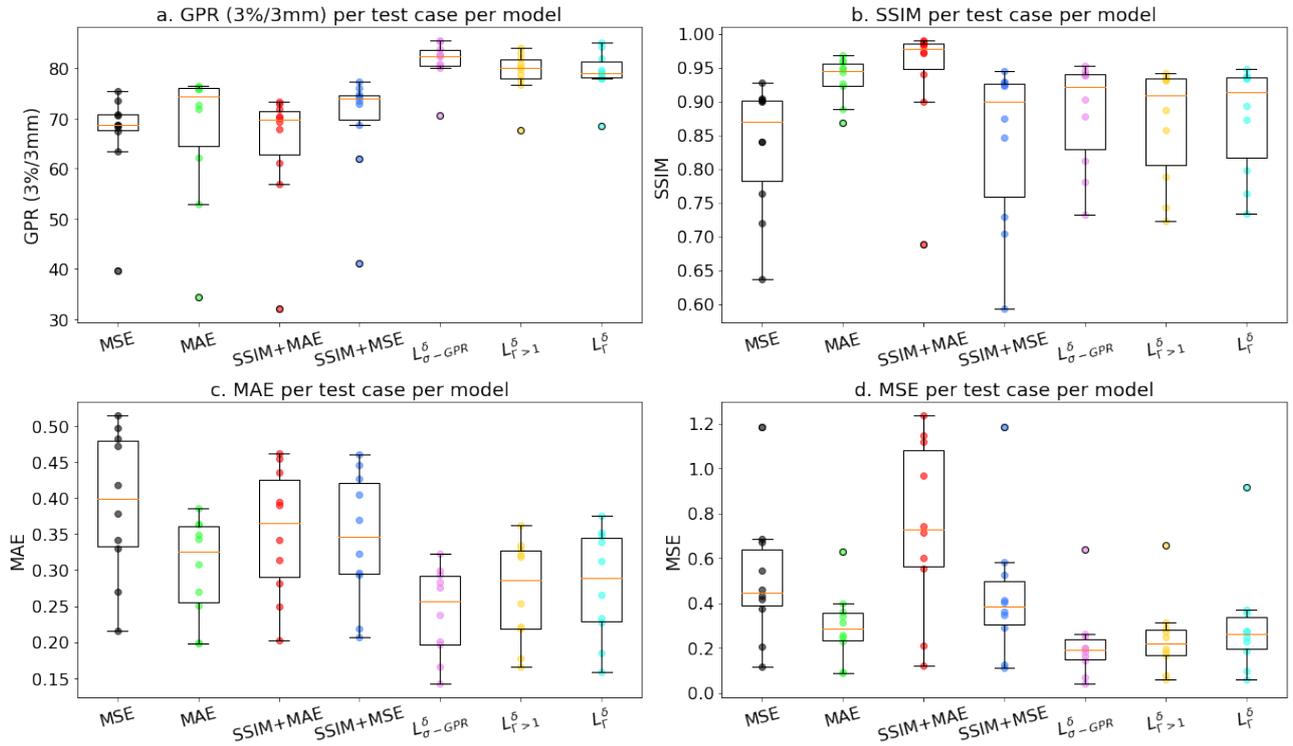


Figure 5.2: **Boxplots representing the evaluation metrics achieved by trained models for each case in the test set depending on the loss function used for training.** The y axis indicates the values of the considered metric. The x axis specifies the loss function with which the corresponding model was trained.

### 5.6.2 Speed-up of GPR-acceleration approach

In an effort to promote the GPR-based loss functions as viable deep learning optimization criteria that allow fast error computations and training, we had to accelerate gamma index computations. To quantify the extent of our acceleration approach, we benchmark against two methods. The first one is a GPU-accelerated exhaustive search approach in a limited vicinity of  $3mm^3$  around the considered reference voxel. The second is an open-source tool from PyMedPhys [Biggs et al. 2022](#) which makes use of acceleration ideas from [Wendling et al. 2007](#) and executes on CPU and is single-threaded. Regarding the latter, we limit the interpolation ratio to 2 to have a fair comparison.

The time estimation was twofold. We timed each evaluation metric and GPR-based loss functions on 3D or 2D distributions stemming from the MC dataset used for the experiments. 3D dose distributions were interpolated to resolution  $1\text{ mm}^3$  and of shape  $128 \times 200 \times 200$ , comprising around  $5 \times 10^6$  voxels. The 2D dose distributions comprised axial slices of the 3D

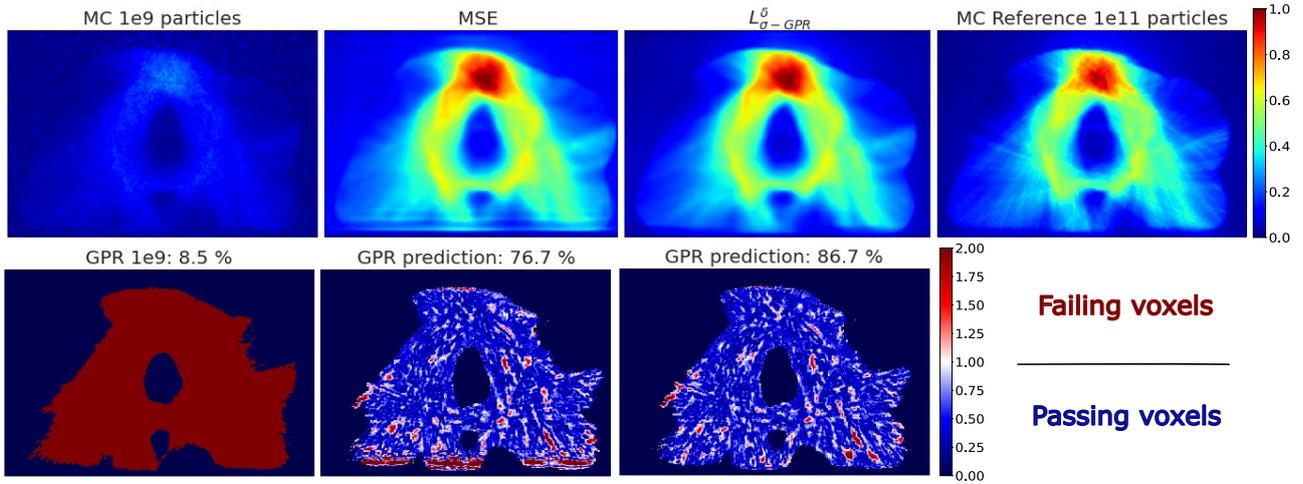


Figure 5.3: **Qualitative comparison of the dose distributions and corresponding gamma index matrices depending on the loss function chosen during training.** First row from left to right: a single slice of the  $1e9$  dose volume, predictions of models trained with MSE and  $L_{\sigma}^{\delta}-GPR$ , and reference  $1e11$  dose. Second row: gamma index maps for the three different representations.

dose distributions and were interpolated to a size  $400 \times 400$ . For the GPR calculations, we set the DTA and  $\Delta$  to respectively  $2mm$  and  $2\%$ . Execution times are displayed in Table 2.

Figure 5.4 and Table 5.2 highlight that our approach has equivalent speed to that of the SSIM. Compared to the exhaustive search method, our approach improves the speed of gamma index computations by a factor of at least 30 in the case of 3D dose distributions and twofold for 2D distributions. Consistently with these results, we note that trainings took around 24 hours for SSIM-MAE, SSIM-MSE and GPR-based loss function, whereas they lasted for 15 hours for experiments with the MAE and the MSE. Results therefore validate our GPR-based loss functions to efficiently train deep neural networks. Nonetheless, our comparison is limited to speed assessment and does not encompass RAM usage and precision considerations. Although our approach highlights significant speed gain in the computation of the GPR metric and, by extension, of the GPR-based loss functions presented in this study, it comes at the price of an increased RAM usage caused by the unfolding operation.

We make the remark that for all loss functions, the obtained GPRs do not meet the 95% GPR threshold indicating clinical validation. Nevertheless, the goal of the experiments was to show the benefits of optimizing directly the clinical metric during training and results support that statement.

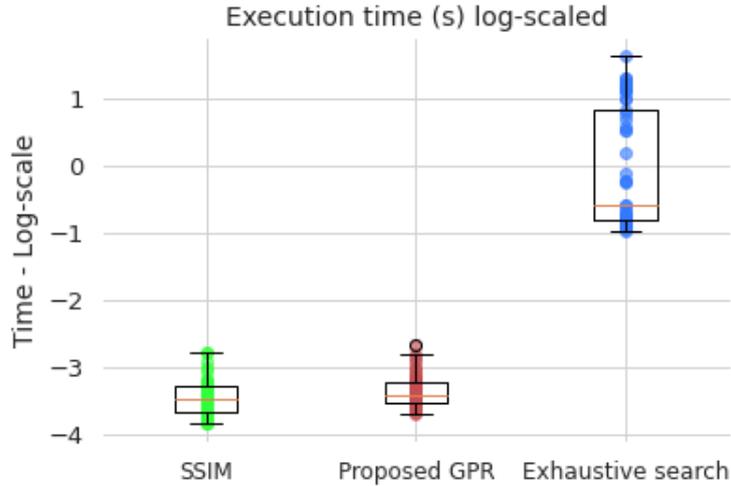


Figure 5.4: **Boxplots of execution times of our proposed approach, of the SSIM and the exhaustive search method on 3D dose distributions.** Our implementation of the gamma index is faster than the open-source exhaustive search method. Its speed is comparable to the execution time of the SSIM, thus ensuring tractable trainings of deep learning models.

## 5.7 Conclusion

Adopting the correct optimization criterion is essential to train deep learning models adequately with the task they are designed to solve. For the task of accelerating MC radiotherapy dose simulation with deep learning, this work proves that directly optimizing models with the clinical validation metric yields significant improvement in predicted dose quality when compared to other loss functions. We provide a fast computation of the GPR to enable such results. Moreover, the GPR is a similarity metric for distributions in general, and may be applied to other tasks such as radiotherapy dose generation or even finding adversarial examples for generative adversarial networks. Future work will focus on addressing the remaining limitations of our approach and assessing the potential of our new class of loss functions in solving other deep learning tasks.

# Chapter 6

## Discussion and Future work

### Contents

---

<b>6.1</b>	<b>Main findings . . . . .</b>	<b>107</b>
6.1.1	Denoising Monte-Carlo . . . . .	108
6.1.2	Clinical oriented training . . . . .	109
<b>6.2</b>	<b>Limitations and future work . . . . .</b>	<b>110</b>
6.2.1	Available datasets for data-driven methods . . . . .	110
6.2.2	Data-driven models . . . . .	111
6.2.3	Clinical oriented priors . . . . .	111

---

### 6.1 Main findings

Simulating the radiation dose is a crucial step to ensure the safety and precision of radiotherapy treatment plans. The MC method, while outperforming other dose simulation methods in terms of precision, is hindered by its prohibitive computation times. AI-based methods offer promising prospects on accelerating MC dose simulations. Moreover, the simulation of radiotherapy dose from VMAT plans is particularly challenging due to the continuous movement of the gantry around the patient and the motion of the MLC leaves. VMAT plans result in more precise and homogeneous dose distributions, but this comes at the cost of increased complexity in the dose simulation process. In the case of MC simulations, this added complexity necessitates the simulation of a larger number of particles, leading to longer computation times. In this thesis, we investigated deep learning-based methods to expedite the simulation of dose distributions

using the gold-standard MC method, specifically in cases involving patients undergoing complex VMAT plans.

### 6.1.1 Denoising Monte-Carlo

As developed in **Chapter 1**, several works try to address MC acceleration with AI-based methods in the case of patients undergoing IMRT or proton therapy. Performances are promising but confined to a small dataset comprising patients receiving treatment for similar anatomies, usually prostate. Moreover, the approach to denoising in these works remains under the one-to-one prediction paradigm. In contrast, we proposed a new approach that falls under the many-to-one setting. As explained in section 1.6.1, we set out to address the following question: **Does training deep learning models on sequences of Monte-Carlo dose simulations lead to more accurate dose distributions?**

To be able to train deep learning models on the task, we first curated a dataset of radiotherapy patients undergoing VMAT plans and receiving radiation in various regions of the body. We provided a detailed account of the patients' specificities, modelization choices and simulation strategy in **Chapter 2**. We described this patient cohort in section 2.1. For each patient we performed MC computations to simulate the associated radiation dose using the machine parameters defined in their treatment plan. Despite adopting several approximations, generating this dataset required the computational power of a supercomputer and underlined the fact that the MC method suffers greatly from heavy computational bottlenecks, as stressed in section 2.2.3.

In **Chapter 3**, we investigated ConvLSTM to set a benchmark for recurrent architectures. We explored its advantages and drawbacks both in 2D and 3D and compared it with several other deep learning architectures. ConvLSTMs show better GPR scores and global robustness than other models while having the lowest complexity. Moreover, we discovered that adding the CT as physical prior did not lead to improved performances of our models. Nonetheless, ablation studies conducted in section 3.4 suggested that our dataset may be too small to properly compare models in 3D. Additionally, our recurrent architectures still required the computational burden of simulating the input sequence of MC dose simulations at inference.

In **Chapter 4**, we set out to address these limitations. First, we conducted our experiments in 2D to augment the potential of our dataset. The first deep learning framework we studied in section 4.2 did not mandate complete sequences of MC simulations to perform prediction at inference. More precisely, we investigated progressive refinement conveyed by cascaded UNets.

Regrettably, the outcomes demonstrated suboptimal performance with susceptibility to error propagation. The subpar results of the cascaded UNets stemmed from aligning models that were initially trained separately, even when end-to-end finetuning was performed afterwards. To mitigate this issue, we then proposed an approach that circumvented this limitation.

To enforce exploitation of the information contained in the MC sequences, we integrated sinusoidal positional embeddings to our UNet architecture as a means to encode the number of simulated particles required for each simulation in section 4.3. This approach based on deep particles embeddings boiled down to associating each low-precision MC simulation with an encoding of its intrinsic noise level, thus encouraging the model to explicitly categorize simulation discrepancies between different noise levels. Results highlighted several key points. First, training this model on various noise levels achieved improved GPR results compared with training the same model on a single noise level. Secondly, our deep particle embedding model outperformed previous experiments from section 3.2 conducted with 2D ConvLSTMs. Indeed, our model scored GPR values reaching the 95% clinically acceptable threshold with an improved robustness on the test set. These outcomes supported our hypothesis that the additional context brought during training by full MC sequences bore significant information on how the radiation is deposited. This approach built upon previous models as it required a single low-precision MC simulation at inference, therefore circumventing the computational bottleneck of computing the complete MC dose sequence.

### 6.1.2 Clinical oriented training

In **Chapter 5**, we tackled another aspect of the optimization of deep learning models on the task of denoising MC simulations. Carefully choosing the loss function used to optimize the weights of deep learning models often leads to improved results. Indeed, an ill-suited loss function can result in suboptimal performances as the model is trained to optimize an inaccurate target and thus to solve a proxy problem. In our case, the GPR is the metric we systematically sought to maximize across all presented experiments to ensure clinical viability of the model’s predictions. Therefore, we wondered: **Can we include the clinical validation goal directly into the target loss function to optimize the model’s parameters? and does this improve the quality of the predictions?**

To address this question, we introduced a new loss function directly derived from the GPR. The GPR is inherently non-differentiable, making it unsuitable for gradient descent in deep learning training. To circumvent this limitation, we proposed an approximation of the metric

that asymptotically converges to the actual GPR values. We implemented the gamma index matrix as tensors to ensure efficient computation of loss values during training. We conducted a performance comparison of a simple UNet model depending on the loss function used during training. The results were exceptional, highlighting the superior performance achieved when optimizing our GPR-based loss function. As a result, our contribution paves the way for robust deep learning training applicable to a wide range of tasks, from MC denoising to dose generation in clinical settings.

## 6.2 Limitations and future work

### 6.2.1 Available datasets for data-driven methods

**Dataset size:** The main limitation that transpires across presented 3D experiments boils down to the limited number of samples available in the dataset. Indeed, the rather small size of the dataset remains the major limit to quantifying the generalization power of the models as well as impeding experiments with more sophisticated and complex neural networks. This limitation was best pinpointed during the ablation studies in section 3.4.

Nonetheless, we carried out experiments in 2D to artificially augment our dataset. Even though each 2D slice of a patient’s dose volume was regarded as an independent sample, we cannot deny that such samples still share evident similarities, both spatially and regarding dose deposition. Therefore, we believe that, although promising, our results will bear a stronger statistical significance when confirmed by applying our methods on a bigger dataset, especially for 3D frameworks.

Future work will therefore focus on curating a more substantial cohort of real-world VMAT patients and corresponding MC dose simulations. Additionally, we argue that the use of transfer learning (Torrey et al. 2010) would expedite our application as it greatly benefited other tasks in computer vision (Brodzicki et al. 2020) and medical imaging (Raghu et al. 2019). Thus, we would also generate an extensive dataset with a more general purpose of performing transfer learning for various MC-based simulations. This dataset would comprise MC simulations of photons’ interactions in random media, without necessarily any medical priors.

**Noise level:** We recall that computing a dose distribution by simulating the behaviour of  $10^8$  particles required 4 CPUs and approximately 2 hours of computation time. While our presented approaches reduce the computational overhead required to compute high-precision simulations,

the highest noise level of the MC simulations given as input to our models remains too high to reach real-time dose simulations. Nonetheless, our work on deep particles embeddings offers promising performances. Future work will therefore focus on pursuing this improvement and reducing further the computational bottleneck by training models on noisier inputs. On another level, we would like to investigate the fact that adding the CT as physical prior does not lead to enhanced results. We hypothesize that the sequence of MC simulations might contain all relevant information the CT brings, such as the absorption capacity of different organs and media.

### 6.2.2 Data-driven models

Presented experiments seem to support our hypothesis that teaching a deep learning model to predict the precise dosed distribution from a sequence of MC simulation yields better learning. Thus, we would like to investigate other well-suited architectures that would benefit from the sequential context, such as Transformers (Vaswani et al. 2017). Also, the performances of the deep particles embeddings model call for studying more closely the inner workings of the model. Current work involves exploring the latent space of the model to better understand how encoding the noise level helps the model better partition low-precision MC simulations.

Moreover, future work will consist in generalizing our architectures and experiments to the three-dimensional setting. Indeed, it is our belief that 3D architectures would better capture the spatial dependencies and specificities of dose deposition witnessed in full 3D dose distributions. Nonetheless, such experiments will undoubtedly call for a substantial dataset.

### 6.2.3 Clinical oriented priors

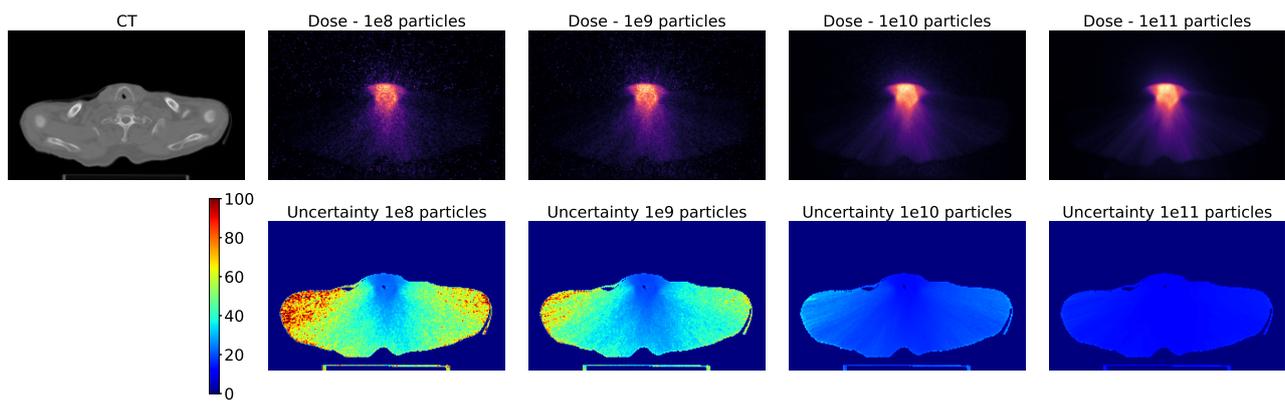
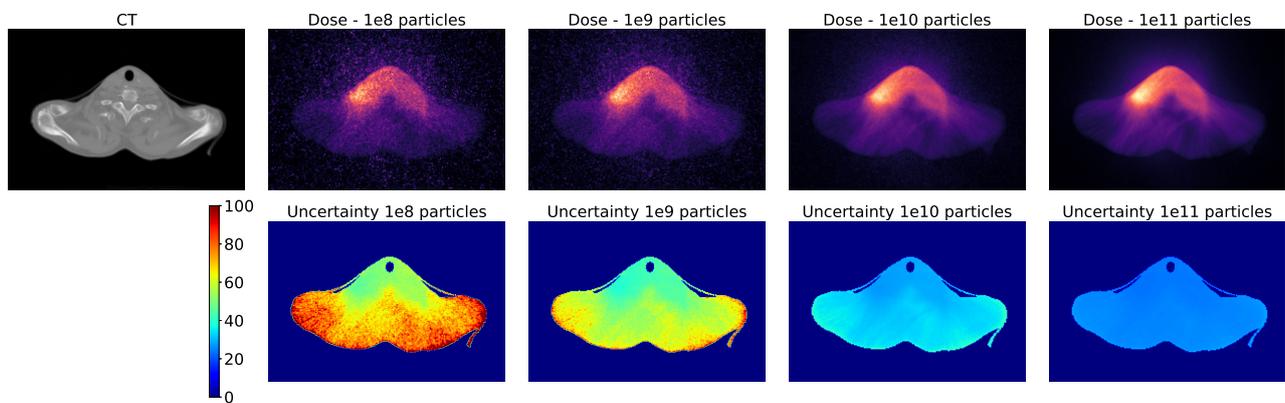
The main limitation of the gamma index-based loss function stems from its implementation. Although vectorized, computing the loss demands a relatively large GPU capacity. This trade-off is emphasized when working with 3D dose distributions. Indeed, computing the gradients of this loss functions during backpropagation brings the GPU memory consumption to a prohibitive level. This limitation entails that batches of samples be confined to small sizes and prohibits whole-image trainings. Current work focuses on reducing the computational tradeoff.

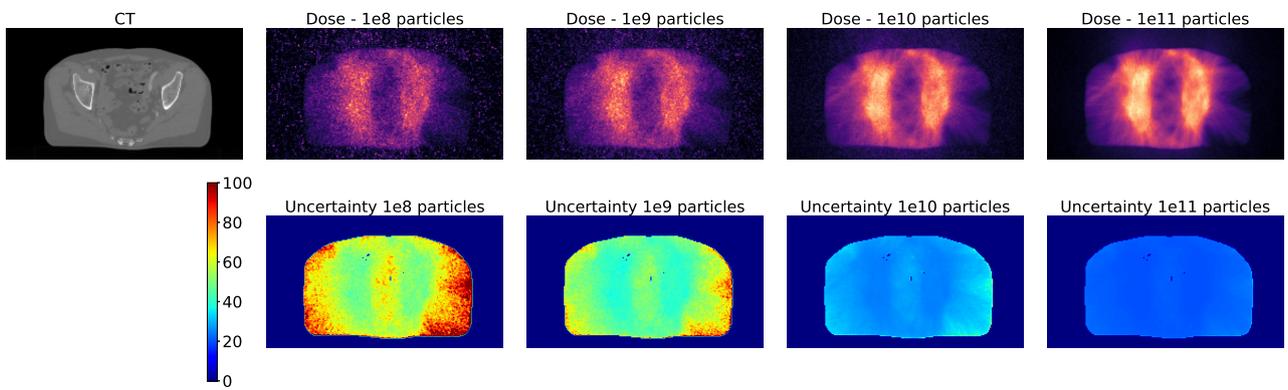
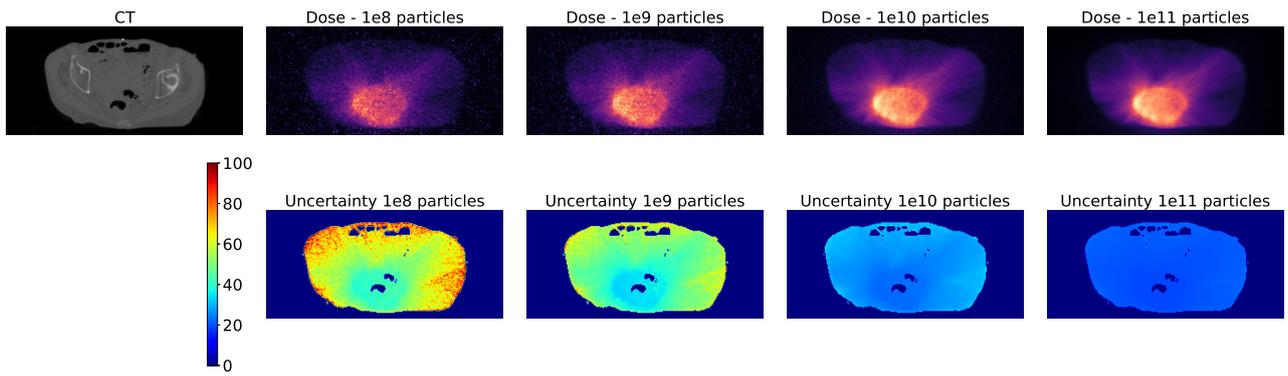
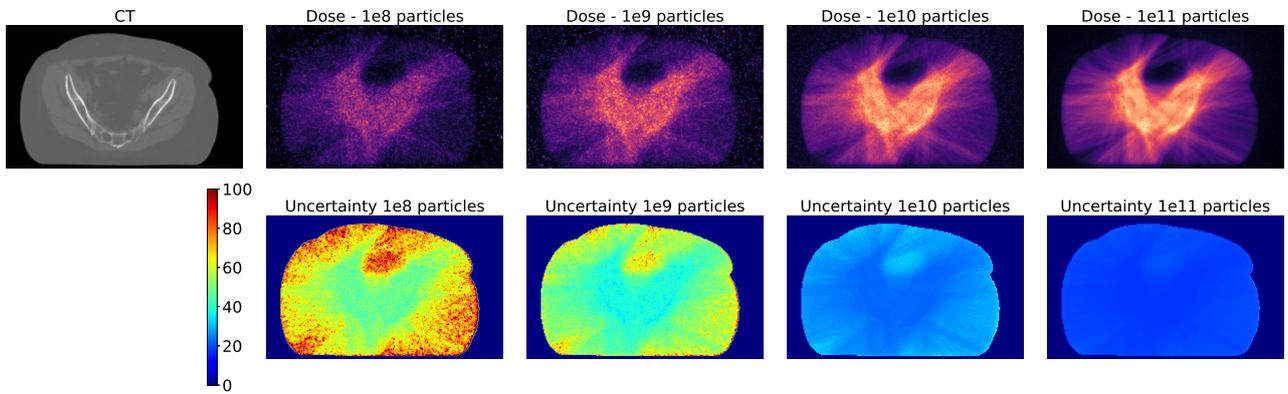
Future work with our gamma-index based loss function will also include evaluating its impact on 3D deep learning training. Moreover, our method opens new avenues for accelerating Monte-Carlo dose computation and with a broader perspective, for the dose generation field and other imaging modalities. Our new family of loss functions offers a more robust way to train neural

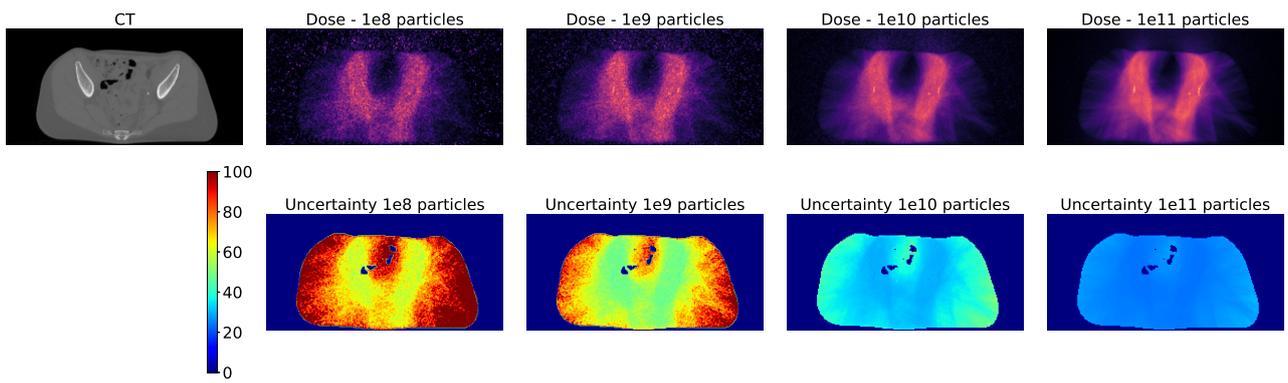
networks as optimization targets directly the quality assessment metric.

# Annexe

Here are some other examples similar to Figure 2.5 from Chapter 2 of simulations and corresponding MC uncertainty of other patients from the cohort:









# Remerciements (Acknowledgements)

Tout d'abord, je tiens à exprimer ma reconnaissance envers mon directeur de thèse, le Prof. Nikos Paragios, pour m'avoir fait confiance et accueillie au sein de Therapanacea. Sa précieuse expertise académique, ses directions scientifiques inspirantes et son assistance ont été déterminantes dans les moments importants de cette aventure.

Mes sincères remerciements vont également à tous les membres du jury, le Prof. Daniella Thornwarth, le Dr. David Sarrut, le Prof. Vincent Lepetit, pour leurs commentaires constructifs et le temps consacré à évaluer mon travail. Je suis honorée d'avoir pu présenter et discuter mes travaux en leur compagnie.

Je remercie le Dr. Charlotte Robert, le Prof. Eric Deutsch, le Dr. Norbert Bus et Christophe R. pour la mise en place de la collaboration scientifique entre l'Institut Gustave Roussy et Therapanacea, ainsi que pour la vérification de la qualité des données qui ont servi tout au long de mes travaux. Leurs conseils et contributions ont été déterminants dès le début de ma thèse, à l'ère du covid et des confinements sans fin. Merci d'avoir été toujours disponibles pour échanger et discuter des nombreux axes de recherche ouverts par mon sujet de thèse, notamment en termes d'applications cliniques réalistes et nécessaires.

Je remercie également le Prof. Nikos Komodakis avec qui les collaborations ont été stimulantes, toujours faciles et véritablement fructueuses pour mon domaine de recherche. Nos échanges m'ont beaucoup appris et ont renforcé ma passion pour le domaine fascinant de l'intelligence artificielle et pour la recherche en général.

Je souhaite exprimer ma gratitude envers mon encadrante, le Dr. Maria Vakalopoulou, pour avoir été le mentor que je cherchais dans le domaine de l'intelligence artificielle appliquée à la médecine. Merci pour ses conseils avisés tant pour la rédaction d'articles que pour surmonter de multiples défis avec finesse, pour les opportunités exceptionnelles qu'elle m'a offertes, pour

sa disponibilité sans faille face aux deadlines imminentes et pour son soutien toujours solaire et motivant tout au long de ce parcours. Thank you for being a truly exceptional and motivating mentor.

Je suis reconnaissante envers mes collègues à Therapanacea, à l'Institut Gustave Roussy et CentraleSupélec avec qui j'ai pu échanger et collaborer, notamment Rafael M., Alexandre C., Marvin L., Amaury L., Baris U., Despoina I., et bien d'autres. Ils ont créé un environnement scientifique de qualité et favorisé des échanges productifs. Nos discussions et multiples collaborations ont grandement enrichi ma compréhension du sujet et mon expérience pendant ces années de thèse.

Un grand merci à mes amis, Martin P., Marion B., Grégoire B., Paul J., Imane C., Rutger F., Cordélia dB., Delphine V., Louis B., Louise B., Adrien C., Zhijin L., Adrien K., Naïs B. et tant d'autres, avec qui j'ai pu partager de nombreuses discussions, scientifiques et non scientifiques, ainsi que de nombreux fous rires.

Merci à Antoine pour son soutien quotidien, y compris pendant les longues nuits de rédaction, son optimisme et sa patience pendant cette période intense.

Enfin, à mes chers parents et ma sœur, je veux exprimer une gratitude infinie. Votre soutien inconditionnel, vos conseils et votre compréhension ont été le socle solide sur lequel reposent mes réussites. Vos encouragements ont été ma source d'inspiration. Je vous suis éternellement reconnaissante pour tout ce que vous avez fait.

Merci à tous ceux qui ont fait de cette aventure académique une expérience inoubliable.

# Bibliography

- Agostinelli, Sea et al. (2003). “GEANT4—a simulation toolkit”. In: *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3, pp. 250–303.
- Ahnesjö, A, P Andreo, and A Brahme (1987). “Calculation and application of point spread functions for treatment planning with high energy photon beams”. In: *Acta Oncologica* 26.1, pp. 49–56.
- Ahnesjö, Anders (1989). “Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media”. In: *Medical physics* 16.4, pp. 577–592.
- Aitkenhead, Adam H et al. (2020). “Automated Monte-Carlo re-calculation of proton therapy plans using Geant4/Gate: implementation and comparison to plan-specific quality assurance measurements”. In: *The British Journal of Radiology* 93.1114, p. 20200228.
- Aparna, PR and TM Libish (2023). “Automatic segmentation and classification of the liver tumor using deep learning algorithms”. In: *2023 3rd International Conference on Advances in Computing, Communication, Embedded and Secure Systems (ACCESS)*. IEEE, pp. 334–339.
- Archambault, Louis et al. (2003). “Overview of Geant4 applications in medical physics”. In: *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No. 03CH37515)*. Vol. 3. IEEE, pp. 1743–1745.
- Aswathy, AL and Vinod Chandra SS (2022). “Cascaded 3D UNet architecture for segmenting the COVID-19 infection from lung CT volume”. In: *Scientific Reports* 12.
- Badal, Andreu and Aldo Badano (2009). “Accelerating Monte Carlo simulations of photon transport in a voxelized geometry using a massively parallel graphics processing unit”. In: *Medical physics* 36.11, pp. 4878–4880.
- Bai, T. et al. (2020). *Deep Dose Plugin Towards Real-time Monte Carlo Dose Calculation Through a Deep Learning based Denoising Algorithm*. arXiv: [2011.14959 \[cs.CV\]](https://arxiv.org/abs/2011.14959).

- Beer, August and P Beer (1852). “Determination of the absorption of red light in colored liquids”. In: *Ann. Phys. Chem* 86, pp. 78–88.
- Berry, Sean L et al. (2016). “Interobserver variability in radiation therapy plan output: results of a single-institution study”. In: *Practical radiation oncology* 6.6, pp. 442–449.
- Biggs, Simon et al. (2022). “PyMedPhys: A community effort to develop an open, Python-based standard library for medical physics applications”. In: *Journal of Open Source Software* 7.78, p. 4555.
- Blumenthal, George R and Robert J Gould (1970). “Bremsstrahlung, synchrotron radiation, and compton scattering of high-energy electrons traversing dilute gases”. In: *Reviews of modern Physics* 42.2, p. 237.
- Bongratz, Fabian, Anne-Marie Rickmann, and Christian Wachinger (2023). “Abdominal organ segmentation via deep diffeomorphic mesh deformations”. In: *Scientific reports* 13.1.
- Bortfeld, Thomas (2006). “IMRT: a review and preview”. In: *Physics in Medicine & Biology* 51.13, R363.
- Bosse, Courtney et al. (2020). “Dose calculation comparisons between three modern treatment planning systems”. In: *Journal of Medical Physics* 45.3, p. 143.
- Boutilier, Justin J et al. (2015). “Models for predicting objective function weights in prostate cancer IMRT”. In: *Medical physics* 42.4, pp. 1586–1595.
- Brodzicki, Andrzej et al. (2020). “Transfer learning methods as a new approach in computer vision tasks with small datasets”. In: *Foundations of Computing and Decision Sciences* 45.3, pp. 179–193.
- Campbell, Warren G et al. (2017). “Neural network dose models for knowledge-based planning in pancreatic SBRT”. In: *Medical physics* 44.12, pp. 6148–6158.
- Çatlı, Serap and Güneş Tanır (2013). “Experimental and Monte Carlo evaluation of Eclipse treatment planning system for effects on dose distribution of the hip prostheses”. In: *Medical Dosimetry* 38.3, pp. 332–336.
- Chai, Tianfeng and Roland R Draxler (2014). “Root mean square error (RMSE) or mean absolute error (MAE)”. In: *Geoscientific model development discussions* 7.1, pp. 1525–1534.
- Chandarana, Hersh et al. (2018). “Emerging role of MRI in radiation therapy”. In: *Journal of Magnetic Resonance Imaging* 48.6, pp. 1468–1478.
- Chen, Jim X (2016). “The evolution of computing: AlphaGo”. In: *Computing in Science & Engineering* 18.4, pp. 4–7.
- Chen, Mingli et al. (2009). “Efficient gamma index calculation using fast Euclidean distance transform”. In: *Physics in Medicine & Biology* 54.7, p. 2037.

- Chen, Xinyuan et al. (2023). “Combining distance and anatomical information for deep-learning based dose distribution predictions for nasopharyngeal cancer radiotherapy planning”. In: *Frontiers in Oncology* 13, p. 1041769.
- Christ, Patrick Ferdinand et al. (2016). “Automatic liver and lesion segmentation in CT using cascaded fully convolutional neural networks and 3D conditional random fields”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19*. Springer, pp. 415–423.
- Çiçek, Özgün et al. (2016a). “3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation”. In: arXiv: [1606.06650 \[cs.CV\]](https://arxiv.org/abs/1606.06650).
- (2016b). “3D U-Net: learning dense volumetric segmentation from sparse annotation”. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part II 19*. Springer, pp. 424–432.
- Clarkson, JR (1941). “A note on depth doses in fields of irregular shape”. In: *The British Journal of Radiology* 14.164, pp. 265–268.
- Constantin, Magdalena et al. (2010). “Linking computer-aided design (CAD) to Geant4-based Monte Carlo simulations for precise implementation of complex treatment head geometries”. In: *Physics in Medicine & Biology* 55.8, N211.
- Cui, Yunfeng et al. (2015). “Contouring variations and the role of atlas in non-small cell lung cancer radiation therapy: Analysis of a multi-institutional preclinical trial planning study”. In: *Practical radiation oncology* 5.2, e67–e75.
- Cunningham, JR (1972). “Scatter-air ratios”. In: *Physics in medicine & biology* 17.1, p. 42.
- De Martino, Fortuna et al. (2021). “Dose calculation algorithms for external radiation therapy: an overview for practitioners”. In: *Applied sciences* 11.15, p. 6806.
- Delpon, Grégory et al. (2016). “Comparison of automated atlas-based segmentation software for postoperative prostate cancer radiotherapy”. In: *Frontiers in oncology* 6, p. 178.
- DeVita Jr, VT, TS Lawrence, and SA Rosenberg (2019). *Cancer: Principles and Practice of Oncology.. [Internet] Philadelphia*.
- Didi, Samir et al. (2022). “Experimental validation of a linac head Geant4 model under a grid computing environment”. In: *Biomedical Physics & Engineering Express* 8.2, p. 025007.
- Dijk, Robert HW van et al. (2022). “A novel multichannel deep learning model for fast denoising of Monte Carlo dose calculations: preclinical applications”. In: *Physics in Medicine & Biology* 67.16, p. 164001.

- Ding, Jie et al. (2022). “Automatic contour refinement for deep learning auto-segmentation of complex organs in MRI-guided adaptive radiation therapy”. In: *Advances in Radiation Oncology* 7.5, p. 100968.
- Drzymala, RE et al. (1991). “Dose-volume histograms”. In: *International Journal of Radiation Oncology\* Biology\* Physics* 21.1, pp. 71–78.
- Duderstadt, James J and William Russell Martin (1979). “Transport theory.” In: *Transport theory*.
- Effects of Atomic Radiation, United Nations Scientific Committee on the et al. (1982). “Ionizing radiation: sources and biological effects. 1982 report to the general assembly, with annexes”. In.
- Failla, Gregory A et al. (2010). “Acuros XB advanced dose calculation for the Eclipse treatment planning system”. In: *Palo Alto, CA: Varian Medical Systems* 20, p. 18.
- Fan, Jiawei et al. (2019). “Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique”. In: *Medical physics* 46.1, pp. 370–381.
- Fiak, Meriem et al. (Jan. 2021). “Monte Carlo Simulation of a 18 MV Medical Linac Photon Beam Using GATE/GEANT4”. In: *Moscow University Physics Bulletin* 76, pp. 15–21. DOI: [10.3103/S0027134921010069](https://doi.org/10.3103/S0027134921010069).
- Gasteuil, Jean et al. (2019). “A GATE/Geant4 Monte Carlo toolkit for surface dose calculation in VMAT breast cancer radiotherapy”. In: *Physica Medica* 61, pp. 112–117.
- Glide-Hurst, C et al. (2013). “Commissioning of the Varian TrueBeam linear accelerator: a multi-institutional study”. In: *Medical physics* 40.3, p. 031719.
- Goodfellow, Ian J. et al. (2014). *Generative Adversarial Networks*. arXiv: [1406.2661](https://arxiv.org/abs/1406.2661) [stat.ML].
- Götz, Th I et al. (2020). “A deep learning approach to radiation dose estimation”. In: *Physics in Medicine & Biology* 65.3, p. 035007.
- Gray, Alison, LD Oliver, and PN Johnston (2009). “The accuracy of the pencil beam convolution and anisotropic analytical algorithms in predicting the dose effects due to attenuation from immobilization devices and large air gaps”. In: *Medical physics* 36.7, pp. 3181–3191.
- Grevillot, L et al. (2011). “Simulation of a 6 MV Elekta Precise Linac photon beam using GATE/GEANT4”. In: *Physics in Medicine & Biology* 56.4, p. 903.
- Gu, Xuejun, Xun Jia, and Steve B Jiang (2011). “GPU-based fast gamma index calculation”. In: *Physics in Medicine & Biology* 56.5, p. 1431.
- Han, Tao et al. (2011). “Dosimetric comparison of Acuros XB deterministic radiation transport method with Monte Carlo and model-based convolution methods in heterogeneous media”. In: *Medical Physics* 38.5, pp. 2651–2664.

- Hasenbalg, Federico et al. (2007). “Collapsed cone convolution and analytical anisotropic algorithm dose calculations compared to VMC++ Monte Carlo simulations in clinical cases”. In: *Physics in Medicine & Biology* 52.13, p. 3679.
- Hendrycks, Dan and Kevin Gimpel (2016). “Gaussian error linear units (gelus)”. In: *arXiv preprint arXiv:1606.08415*.
- Hissoiny, S et al. (2011). “Fast dose calculation in magnetic fields with GPUMCD”. In: *Physics in Medicine & Biology* 56.16, p. 5119.
- Ho, Jonathan, Ajay Jain, and Pieter Abbeel (2020). “Denoising diffusion probabilistic models”. In: *Advances in neural information processing systems* 33, pp. 6840–6851.
- Hrinivich, William Thomas and Junghoon Lee (2020). “Artificial intelligence-based radiotherapy machine parameter optimization using reinforcement learning”. In: *Medical physics* 47.12, pp. 6140–6150.
- Hrinivich, WT, D Song, and J Lee (2023). “Machine Parameter Optimization of a Clinical Linear Accelerator Using Deep Reinforcement Learning for Automatic Generation of Deliverable Prostate VMAT Plans”. In: *International Journal of Radiation Oncology, Biology, Physics* 117.2, S56.
- Huynh, Elizabeth et al. (2020). “Artificial intelligence in radiation oncology”. In: *Nature Reviews Clinical Oncology* 17.12, pp. 771–781.
- Ibragimov, Bulat and Lei Xing (2017). “Segmentation of organs-at-risks in head and neck CT images using convolutional neural networks”. In: *Medical physics* 44.2, pp. 547–557.
- Ioffe, Sergey and Christian Szegedy (2015). *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*. arXiv: [1502.03167 \[cs.LG\]](https://arxiv.org/abs/1502.03167).
- Isola, Phillip et al. (2017). “Image-to-Image Translation with Conditional Adversarial Networks”. In: *CVPR*.
- Jochems, Arthur et al. (2017). “Developing and validating a survival prediction model for NSCLC patients through distributed learning across 3 countries”. In: *International Journal of Radiation Oncology\* Biology\* Physics* 99.2, pp. 344–352.
- Johnstone, Emily et al. (2018). “Systematic review of synthetic computed tomography generation methodologies for use in magnetic resonance imaging-only radiation therapy”. In: *International Journal of Radiation Oncology\* Biology\* Physics* 100.1, pp. 199–217.
- Kandalan, Roya Norouzi et al. (2020). “Dose prediction with deep learning for prostate cancer radiation therapy: model adaptation to different treatment planning practices”. In: *Radiotherapy and Oncology* 153, pp. 228–235.

- Kase, Kenneth R and Walter Ralph Nelson (1978). “Concepts of radiation dosimetry”. In: (*No Title*).
- Kearney, Vasant et al. (2018). “DoseNet: a volumetric dose prediction algorithm using 3D fully-convolutional neural networks”. In: *Physics in Medicine & Biology* 63.23, p. 235022.
- Khalil, Muhammad Ibrahim et al. (2022). “Multi-Scale Network for Thoracic Organs Segmentation.” In: *Computers, Materials & Continua* 70.2.
- Khan, Faiz M et al. (1973). “Computer and approximation methods of calculating depth dose in irregularly shaped fields”. In: *Radiology* 106.2, pp. 433–436.
- Khoo, Vincent S et al. (2000). “A comparison of clinical target volumes determined by CT and MRI for the radiotherapy planning of base of skull meningiomas”. In: *International Journal of Radiation Oncology\* Biology\* Physics* 46.5, pp. 1309–1317.
- Khoo, VS et al. (1999). “Comparison of MRI with CT for the radiotherapy planning of prostate cancer: a feasibility study.” In: *The British journal of radiology* 72.858, pp. 590–597.
- Kim, Y et al. (2016). “Impact of contouring accuracy on expected tumor control probability for head and neck cancer: semiautomated segmentation versus manual contouring”. In: *International journal of radiation oncology, biology, physics* 96.2, E545.
- Kolesnikov, Alexander et al. (2020). “Big transfer (bit): General visual representation learning”. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*. Springer, pp. 491–507.
- Krim, Deae-Eddine et al. (2020). “GATE Simulation of 6 MV Photon Beam Produced by Elekta Medical Linear Accelerator”. In: *International Conference on Electronic Engineering and Renewable Energy*. Springer, pp. 301–307.
- Kutcher, Gerald J et al. (1994). “Comprehensive QA for radiation oncology: report of AAPM radiation therapy committee task group 40”. In: *MEDICAL PHYSICS-LANCASTER PA-21*, pp. 581–581.
- Lachinov, Dmitry, Evgeny Vasiliev, and Vadim Turlapov (2018). “Glioma segmentation with cascaded UNet”. In: *International MICCAI Brainlesion Workshop*. Springer, pp. 189–198.
- Landry, Guillaume et al. (2010). “Sensitivity of low energy brachytherapy Monte Carlo dose calculations to uncertainties in human tissue composition”. In: *Medical physics* 37.10, pp. 5188–5198.
- Lee, Min Sun et al. (2019). “Deep-dose: a voxel dose estimation method using deep convolutional neural network for personalized internal dosimetry”. In: *Scientific reports* 9.1, p. 10308.
- Lehtinen, J. et al. (2018). “Noise2noise: Learning image restoration without clean data”. In: *arXiv*.

- Li, Ruikun et al. (2023). “MRI-based two-stage deep learning model for automatic detection and segmentation of brain metastases”. In: *European Radiology*, pp. 1–11.
- Liang, Zhi-Pei and Paul C Lauterbur (2000). *Principles of magnetic resonance imaging*. SPIE Optical Engineering Press Bellingham, WA.
- Lin, Tsung-Yi et al. (2017). “Focal Loss for Dense Object Detection”. In: *2017 IEEE ICCV*, pp. 2999–3007.
- Liu, Hongying et al. (2019). “CU-Net: Cascaded U-Net with loss weighted sampling for brain tumor segmentation”. In: *Multimodal Brain Image Analysis and Mathematical Foundations of Computational Anatomy: 4th International Workshop, MBIA 2019, and 7th International Workshop, MFCA 2019, Held in Conjunction with MICCAI 2019, Shenzhen, China, October 17, 2019, Proceedings 4*. Springer, pp. 102–111.
- Liu, Zhuang et al. (2022). “A convnet for the 2020s”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11976–11986.
- Loshchilov, Ilya and Frank Hutter (2017a). “Decoupled weight decay regularization”. In: *arXiv preprint arXiv:1711.05101*.
- (2017b). “Decoupled weight decay regularization”. In: *arXiv preprint arXiv:1711.05101*.
- Lou, Bin et al. (2019). “An image-based deep learning framework for individualising radiotherapy dose: a retrospective analysis of outcome prediction”. In: *The Lancet Digital Health* 1.3, e136–e147.
- Low, D. et al. (1998). “A technique for the quantitative evaluation of dose distributions.” In: *Medical physics*.
- Lustberg, Tim et al. (2018). “Clinical evaluation of atlas and deep learning based automatic contouring for lung cancer”. In: *Radiotherapy and Oncology* 126.2, pp. 312–317.
- Ma, CM et al. (2005). “Effect of statistical uncertainties on Monte Carlo treatment planning”. In: *Physics in Medicine & Biology* 50.5, p. 891.
- Maas, Andrew L. (2013). “Rectifier Nonlinearities Improve Neural Network Acoustic Models”. In.
- Martini, Stefania et al. (2020). “Radiation therapy for oligometastatic oropharyngeal cancer”. In: *BJR— case reports* 6.1, p. 20190021.
- Martinot, S et al. (2021). “OC-0308 Fast Monte-Carlo dose simulation with recurrent deep learning”. In: *Radiotherapy and Oncology* 161, S216–S217.
- Martinot, Sonia et al. (2021). “High-Particle Simulation of Monte-Carlo Dose Distribution with 3D ConvLSTMs”. In: *MICCAI*, pp. 499–508.

- Martinsen, Paul et al. (2009). “Accelerating Monte Carlo simulations with an NVIDIA® graphics processor”. In: *Computer Physics Communications* 180.10, pp. 1983–1989.
- Masi, Laura et al. (2013). “Impact of plan parameters on the dosimetric accuracy of volumetric modulated arc therapy”. In: *Medical physics* 40.7, p. 071718.
- Mayles, Philip, Alan Nahum, and Jean-Claude Rosenwald (2007). *Handbook of radiotherapy physics: theory and practice*. CRC Press.
- Memory, Long Short-Term (2010). “Long short-term memory”. In: *Neural computation* 9.8, pp. 1735–1780.
- Merwe, Debbie van der et al. (2017). “Accuracy requirements and uncertainties in radiotherapy: a report of the International Atomic Energy Agency”. In: *Acta oncologica* 56.1, pp. 1–6.
- Metropolis, N (1987). “The beginning”. In: *Los Alamos Science* 15, pp. 125–130.
- Mettler, Fred A (1985). “Medical effects of ionizing radiation”. In.
- Mnih, Volodymyr et al. (2015). “Human-level control through deep reinforcement learning”. In: *nature* 518.7540, pp. 529–533.
- Mohan, Radhe, C Chui, and L Lidofsky (1986). “Differential pencil beam dose computation model for photons”. In: *Medical Physics* 13.1, pp. 64–73.
- Nelms, Benjamin E et al. (2013). “Evaluating IMRT and VMAT dose accuracy: practical examples of failure to detect systematic errors when applying a commonly used metric and action levels”. In: *Medical physics* 40.11, p. 111722.
- Neph, Ryan et al. (2019a). “Deepmcdose: A deep learning method for efficient monte carlo beamlet dose calculation by predictive denoising in mr-guided radiotherapy”. In: *Workshop on Artificial Intelligence in Radiation Therapy*. Springer, pp. 137–145.
- Neph, Ryan et al. (2019b). “Parallel beamlet dose calculation via beamlet contexts in a distributed multi-GPU framework”. In: *Medical physics* 46.8, pp. 3719–3733.
- Nikolov, Stanislav et al. (2021). “Clinically applicable segmentation of head and neck anatomy for radiotherapy: deep learning algorithm development and validation study”. In: *Journal of medical Internet research* 23.7, e26151.
- Nowak, Michel (2018). “Accelerating Monte Carlo particle transport with adaptively generated importance maps”. PhD thesis. Université Paris-Saclay (ComUE).
- Oberije, Cary et al. (2015). “A validated prediction model for overall survival from stage III non-small cell lung cancer: toward survival prediction for individual patients”. In: *International Journal of Radiation Oncology\* Biology\* Physics* 92.4, pp. 935–944.
- Oldendorf, William H (1978). “The quest for an image of brain: a brief historical and technical review of brain imaging techniques”. In: *Neurology* 28.6, pp. 517–517.

- Otto, K. (2008). “Volumetric modulated arc therapy: IMRT in a single gantry arc”. In: *Medical physics*.
- Paganetti, Harald (2012). “Range uncertainties in proton therapy and the role of Monte Carlo simulations”. In: *Physics in Medicine & Biology* 57.11, R99.
- Pastor-Serrano, Oscar and Zoltán Perkó (2022). “Millisecond speed deep learning based proton dose calculation with Monte Carlo accuracy”. In: *Physics in Medicine & Biology* 67.10, p. 105006.
- Peng, Zhao et al. (2019). “Deep learning for accelerating Monte Carlo radiation transport simulation in intensity-modulated radiation therapy”. In: *arXiv preprint arXiv:1910.07735*.
- Pflüger, Irada et al. (2022). “Automated detection and quantification of brain metastases on clinical MRI data using artificial neural networks”. In: *Neuro-oncology advances* 4.1, vdac138.
- Podgorsak, Ervin B (2005). *Radiation oncology physics*.
- Qilin, Zhang et al. (2022). “The feasibility study on the generalization of deep learning dose prediction model for volumetric modulated arc therapy of cervical cancer”. In: *Journal of Applied Clinical Medical Physics* 23.6, e13583.
- Quan, Enzhuo et al. (July 2012). “A Comprehensive Comparison of IMRT and VMAT Plan Quality for Prostate Cancer Treatment”. In: *International journal of radiation oncology, biology, physics* 83, pp. 1169–78. DOI: [10.1016/j.ijrobp.2011.09.015](https://doi.org/10.1016/j.ijrobp.2011.09.015).
- Raeside, DE (1976). “Monte Carlo principles and applications”. In: *Physics in Medicine & Biology* 21.2, p. 181.
- Raffuzzi, Valeria, Eugene Shwageraus, and Lee Morgan (2022). “Accelerating Monte Carlo neutron transport by approximating thermal cross sections with functional forms”. In: *Annals of Nuclear Energy* 169, p. 108819.
- Raghu, Maithra et al. (2019). “Transfusion: Understanding transfer learning for medical imaging”. In: *Advances in neural information processing systems* 32.
- Ramachandran, Prajit, Barret Zoph, and Quoc V Le (2017). “Searching for activation functions”. In: *arXiv preprint arXiv:1710.05941*.
- Richmond, Neil et al. (2021). “Comparison of the RayStation photon Monte Carlo dose calculation algorithm against measured data under homogeneous and heterogeneous irradiation geometries”. In: *Physica Medica* 82, pp. 87–99.
- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox (2015a). “U-net: Convolutional networks for biomedical image segmentation”. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.

- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox (2015b). “U-net: Convolutional networks for biomedical image segmentation”. In: *MICCAI*, pp. 234–241.
- Rosenberg, I (2008). “Radiation oncology physics: a handbook for teachers and students”. In: *British journal of cancer* 98.5, pp. 1020–1020.
- Sardari, D et al. (2010). “Measurement of depth-dose of linear accelerator and simulation by use of Geant4 computer code”. In: *Reports of practical oncology and radiotherapy* 15.3, pp. 64–68.
- Sarrut, David et al. (2022). “The OpenGATE ecosystem for Monte Carlo simulation in medical physics”. In: *Physics in Medicine & Biology* 67.18, p. 184001.
- Savova, Guergana K et al. (2017). “DeepPhe: a natural language processing system for extracting cancer phenotypes from clinical records”. In: *Cancer research* 77.21, e115–e118.
- Schmidt, Maria A and Geoffrey S Payne (2015). “Radiotherapy planning using MRI”. In: *Physics in Medicine & Biology* 60.22, R323.
- Schrittwieser, Julian et al. (2020). “Mastering atari, go, chess and shogi by planning with a learned model”. In: *Nature* 588.7839, pp. 604–609.
- Schuster, Mike and Kuldip K Paliwal (1997). “Bidirectional recurrent neural networks”. In: *IEEE transactions on Signal Processing* 45.11, pp. 2673–2681.
- Shen, Chenyang et al. (2019). “Intelligent inverse treatment planning via deep reinforcement learning, a proof-of-principle study in high dose-rate brachytherapy for cervical cancer”. In: *Physics in Medicine & Biology* 64.11, p. 115013.
- Shi, Xingjian et al. (2015). “Convolutional LSTM network: A machine learning approach for precipitation nowcasting”. In: *Advances in neural information processing systems* 28.
- Sudre, Carole H et al. (2017). “Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations”. In: *Deep learning in medical image analysis and multi-modal learning for clinical decision support*, pp. 240–248.
- Teixeira, MS et al. (2019). “Monte Carlo simulation of Novalis Classic 6 MV accelerator using phase space generation in GATE/Geant4 code”. In: *Progress in Nuclear Energy* 110, pp. 142–147.
- Thwaites, David and John Tuohy (Aug. 2006). “Back to the future: The history and development of the clinical linear accelerator”. In: *Physics in medicine and biology* 51, R343–62. DOI: [10.1088/0031-9155/51/13/R20](https://doi.org/10.1088/0031-9155/51/13/R20).
- Thwaites, DAVID I, BJ Mijnheer, and JOHN A Mills (2005). “Quality assurance of external beam radiotherapy”. In: *Radiation oncology physics: a handbook for teachers and students*. Vienna: International Atomic Energy Agency, pp. 407–450.

- Tillikainen, L et al. (2008). “A 3D pencil-beam-based superposition algorithm for photon dose calculation in heterogeneous media”. In: *Physics in Medicine & Biology* 53.14, p. 3821.
- Torrey, Lisa and Jude Shavlik (2010). “Transfer learning”. In: *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*. IGI global, pp. 242–264.
- Udagedara, Indika et al. (2015). “Reduced order modeling for accelerated Monte Carlo simulations in radiation transport”. In: *Applied Mathematics and Computation* 267, pp. 237–251.
- Ullah, Ihsan et al. (2023). “A deep learning based dual encoder–decoder framework for anatomical structure segmentation in chest X-ray images”. In: *Scientific Reports* 13.1, p. 791.
- Vandewinckele, Liesbeth et al. (2022). “Treatment plan prediction for lung IMRT using deep learning based fluence map generation”. In: *Physica Medica* 99, pp. 44–54.
- Vangvichith, M et al. (2019). “Comparison of five dose calculation algorithms in a heterogeneous media using design of experiment”. In: *Physica Medica* 61, pp. 103–111.
- Vanneste, Ben et al. (Dec. 2016). “Prostate Cancer Radiation Therapy: What Do Clinicians Have to Know?” In: *BioMed Research International* 2016, pp. 1–14. DOI: [10.1155/2016/6829875](https://doi.org/10.1155/2016/6829875).
- Vassiliev, Oleg N et al. (2010). “Validation of a new grid-based Boltzmann equation solver for dose calculation in radiotherapy with photon beams”. In: *Physics in Medicine & Biology* 55.3, p. 581.
- Vasudevan, V et al. (2020). “Combining Monte Carlo with Deep Learning: Predicting High-resolution, Low-noise Dose Distributions Using a Generative Adversarial Network for Fast and Precise Monte Carlo Simulations”. In: *International Journal of Radiation Oncology, Biology, Physics* 108.3, S44–S45.
- Vaswani, Ashish et al. (2017). “Attention is all you need”. In: *Advances in neural information processing systems* 30.
- Voss, Luke et al. (2023). “BayesDose: Comprehensive proton dose prediction with model uncertainty using Bayesian LSTMs”. In: *arXiv preprint arXiv:2307.01151*.
- Walker, Gary V et al. (2014). “Prospective randomized double-blind study of atlas-based organ-at-risk autosegmentation-assisted radiation planning in head and neck cancer”. In: *Radiotherapy and Oncology* 112.3, pp. 321–325.
- Wang, Tonghe et al. (2020). “A learning-based automatic segmentation and quantification method on left ventricle in gated myocardial perfusion SPECT imaging: A feasibility study”. In: *Journal of Nuclear Cardiology* 27, pp. 976–987.

- Wen, Lu et al. (2023). “Multi-level progressive transfer learning for cervical cancer dose prediction”. In: *Pattern Recognition* 141, p. 109606.
- Wendling, Markus et al. (2007). “A fast algorithm for gamma evaluation in 3D”. In: *Medical physics* 34.5, pp. 1647–1654.
- Wuthrick, Evan J et al. (2015). “Institutional clinical trial accrual volume and survival of patients with head and neck cancer”. In: *Journal of Clinical Oncology* 33.2, p. 156.
- Xiang, T. et al. (2020). “BiO-Net: Learning Recurrent Bi-directional Connections for Encoder-Decoder Architecture”. In.
- Xing, Yixun et al. (2020). “A feasibility study on deep learning-based radiotherapy dose calculation”. In: *Medical physics* 47.2, pp. 753–758.
- Yu, Cenji et al. (2022). “Multi-organ segmentation of abdominal structures from non-contrast and contrast enhanced CT images”. In: *Scientific reports* 12.1, p. 19093.
- Zhan, Bo et al. (2022). “Multi-constraint generative adversarial network for dose prediction in radiotherapy”. In: *Medical Image Analysis* 77, p. 102339.
- Zhao, Hang et al. (2016). “Loss functions for image restoration with neural networks”. In: *IEEE Transactions on computational imaging* 3.1, pp. 47–57.
- Zhong, Zisha et al. (2019). “Simultaneous cosegmentation of tumors in PET-CT images using deep fully convolutional networks”. In: *Medical physics* 46.2, pp. 619–633.
- Zhou Wang et al. (2004). “Image quality assessment: from error visibility to structural similarity”. In: *IEEE Transactions on Image Processing* 13.4, pp. 600–612. DOI: [10.1109/TIP.2003.819861](https://doi.org/10.1109/TIP.2003.819861).
- Zhu, Wentao et al. (2019). “AnatomyNet: deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy”. In: *Medical physics* 46.2, pp. 576–589.