Proton computed tomography from sandwich time-of-flight data

Aurélien Coussat

CREATIS

29/09/2025

Outline

Outline

Medical imaging in a few words

Energy-loss proton computed tomography
- Motivation
- Scanner design

Time-of-flight proton computed tomography

Sandwich Time-of-flight proton computed tomography
- First proposal: 1D calibration curve
- Second proposal: 2D calibration curve
- Third proposal: formulation as an optimization problem

Conclusion

Goal of medical imaging

- *Medical imaging refers to techniques and process of imaging the interior of a body for clinical analysis and medical intervention*[1]
- Many techniques exist:
  - Magnetic resonance imaging (MRI)
  - X-ray computed tomography (CT)
  - Positron emission tomography (PET)
  - Single photon emission Computed tomography (SPECT)
  - ... and many more, including proton computed tomography (pCT)!

[1] https://en.wikipedia.org/wiki/Medical_imaging

Goal of medical imaging

- *Medical imaging refers to techniques and process of imaging the interior of a body for clinical analysis and medical intervention*[1]
- Many techniques exist:
  - Magnetic resonance imaging (MRI)
  - X-ray computed tomography (CT)
  - Positron emission tomography (PET)
  - Single photon emission Computed tomography (SPECT)
  - ...and many more, including proton computed tomography (pCT)!

[1] https://en.wikipedia.org/wiki/Medical_imaging

How does X-ray CT work?

- A rotating X-ray tube emits X-rays towards the patient

How does X-ray CT work?

- A rotating X-ray tube emits X-rays towards the patient
- X-rays are attenuated when traversing the various tissues (described by $\mu(x)$)

How does X-ray CT work?

- A rotating X-ray tube emits X-rays towards the patient
- X-rays are attenuated when traversing the various tissues (described by $\mu(x)$)
- Residual energy is detected in order to estimate the integral of the attenuation along the ray (Beer-Lambert law: $-\log\left(\frac{I}{I_0}\right) = \int \mu(x)\mathrm{d}x$)

How does X-ray CT work?

- A rotating X-ray tube emits X-rays towards the patient
- X-rays are attenuated when traversing the various tissues (described by $\mu(x)$)
- Residual energy is detected in order to estimate the integral of the attenuation along the ray (Beer-Lambert law: $-\log\left(\frac{I}{I_0}\right) = \int \mu(x)\mathrm{d}x$)
- For each gantry position, a projection is acquired, eventually forming a sinogram

Tomographic image reconstruction

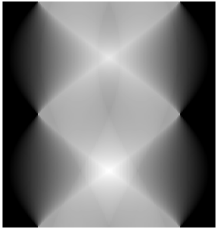

Tomographic image reconstruction

- Goal: turn sinogram data into an image understandable by the human eye

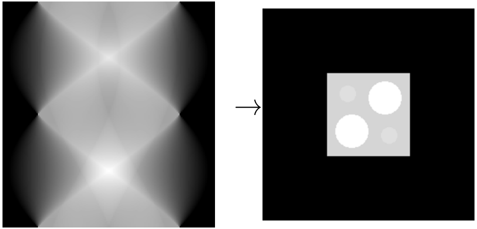

Tomographic image reconstruction

- Goal: turn sinogram data into an image understandable by the human eye
- Many algorithms exist, that fall into two categories:
  - Analytic algorithms (FBP, FDK)
  - Iterative algorithms (ART, SART, CG...)

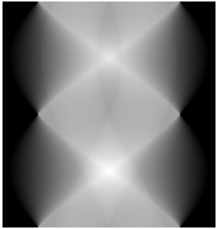

$\rightarrow$

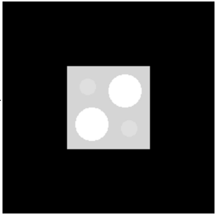

Outline

Outline

Proton therapy

- When traversing a material, protons deposit their energy in a well-defined area (Bragg peak phenomenon)
- That makes protons good candidates to destroy tumoral cells
- This process is known as proton therapy

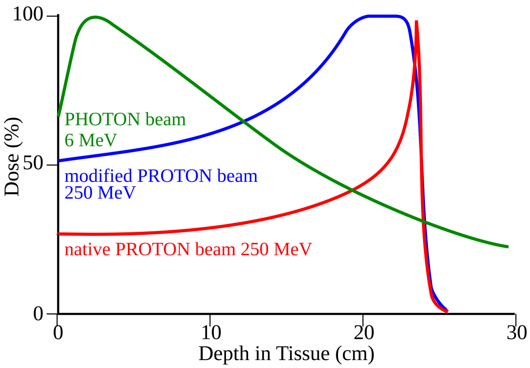

Proton range

- Location and shape of the Bragg peak must be precisely determined

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton range

- Location and shape of the Bragg peak must be precisely determined
- Treatment planning requires a map of the RSP distribution within the patient's body

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton range

- Location and shape of the Bragg peak must be precisely determined
- Treatment planning requires a map of the RSP distribution within the patient's body
  - SP: ability for a material to stop protons

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton range

- Location and shape of the Bragg peak must be precisely determined
- Treatment planning requires a map of the RSP distribution within the patient's body
  - SP: ability for a material to stop protons
  - $\text{RSP} = \frac{\text{SP}_{\textbf{material}}}{\text{SP}_{\textbf{water}}}$

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton range

- Location and shape of the Bragg peak must be precisely determined
- Treatment planning requires a map of the RSP distribution within the patient's body
  - SP: ability for a material to stop protons
  - $\text{RSP} = \frac{\text{SP}_{\textbf{material}}}{\text{SP}_{\textbf{water}}}$
- Currently determined by converting HU from X-ray CT

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton range

- Location and shape of the Bragg peak must be precisely determined
- Treatment planning requires a map of the RSP distribution within the patient's body
  - SP: ability for a material to stop protons
  - RSP $= \frac{SP_{\textbf{material}}}{SP_{\textbf{water}}}$
- Currently determined by converting HU from X-ray CT
- This conversion introduces uncertainties ranging from 1 % to 5 %[2], resulting in increased saftey margins

[2] Ming Yang et al. In: *Physics in Medicine & Biology* 57.13 (2012)

Proton Computed tomography

- Goal: use protons to directly recover the RSP distribution
- Shoot protons with enough energy to completely traverse the patient and detect their residual energy
- Use tomographic reconstruction techniques in order to reconstruct an image of the RSP distribution

Proton physics

In the patient, protons undergo several interactions:

[3] Catherine Thérèse Quiñones. PhD thesis. Université de Lyon, 2016

Proton physics

In the patient, protons undergo several interactions:

- Multiple Coulomb scattering (MCS): stochastic collision of protons onto atomic eletrons

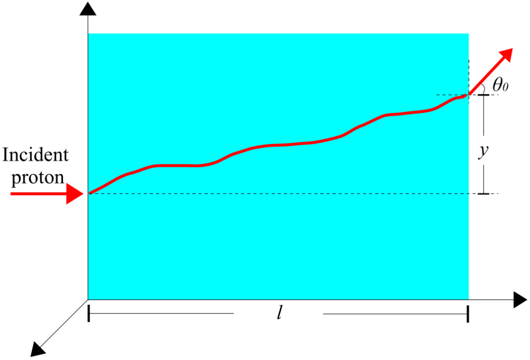

Courtesy of Catherine Thérèse Quiñones[3].

[3] Catherine Thérèse Quiñones. PhD thesis. Université de Lyon, 2016

Proton physics

In the patient, protons undergo several interactions:

- Multiple Coulomb scattering (MCS): stochastic collision of protons onto atomic eletrons
- Nuclear interactions, that can be...

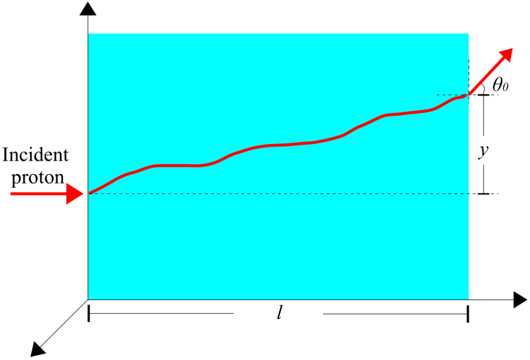

Courtesy of Catherine Thérèse Quiñones[3].

[3] Catherine Thérèse Quiñones. PhD thesis. Université de Lyon, 2016

Proton physics

In the patient, protons undergo several interactions:

- Multiple Coulomb scattering (MCS): stochastic collision of protons onto atomic eletrons
- Nuclear interactions, that can be...
  - Elastic: the indicent proton scatters off a nucleus

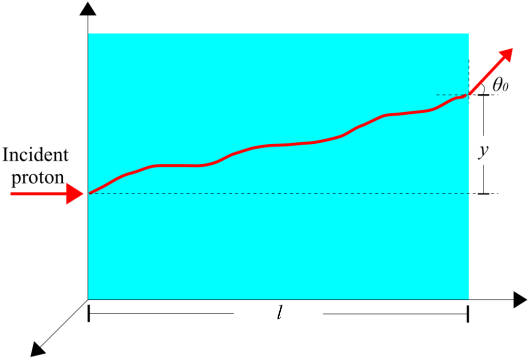

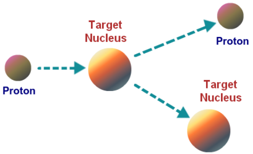

Courtesy of Catherine Thérèse Quiñones[3].

[3] Catherine Thérèse Quiñones. PhD thesis. Université de Lyon, 2016

Proton physics

In the patient, protons undergo several interactions:

- Multiple Coulomb scattering (MCS): stochastic collision of protons onto atomic eletrons
- Nuclear interactions, that can be...
  - Elastic: the indicent proton scatters off a nucleus
  - Inelastic: the proton is absorbed then re-emitted by a nucleus

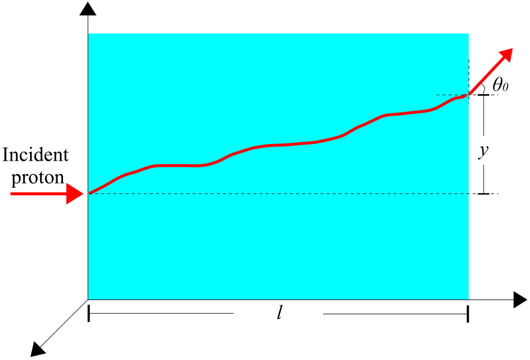

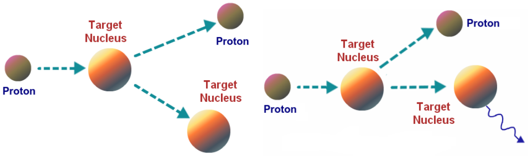

Courtesy of Catherine Thérèse Quiñones[3].

[3] Catherine Thérèse Quiñones. PhD thesis. Université de Lyon, 2016

Outline

Energy-loss pCT scanner

- Up- and downstream detectors track the protons' position and direction to estimate their most likely path (MLP)
- The calorimeter measures the residual energy to estimate the integral of the SP along the MLP
- The scanner (or the object) rotates

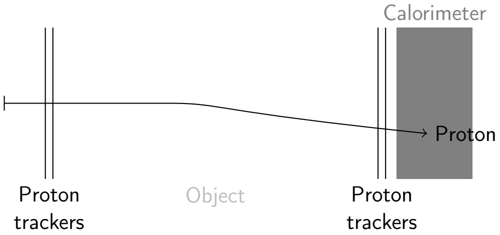

List-mode pCT data

- Detected protons's information are stored in list-mode files
  - Incoming position and direction
  - Outgoing position and direction
  - Residual energy

[4] Simon Rit et al. In: *Medical Physics* 40.3 (2013)

List-mode pCT data

- Detected protons's information are stored in list-mode files
  - Incoming position and direction
  - Outgoing position and direction
  - Residual energy
- "Distance-driven projections" are formed by properly binning these data

[4] Simon Rit et al. In: *Medical Physics* 40.3 (2013)

List-mode pCT data

- Detected protons's information are stored in list-mode files
  - Incoming position and direction
  - Outgoing position and direction
  - Residual energy
- "Distance-driven projections" are formed by properly binning these data
- Reconstruction can be achieved using "distance-driven FBP"[4]

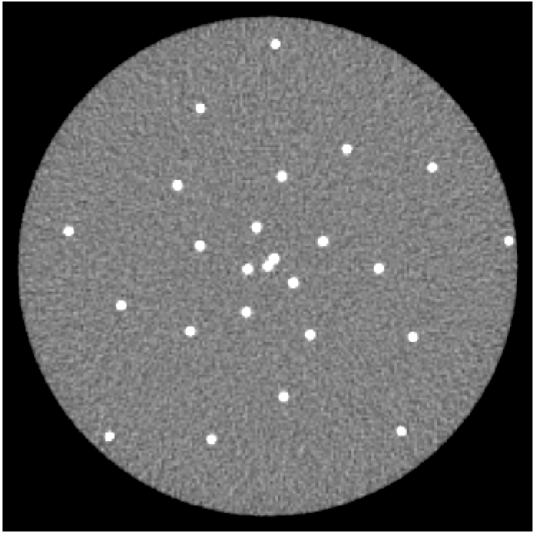

[4] Simon Rit et al. In: *Medical Physics* 40.3 (2013)

Outline

Time-of-flight pCT

- Calorimeter is replaced by two time-tracking low gain avalanche detectors[5]
- The time-of-flight (TOF) is converted to the corresponding energy
- The reconstruction can then be carried out as in energy-loss pCT

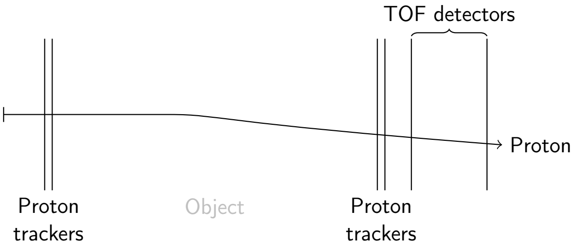

[5] Felix Ulrich-Pur et al. In: *Physics in Medicine & Biology* 67.9 (2022), Nils Krah et al. In: *Physics in Medicine & Biology* 67.16 (2022)

Outline

Medical imaging in a few words

Energy-loss proton computed tomography
- Motivation
- Scanner design

Time-of-flight proton computed tomography

Sandwich Time-of-flight proton computed tomography
- First proposal: 1D calibration curve
- Second proposal: 2D calibration curve
- Third proposal: formulation as an optimization problem

Conclusion

Sandwich Time-of-flight pCT

- The residual energy is not measured anymore
- Instead, the time is measured in the up- and downstream detectors[6]
- How to reconstruct the RSP distribution from TOF measurements only?

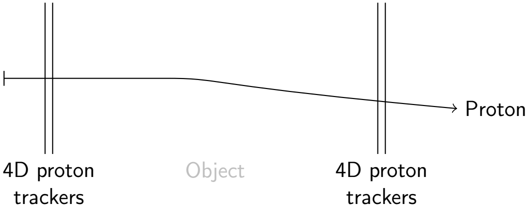

[6] F. Ulrich-Pur et al. In: *Journal of Instrumentation* 18.02 (2023)

Difficulty of sandwich time-of-flight pCT

- Protons traversing different materials with equivalent water-equivalent path lengths (WEPLs) result in
  - equal energy loss
  - different TOFs
- The conversion between TOF and WEPL is non-bijective
- The forward problem is non-linear and non-trivial to invert

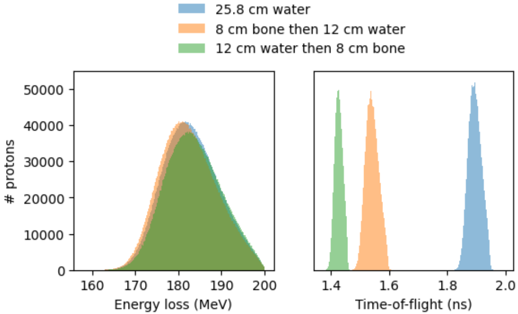

Outline

Monte Carlo simulation

- Perfect detectors of 40 cm in the x and y directions located at $z = \{-110\,\text{mm}, 110\,\text{mm}\}$
- Parallel, 40 cm wide 200 MeV proton beam located at $z = -1000\,\text{mm}$
- The scanned object is located between the detectors
- At each of the 720 steps...
  - The object rotates around the y axis by 0.5°
  - 72 000 protons are shot
- Implemented with GATE [7] version 10

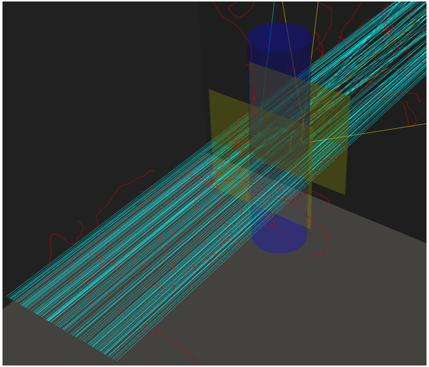

[7] Nils Krah et al. In: *XXth International Conference on the use of Computers in Radiation therapy*. 2024

TOF to WEPL conversion

- Based on the Monte Carlo simulation explained before
- Object is a water box of $40 \times 40 \times L$ cm, where L is the length of the box in the z direction
- The simulation is repeated 100 times for L between 0 mm and 220 mm
- In each simulation, 10^4 protons are emitted from a point-like source
- For each simulation, the median energy loss and TOF are stored
- A polynomial curve is finally fitted

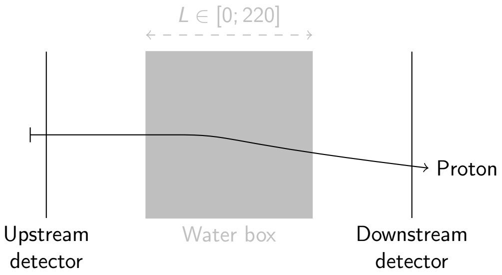

TOF and energy-loss to WEPL

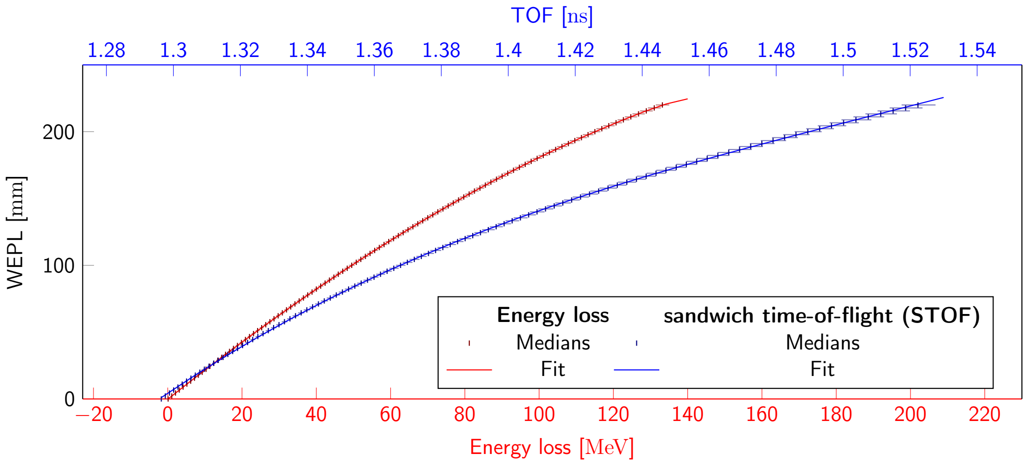

Tomographic reconstructions

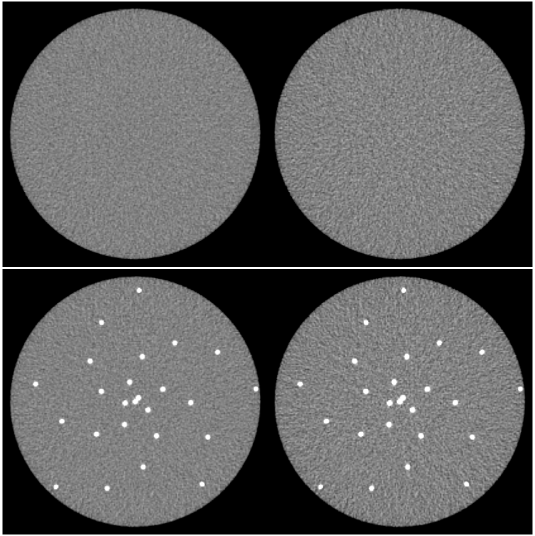

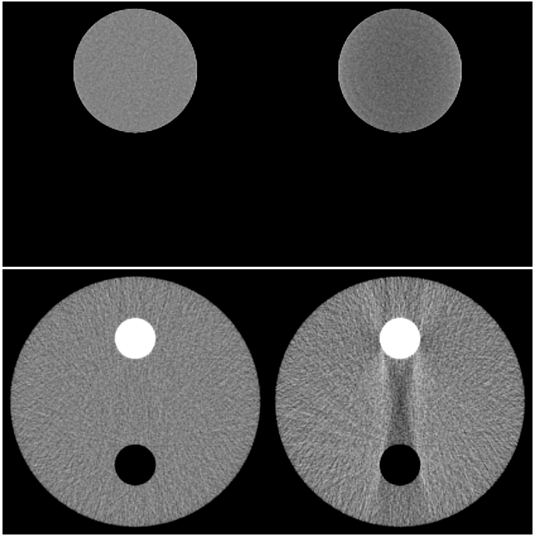

Profiles

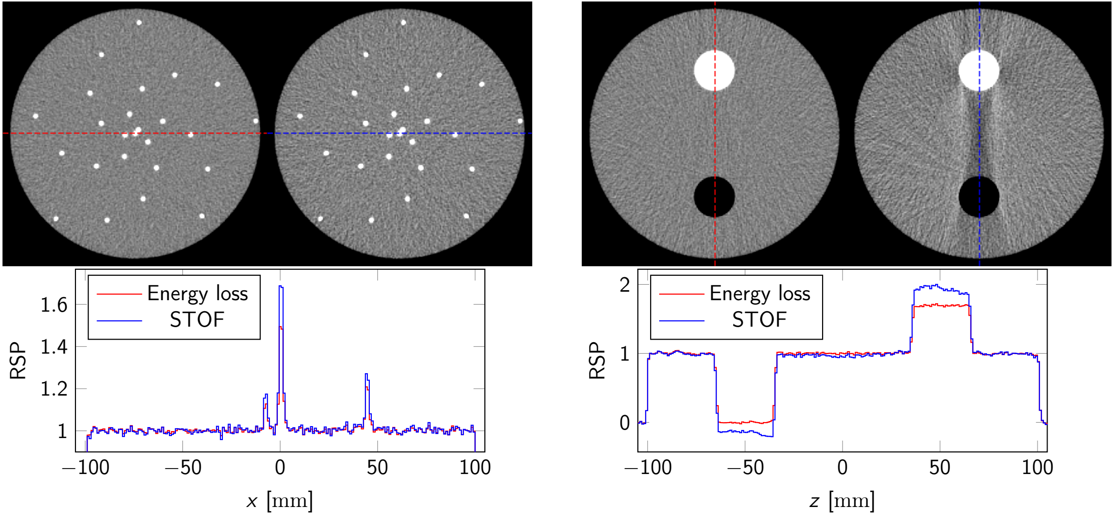

Energy loss
STOF
RSP
1.6
1.4
1.2
1
−100
−50
0
50
100
x [mm]
Energy loss
STOF
RSP
2
1
0
−100
−50
0
50
100
z [mm]

The abstract

Proton computed tomography reconstruction from sandwich time-of-flight using a lookup table approach

Proton computed tomography reconstruction from sandwich time-of-flight using a lookup table approach

Aurélien Coussat[1], Nils Krah[1], Jean Michel Létang[1], and Simon Rit[1]

[1]INSA-Lyon, Université Claude Bernard Lyon 1, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69373, Lyon, France

Abstract
Proton computed tomography (pCT) is an imaging modality which offers numerous advantages for proton therapy, including more precise estimation of the proton stopping power map compared to its current estimation based on X-ray computed tomography (CT). List-mode pCT scanners measure the energy loss of each proton to estimate the integral of the proton stopping power along its path. However, there is currently no list-mode pCT scanner which is fast enough for clinical use. A time-of-flight (TOF) pCT scanner was recently proposed as a potential solution, referred to as *sandwich TOF*, where the stopping power integral must be estimated using the TOF of protons between two detectors located before and after the scanned object along the proton beam. The advantage of this kind of set-up is its compactness. Conversion from sandwich TOF to the line integral of the relative proton stopping power is not trivial as the proton speed and energy change differently depending on the sought stopping power of the traversed materials. In this work, we propose a simple reconstruction procedure using a lookup table approach. The method is assessed using Monte Carlo simulations of several objects and compared to images obtained with energy-based pCT.

1 Introduction

Proton therapy is a radiotherapy technique that leverages the localized dose deposition of protons in matter to target deep-seated volumes while minimizing damages to surrounding healthy tissues. The location of the end of the proton beam range, where it deposits the maximum energy and known as the Bragg peak, depends on the relative stopping power (RSP) of the traversed materials. As a consequence, treatment planning for proton therapy requires as input a map of the RSP distribution within the patient's body. This distribution is currently determined by converting Hounsfield units obtained from an X-ray computed tomography (CT) to RSP values. However, this conversion introduces uncertainties ranging from 1.6 % to 5 % [1], resulting in increased safety margins around the treated volume and higher normal tissue complication probability.

Proton computed tomography (pCT) is an imaging method that can be used to directly characterize the RSP distribution in the treatment room. In pCT, protons travel through the patient and lose energy depending on the traversed tissues. The residual energy of each proton is then measured using an energy detector located after the patient, and then converted into the corresponding water-equivalent path length (WEPL) which is a line integral of the sought RSP. The RSP distribution can then be recovered with a tomographic reconstruction algorithm.

Protons undergo multiple deflections when traversing the patient due to multiple Coulomb scattering and nuclear interactions, and hence follow curved trajectories, which limits the spatial resolution. List-mode pCT scanner designs address this issue by tracking each proton individually using two detectors located before (upstream) and after (downstream) the scanned object. The measured positions and directions are then used to estimate the most-likely path of each proton individually. To keep the acquisition time reasonable, data should be acquired at a rate of at least a few MHz which has not been achieved so far.

A potential solution is to infer the energy loss indirectly from the times-of-flight (TOFs) of the protons between two detectors located after the object as the proton speed depends on the proton energy [2]. However, placing these two detectors with sufficient distance is impractical and an alternative solution [3], referred to as "sandwich TOF (STOF) pCT" in the following, is to place two 4D-tracking detectors that sandwich the object to estimate both the path (3D) and the TOF (+1D) of each proton. The TOF depends on the materials traversed by the proton, therefore the WEPL should be directly inferable from the TOF, removing the need for a residual energy detector. The 4D-tracking detectors could be implemented using low gain avalanche detectors as they are capable of sufficient timing and spatial resolutions [4]. One of the many challenges in implementing STOF pCT stems from the fact that protons traversing different materials with equivalent WEPLs result in equal energy losses, but significantly different TOFs, hence the conversion between WEPL and TOF is non-bijective. Consequently, the forward problem is non-linear and non-trivial to inverse, and the solution might be non-unique.

This work investigates the image quality, limitations, and artifacts of STOF when using a naive reconstruction method. The intention is to create a baseline for more advanced reconstruction techniques in the future. We are aware of the simplifying assumptions that restrict the set of objects for which the RSP distribution can be accurately reconstructed. More precisely, a Monte Carlo simulation of STOF pCT was implemented (Section 2.1), TOF and energy loss data were fitted to WEPL to convert measurements to WEPL (Section 2.2) and several phantoms were reconstructed (Section 2.3). The fits and the reconstructions are displayed in Section 3. The results and outcomes are discussed in Sections 4 and 5, respectively.

Outline

The idea

- Use the convex hull of the object to account for the air around the object
- A two-dimensional calibration curve is constructed on two variables:
  1. the TOF between the beginning of the object and the downstream detectors ($\tau_{\text{object}} + \tau_2$)
  2. the amount of air between the object and the downstream detector (d_2)

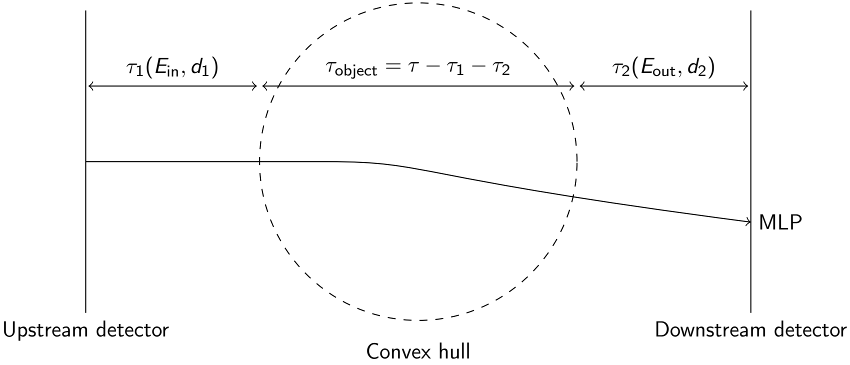

The fit

The Monte Carlo is repeated for…

- 10 realisations of d_2 from 0 to 220 mm
- 10 realisations of L from 0 to $220 - d_2$

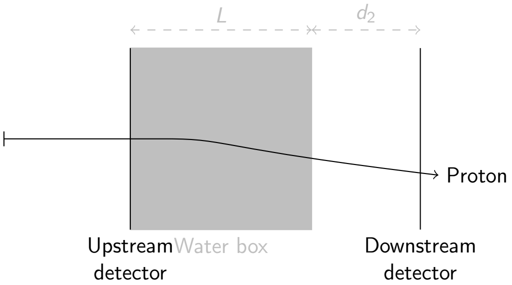

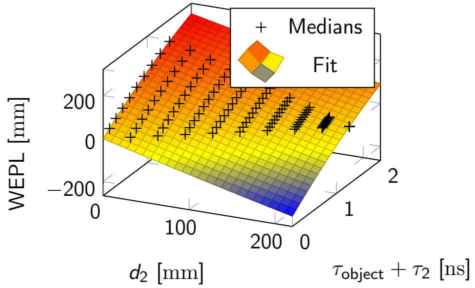

Reconstructions

Previous method:

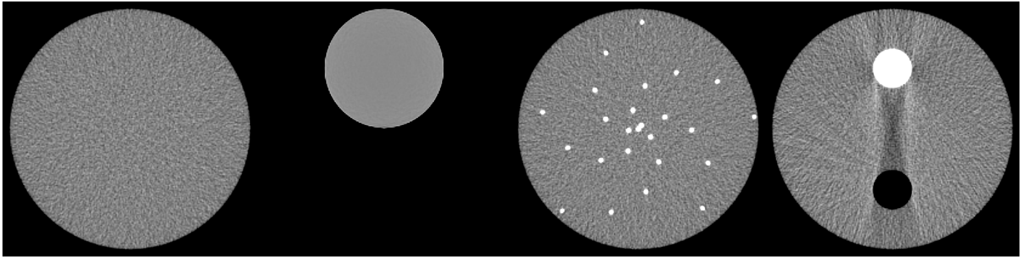

New method:

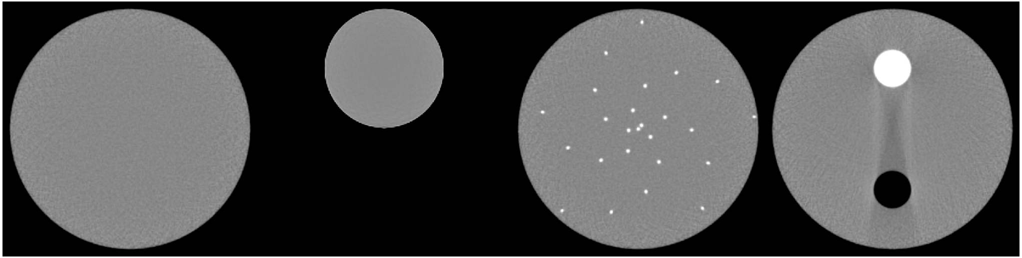

Outline

- Medical imaging in a few words
- Energy-loss proton computed tomography
  - Motivation
  - Scanner design
- Time-of-flight proton computed tomography
- Sandwich Time-of-flight proton computed tomography
  - First proposal: 1D calibration curve
  - Second proposal: 2D calibration curve
  - Third proposal: formulation as an optimization problem
- Conclusion

Forward model

Continuous version

- TOF of proton p:

$$\tau_p = \int_{\Gamma_p} \frac{\mathrm{d}\boldsymbol{x}}{v_p(\boldsymbol{x})} \tag{1}$$

- Velocity of the proton p at point $\boldsymbol{x}$:

$$v_p(\boldsymbol{x}) = c^2 \frac{E_p(\boldsymbol{x})}{E_p(\boldsymbol{x}) + mc^2} \sqrt{1 + 2\frac{mc^2}{E_p(\boldsymbol{x})}} \tag{2}$$

- Energy of the proton p at point $\boldsymbol{x} = \Gamma_p(s_k)$:

$$E_p(\Gamma_p(s)) = E_{\text{in}} - \int_{-\infty}^{s} \rho(\Gamma_p(s'))\sigma_{\mathrm{H_2O}}(E_p(\Gamma_p(s')))\mathrm{d}s' \tag{3}$$

Forward model

Continuous version

- TOF of proton p:

$$\tau_p = \int_{\Gamma_p} \frac{\mathrm{d}\boldsymbol{x}}{v_p(\boldsymbol{x})} \tag{1}$$

- Velocity of the proton p at point $\boldsymbol{x}$:

$$v_p(\boldsymbol{x}) = c^2 \frac{E_p(\boldsymbol{x})}{E_p(\boldsymbol{x}) + mc^2} \sqrt{1 + 2\frac{mc^2}{E_p(\boldsymbol{x})}} \tag{2}$$

- Energy of the proton p at point $\boldsymbol{x} = \Gamma_p(s_k)$:

$$E_p(\Gamma_p(s)) = E_{\text{in}} - \int_{-\infty}^{s} \rho(\Gamma_p(s'))\sigma_{\mathrm{H_2O}}(E_p(\Gamma_p(s')))\mathrm{d}s' \tag{3}$$

Discrete version

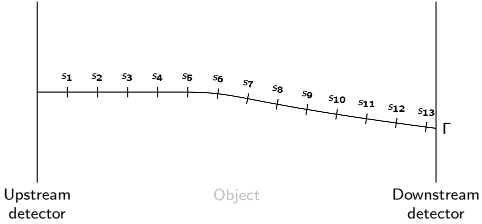

Forward model

Continuous version

- TOF of proton p:

$$\tau_p = \int_{\Gamma_p} \frac{\mathrm{d}\boldsymbol{x}}{v_p(\boldsymbol{x})} \tag{1}$$

- Velocity of the proton p at point $\boldsymbol{x}$:

$$v_p(\boldsymbol{x}) = c^2 \frac{E_p(\boldsymbol{x})}{E_p(\boldsymbol{x}) + mc^2} \sqrt{1 + 2\frac{mc^2}{E_p(\boldsymbol{x})}} \tag{2}$$

- Energy of the proton p at point $\boldsymbol{x} = \Gamma_p(s_k)$:

$$E_p(\Gamma_p(s)) = E_{\text{in}} - \int_{-\infty}^{s} \rho(\Gamma_p(s'))\sigma_{\mathrm{H_2O}}(E_p(\Gamma_p(s')))\mathrm{d}s' \tag{3}$$

Discrete version

- TOF of proton p:

$$\tau_p \approx \sum_{k=1}^{K} \frac{d_p^{s_k}}{v_p^{s_k}} \tag{4}$$

- Velocity of the proton p at point $\boldsymbol{x}$:

$$v_p^{s_k} = c^2 \frac{E_p^{s_k}}{E_p^{s_k} + mc^2} \sqrt{1 + 2\frac{mc^2}{E_p^{s_k}}} \tag{5}$$

- Energy of the proton p at point $\boldsymbol{x} = \Gamma_p(s_k)$:

$$E_p^{s_k} = E_p^{s_{k-1}} - d_p^{s_k}\tilde{\rho}_\rho(\Gamma_p(s_{k-1}))\sigma_{\mathrm{H_2O}}(E_p^{s_{k-1}}) \tag{6}$$

Optimization scheme

- Let $\hat{\boldsymbol{\tau}}$ the TOFs produced by the forward model, and $\boldsymbol{\tau}$ the measurements
- The cost function can be defined as

$$\mathcal{L}(\boldsymbol{\rho}) = \|\hat{\boldsymbol{\tau}}(\boldsymbol{\rho}) - \boldsymbol{\tau}\|_2^2 \tag{7}$$

Minimization of the cost function

$$\mathcal{L}(\boldsymbol{\rho}) = \|\hat{\boldsymbol{\tau}}(\boldsymbol{\rho}) - \boldsymbol{\tau}\|_2^2$$

- The cost function can be minimized using gradient descent
- In its simplest form, the update step is given by $\rho_i' = \rho_i - \alpha \dfrac{\partial \mathcal{L}}{\partial \rho_i}$
- The derivative can be computed using
  - Finite differences
  - Analytical differentiation
  - Automatic differenciation

Calculating the derivative

Finite differences

$$\frac{\partial \mathcal{L}}{\partial \rho_i} = \lim_{h \to 0} \frac{\mathcal{L}\left(\begin{bmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_i + h \\ \vdots \\ \rho_J \end{bmatrix}\right) - \mathcal{L}\left(\begin{bmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_i \\ \vdots \\ \rho_J \end{bmatrix}\right)}{h} \tag{8}$$

Calculating the derivative

Finite differences

$$\frac{\partial \mathcal{L}}{\partial \rho_i} = \lim_{h \to 0} \frac{\mathcal{L}\left(\begin{bmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_i + h \\ \vdots \\ \rho_J \end{bmatrix}\right) - \mathcal{L}\left(\begin{bmatrix} \rho_0 \\ \rho_1 \\ \vdots \\ \rho_i \\ \vdots \\ \rho_J \end{bmatrix}\right)}{h} \tag{8}$$

Super slow but useful nonetheless to check other implementations!

Calculating the derivative

Analytical differenciation

By repeatedly applying the chain rule:

- $$\frac{\partial \mathcal{L}}{\partial \rho_i} = \left(\frac{\partial \mathcal{L}}{\partial \hat{\boldsymbol{\tau}}}\right)^\top \frac{\partial \hat{\boldsymbol{\tau}}}{\partial \rho_i}$$
- $$\frac{\partial \hat{\tau}_p}{\partial \rho_i} = -\sum_{k=1}^{K} \frac{d_p^{s_k}}{(v_p^{s_k})^2} \frac{\partial v_p^{s_k}}{\partial \rho_i}$$
- $$\frac{\partial v_p^{s_k}}{\partial \rho_i} = c^2 \left(\frac{mc^2}{(E_p^{s_k} + mc^2)^2} \sqrt{1 + 2\frac{mc^2}{E_p^{s_k}}} + \frac{mc^2}{\sqrt{1 + 2\frac{mc^2}{E_p^{s_k}}}} \frac{E_p^{s_k}}{E_p^{s_k} + mc^2} \right) \frac{\partial E_p^{s_k}}{\partial \rho_i}$$
- $$\frac{\partial E_p^{s_k}}{\partial \rho_i} = \frac{\partial E_p^{s_{k-1}}}{\partial \rho_i} - d_p^{s_k} \left(\frac{\partial \tilde{\rho}_\rho(\Gamma_p(s_{k-1}))}{\partial \rho_i} \sigma_{\mathrm{H_2O}}(E_p^{s_{k-1}}) + \tilde{\rho}_\rho(\Gamma_p(s_{k-1})) \frac{\partial \sigma_{\mathrm{H_2O}}(E_p^{s_{k-1}})}{\partial E_p^{s_{k-1}}} \frac{\partial E_p^{s_{k-1}}}{\partial \rho_i} \right)$$

Calculating the derivative

Analytical differenciation

$E_p \leftarrow E_{\text{in}}$ ▷ Current energy

$\dfrac{\partial E_p}{\partial \rho_i} \leftarrow 0$ ▷ Current derivative of the energy

$\dfrac{\partial \hat{\tau}_p}{\partial \rho_i} \leftarrow 0$ ▷ Current TOF derivative

for $k \leftarrow 2, \dots, K$ **do**

$$\frac{\partial E_p}{\partial \rho_i} \leftarrow \frac{\partial E_p}{\partial \rho_i} - d_p^{s_k} \left(\phi_i(\Gamma_p(s_k))\sigma_{\mathrm{H_2O}}(E_p) + \tilde{\rho}_\rho(\Gamma_p(s_{k-1})) \frac{\partial \sigma_{\mathrm{H_2O}}(E_p)}{\partial E_p} \frac{\partial E_p}{\partial \rho_i} \right)$$

$$\frac{\partial v_p}{\partial \rho_i} \leftarrow c^2 \left(\frac{mc^2}{(E_p + mc^2)^2} \sqrt{1 + 2\frac{mc^2}{E_p}} + \frac{mc^2}{\sqrt{1 + 2\frac{mc^2}{E_p}}} \frac{E_p}{E_p + mc^2} \right) \frac{\partial E_p}{\partial \rho_i}$$

$\dfrac{\partial \hat{\tau}_p}{\partial \rho_i} \leftarrow \dfrac{\partial \hat{\tau}_p}{\partial \rho_i} - \dfrac{d_p^{s_k}}{v_p^2} \dfrac{\partial v_p}{\partial \rho_i}$ ▷ Update TOF derivative

$E_p \leftarrow E_p - d_p^{s_k} \tilde{\rho}_\rho(\Gamma_p(s_{k-1}))\sigma_{\mathrm{H_2O}}(E_p^{s_{k-1}})$ ▷ Update current energy

end for

return $\dfrac{\partial \hat{\tau}_p}{\partial \rho_i}$

Calculating the derivative

Automatic differenciation

- The chain rule can be automatically applied
- PyTorch's `autograd` does so by building a graph representing the computation

Calculating the derivative

Automatic differenciation

- The chain rule can be automatically applied
- PyTorch's `autograd` does so by building a graph representing the computation

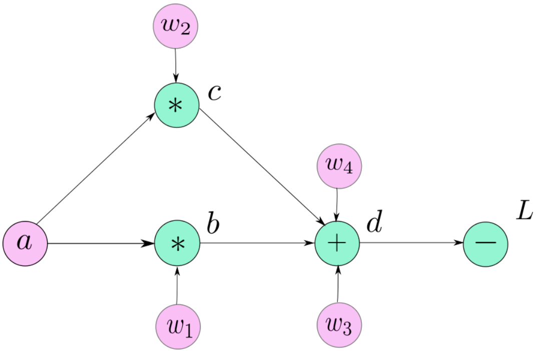

https://www.digitalocean.com/community/tutorials/pytorch-101-understanding-graphs-and-automatic-differentiation

```python
import torch
a = 5.
w = torch.tensor([1., 2., 3., 4.],
↪   requires_grad=True)
b = a * w[0]
c = a * w[1]
d = b + c + w[2] + w[3]
L = -d
print(L)  # tensor(-22.,
↪   grad_fn=<NegBackward0>)
L.backward()
print(w.grad)  # tensor([-5., -5.,
↪   -1., -1.])
```

Calculating the derivative

Automatic differenciation

- The chain rule can be automatically applied
- PyTorch's `autograd` does so by building a graph representing the computation

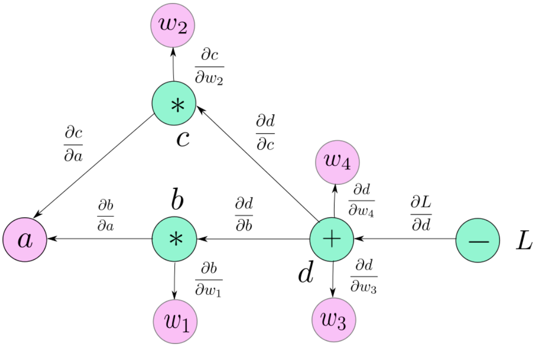

https://www.digitalocean.com/community/tutorials/pytorch-101-understanding-graphs-and-automatic-differentiation

```python
import torch
a = 5.
w = torch.tensor([1., 2., 3., 4.],
↪  requires_grad=True)
b = a * w[0]
c = a * w[1]
d = b + c + w[2] + w[3]
L = -d
print(L)  # tensor(-22.,
↪  grad_fn=<NegBackward0>)
L.backward()
print(w.grad)  # tensor([-5., -5.,
↪  -1., -1.])
```

Results

- Number of projections: 90
- Protons per projection: about 6000
- Reconstruction grid: $220 \times 4 \times 220$
- Voxel size: $1\,\mathrm{mm} \times 100\,\mathrm{mm} \times 1\,\mathrm{mm}$
- Number of iterations: 100
- Step size: $2\,\mathrm{mm}$
- Optimizer: Adam
- Color scale: $[0.75, 1.25]$

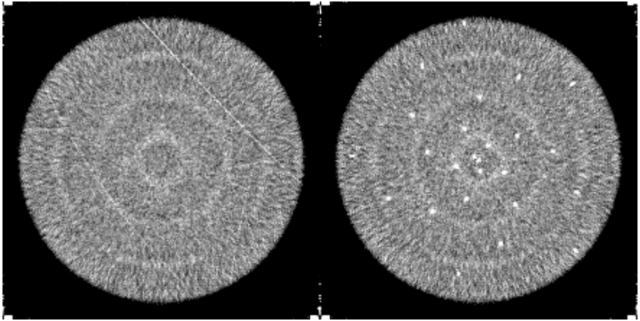

Outline

Energy-loss, TOF and sandwich TOF pCT

Energy-loss pCT

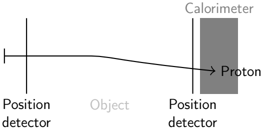

- Energy is detected using a calorimeter
- Hard to keep a reasonable acquisition time

TOF pCT[8]

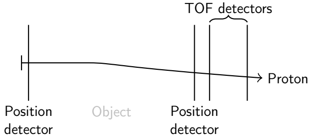

- Energy is estimated using TOF between two downstream detectors
- Placing TOF detectors with sufficient distance is impractical

STOF pCT[9]

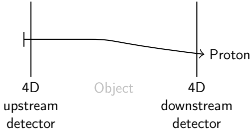

- WEPL is directly estimated from the TOF between up- and downstream detectors
- Can be implemented with two 4D low gain avalanche detectors

[8] Nils Krah et al. In: *Physics in Medicine & Biology* 67.16 (2022)
[9] F. Ulrich-Pur et al. In: *Journal of Instrumentation* 18.02 (2023)

Conclusion

- Reconstruction from TOF pCT data is challenging
- We proposed methods based on calibration curves and on direct optimization of the image
- Proposed methods are contributed to the PCT toolkit[10]

[10] https://github.com/RTKConsortium/PCT

Conclusion

- Reconstruction from TOF pCT data is challenging
- We proposed methods based on calibration curves and on direct optimization of the image
- Proposed methods are contributed to the PCT toolkit[10]

Dziękuję!

[10] https://github.com/RTKConsortium/PCT

Backup slides

Computation costs

	Speed	**Memory usage**
`autograd`	16 h	20 GB
`autograd` + GPU	1 h	20 GB
`autograd` + functions	16 h	10 GB
`autograd` + GPU + functions	1 h	10 GB

Why not have the whole forward model as a single `autograd` operator?

How to compute $\frac{\partial \rho_i}{\partial \hat{\tau}_p}$?

autograd custom functions

- Possibility to replace some parts of the gradient graph with custom analytical derivatives
- Some advantages:
  - Reduced memory footprint
  - Reduced computation time (?)
  - Validation of the analytical derivatives using `torch.autograd.gradcheck`
- Main drawback: additional complexity

```
class Velocity(torch.autograd.Function):

    @staticmethod
    def forward(ctx, e):
        ctx.save_for_backward(e)
        return c * (e / (e + m0)) *
        ↪  torch.sqrt(1 + 2 * (m0 / e))

    @staticmethod
    def backward(ctx, grad_output):
        e, = ctx.saved_tensors
        dv_de = -c*e*torch.sqrt(1 +
        ↪  2*m0/e)/(e + m0)**2 +
        ↪  c*torch.sqrt(1 + 2*m0/e)/(e +
        ↪  m0) - c*m0/(e*torch.sqrt(1 +
        ↪  2*m0/e)*(e + m0))
        return grad_output * dv_de
```