

# Regularised patient-specific stopping power calibration for proton therapy planning based on proton radiographic images

Accepted manuscript

N. Krah<sup>1</sup>, V. Patera<sup>2,3</sup>, S. Rit<sup>1</sup>, A. Schiavi<sup>2,3</sup>, I. Rinaldi<sup>4,5</sup>

<sup>1</sup>Lyon University, CNRS, CREATIS UMR5220, Centre Léon Bérard, Lyon, France

<sup>2</sup>Department of Basic and Applied Sciences for Engineering at Sapienza University of Rome, Rome, Italy

<sup>3</sup>National Institute for Nuclear Physics, INFN Section of Rome, Rome, Italy

<sup>4</sup>CNRS/IN2P3 and Lyon 1 University, UMR 5822, Villeurbanne, France

<sup>5</sup>MAASTRO Clinic, Maastricht, The Netherlands

*nils.krah@creatis.insa-lyon.fr*

## Abstract

Proton transmission imaging has been proposed and investigated as imaging modality complementary to X-ray based techniques in proton beam therapy. In particular, it addresses the issue of range uncertainties due to the conversion of an X-ray patient computed tomography (CT) image expressed in Hounsfield Units (HU) to relative stopping power (RSP) needed as input to the treatment planning system. One approach to exploit a single proton radiographic projection is to perform a patient-specific calibration of the CT to RSP conversion curve by optimising the match between a measured and a numerically integrated proton radiography.

In this work, we develop the mathematical tools needed to perform such an optimisation in an efficient and robust way. Our main focus lies on set-ups which combine pencil beam scanning with a range telescope detector, although most of our methods can be employed in combination with other set-ups as well. Proton radiographies are simulated in Monte Carlo using an idealised detector and applying the same data processing chain used with experimental data. This approach allows us to have a ground truth CT-RSP curve to compare the optimisation results with.

Our results show that the parameters of the CT-RSP curve are strongly correlated when using a pencil beam based set-up, which leads to unrealistic variation in the optimised CT-RSP curves. To address this issue, we introduce a regularisation procedure which guarantees a plausible degree of smoothness in the optimised CT-RSP curves. We investigate three different methods to perform the numerical projection operation needed to generate a proton digitally reconstructed radiography. We find that the approximate and computationally faster method performs as well as the more accurate but more demanding method. We perform a Monte Carlo experiment based on a head and neck patient to evaluate the range accuracy achievable with the optimised CT-RSP curves and find an agreement with the ground truth expectation of better than 0.5%. Our results further indicate that the region in the patient in which the proton radiography is acquired does not necessarily have to correspond to the treatment volume to achieve this accuracy. This is important as the imaged region could be freely chosen, e.g. in order to spare organs at risk.

## 1 INTRODUCTION

Proton beam therapy offers high selectivity in dose deposition due to the protons' distinct depth dose profile, commonly referred to as Bragg curve. This feature potentially allows for high dose to the tumour while sparing healthy surrounding tissue (Durante & Paganetti 2016). To fully exploit this advantage, precise knowledge of the relative stopping power (RSP) of the patient tissue is needed to correctly predict the proton range in the treatment planning system (TPS), in the best case on the day of treatment.

Currently, proton treatment planning is performed based on X-ray computed tomography (CT) data typically acquired some days before the first delivered fraction. The X-ray CT is expressed in Hounsfield Unit (HU) and approximately converted to RSP (Schneider et al. 2005). The uncertainties associated with the conversion process require additional safety margins around the contoured treatment volume (Jäkel et al. 2001, Paganetti 2012b). Proton radiography and proton CT have already been investigated a long time ago (West & Sherwood 1972, Koehler 1968, Steward & Koehler 1973, Takada et al. 1988) and more recently been brought up in the context of proton beam therapy (Schneider & Pedroni 1995, Schulte et al. 2004). In analogy to X-ray based transmission imaging, protons are shot through the patient and an image is formed based on the residual energy or residual range measured by a suitable detector positioned behind the patient (Parodi 2014, Poludniowski et al. 2015, Johnson 2018). A full tomographic scan directly provides a volumetric RSP map of the patient and therefore bypasses any conversion. It could thereby allow for reduced safety margins (Arbor et al. 2015).

Schneider et al. (2005) proposed an approach based on a single proton radiography expressing the water equivalent thickness (WET) of the patient which is arguably easier to achieve in practice and exposes the patient to a lower dose than a full proton CT scan. The idea is not to circumvent the conversion, but to perform a patient-specific calibration of the CT-RSP curve by comparing a proton radiographic image with a proton digitally reconstructed radiography (pDRR). This optimised curve would then be used for the treatment planning instead of the clinical one.

More recently, the method has been studied in combination with two types of proton imaging set-ups. Doolan et al. (2015) used a single plane detector in double scattering beam delivery. In this case, the beam is gradually pulled back by a range modulator device while the pixel detector behind the patient records the signal created by many protons at once. On the other hand, in single proton tracking, the imaging set-up records information proton by proton and an estimate of individual proton trajectories through the patient can be derived, the so-called most likely path (Williams 2004, Schulte et al. 2008). Collins-Fekete et al. (2017) have investigated the patient-specific calibration in combination with this kind of system.

In this work, we consider a proton imaging set-up which combines pencil beam scanning with a range telescope (Rinaldi et al. 2013, 2014, Farace et al. 2016). It can be employed in clinical practice potentially using available quality assurance (QA) equipment and without further hardware modifications in the treatment room (Krah et al. 2018). It is therefore of interest to develop the mathematical tools necessary to perform a patient-specific CT-RSP calibration in combination with such a set-up. In particular, we develop a regularised optimisation procedure which, as it turns out, is necessary to obtain plausible calibration curves. We perform a Monte Carlo study to assess the accuracy and robustness of the proposed methods. Although we focused on a specific kind of imaging system, most of our methods can be adapted and applied to other proton imaging set-ups as well.

## 2 MATERIALS AND METHODS

### 2.1 Patient specific calibration - general principle

The CT-RSP conversion curve is typically modelled as a series of linear segments parametrised by a set of parameter tuples  $(H_\alpha, R_\alpha)$ , as illustrated in the centre of figure 1 (Schneider et al. 1996, Taasti et al. 2018). The patient-specific calibration optimises the values of  $R_\alpha$  by minimising the difference between a pDRR and a measured proton radiography. A brute force implementation would start from an initial guess for the CT-RSP curve, convert a planning CT (expressed in Hounsfield Units [HU]) to a three-dimensional (3D) RSP map, project it to produce the pDRR, compare with the proton radiography (e.g., by calculating the mean squared difference), update the set of  $R_\alpha$  values, and iterate again. We will present a method to separate the projection operation from the actual optimisation step.

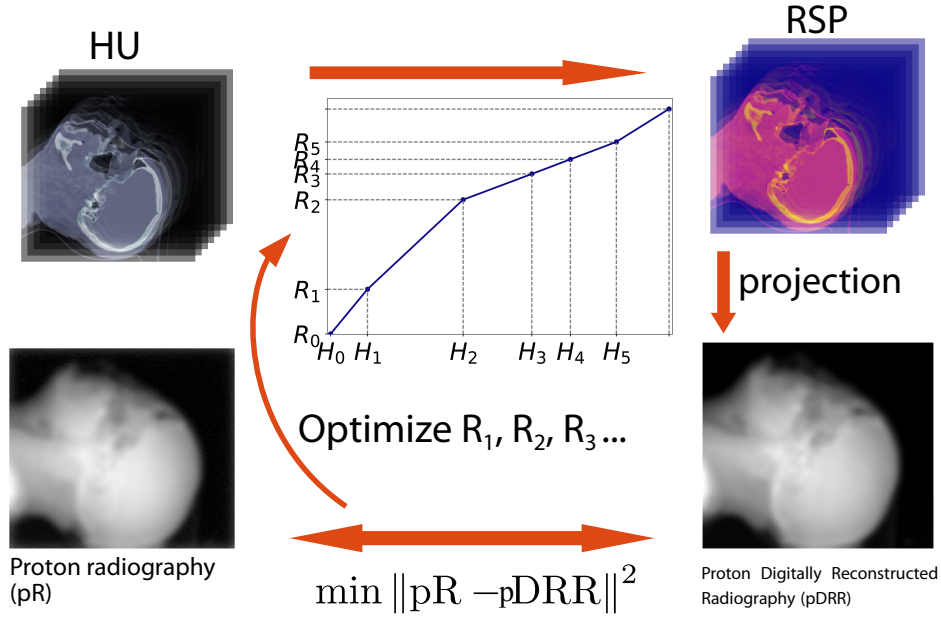


Figure 1: Optimisation scheme for the patient-specific CT-RSP calibration. The plot in the center is a schematic illustration of the CT-RSP curve defined as set of linear interpolations.

## 2.2 Framework of this study

Naturally, an important question to be asked would be: “How accurate does the patient-specific calibration method reproduce the HU to RSP relation of the patient’s tissue?” Answering this experimentally on the basis of phantom measurements would require precise knowledge of the phantom’s spatial material distribution. For realistically complex anthropomorphic or animal tissue phantoms, this information is difficult to have, although proton radiographic measurements have been performed in phantoms of known material composition in some studies (Farace et al. 2016, Wohlfahrt et al. 2018). Phantoms which are geometrically less complex than a patient, on the other hand, might lead to unrealistic conclusions. For this reason, we remained entirely within a Monte Carlo framework for our study where we could use a digital patient model as ground truth. It was derived by converting a patient CT into RSP before importing it into Monte Carlo (see section 2.4). The CT-RSP curve used for this conversion was therefore the ground truth reference curve and we evaluated how well it was recovered by the optimisation procedure. At no point do we attempt to directly compare the calibration outcome with some clinical curve. This would require quantifying or modelling the typical discrepancies between the true tissue properties and those predicted by a clinical curve, and it was not the scope of this work.

## 2.3 Proton radiographies

We used *Fred*, a GPU-accelerated Monte Carlo code developed at University Rome La Sapienza, to simulate proton radiographies. The code has been validated against other MC toolkits, Geant4 (Agostinelli et al. 2003, Allison et al. 2016) and FLUKA (Ferrari et al. 2005, Böhlen et al. 2014)), as well as experimental data from the CNAO proton therapy facility in Italy (Schiavi et al. 2017, ?) and, more recently, against measurements from the Krakow proton therapy facility (Garbacz et al. 2019). Its advantage for our study was the high computation speed.

A patient CT image was imported into *Fred* and to each voxel, an elemental composition, density, RSP, and radiation length was assigned as explained in section 2.4. No segmentation into tissue compartments was performed.

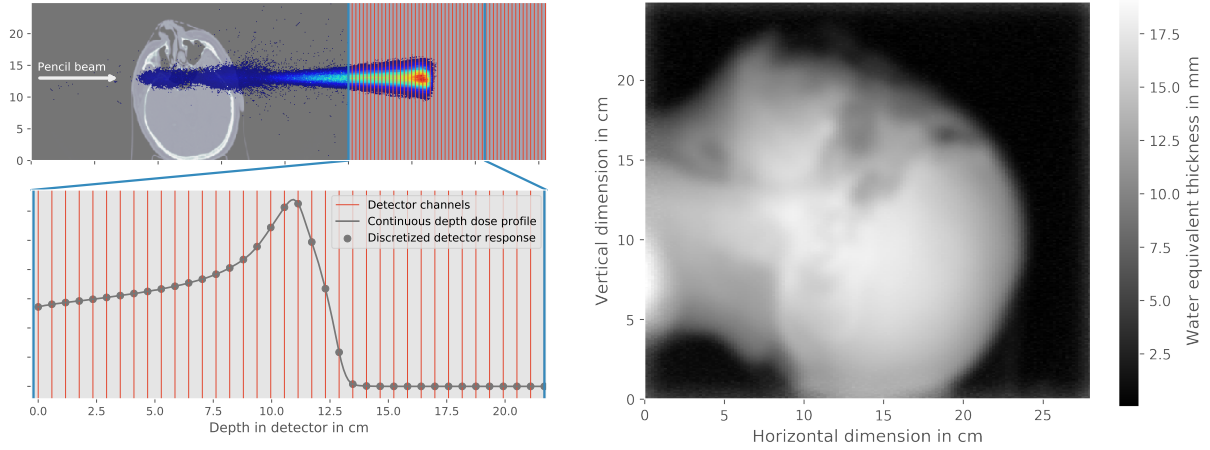


Figure 2: Left: Geometry used in Monte Carlo to simulate proton radiographies. The detector is idealised as solid block of water and the dose scored in 2.3 mm depth intervals. Right: Example of a simulated proton radiography.

We modelled the range telescope as a homogeneous box of water in which the integrated depth dose profiles were scored in depth increments of 2.3 mm. The geometry used in *Fred* is shown in figure 2 as well as a proton radiography of the patient used in this study.

The irradiation plan contained a two dimensional array of pencil beams with  $10^3$  protons per spot and a beam energy of 200 MeV, i.e., high enough so that the protons fully traversed the patient and partly penetrated the range telescope. For simplicity, we used a parallel beam geometry and kept the beam axis aligned with the voxel grid. Furthermore, we chose the beam spot positions to be centred on the voxel surfaces. The initial beam profile was assumed to be a symmetric Gaussian distribution characterised by the full width at half maximum (FWHM) parameter.

For each scanned beam spot, the range telescope recorded a (discretised) Bragg curve from which the water equivalent thickness (WET) was determined using a data processing procedure detailed in (Krah et al. 2018). From the list of beam spot coordinates and WET values, the proton radiography was constructed as 2D image so that each spot corresponded to one pixel in the proton radiography.

## 2.4 Conversion of CT to RSP, density, and elemental composition

*Fred* interpolates all parameters (density, RSP, radiation length, elemental composition) linearly between the materials provided in a lookup table according to the CT number in each voxel, i.e., it treats each tissue as a composition of two out of the given set of reference materials. *Fred* uses RSP, density, and radiation length to calculate electromagnetic energy loss and multiple Coulomb scattering (MCS), and the elemental composition to obtain cross sections of nuclear interactions (Schiavi et al. 2017). When simulating the proton radiographies, we used *Fred*'s default set of materials reported in table 1 which is essentially table 4 in Schneider et al. (1996) with an additional column containing the calculated radiation length  $X_0$ . Following the arguments in Sec. 2.2, we defined the relationship between CT numbers and RSP in this table as ground truth conversion curve. Some materials in the table refer to tissues which do not appear in a patient head and which would be disregarded when using the table in a TPS. For the purpose of providing a ground truth information, we kept them to have a relatively granular lookup table in the simulation.

We built an additional reference lookup table (table 2) using the material compositions and HU intervals reported in table 6 in Schneider et al. (2000). The density was determined using

Table 1: Material table from Schneider et al. (1996) used to simulate the proton radiographies. It contains the CT-RSP relationship which we used as ground truth knowledge.

	Elemental fraction in %														$\rho$ [g/cm <sup>3</sup> ]	CT No. [HU]	RSP	X <sub>0</sub> [cm]
	O	C	H	N	Ca	P	Na	Mg	S	Cl	K	Fe	I					
Air	23.5	0	0	76.5	0	0	0	0	0	0	0	0	0	0.0012	-1000	0.001062	30862.6	
Lung infl.	74.9	10.5	10.3	3.1	0	0.2	0.2	0	0.3	0.3	0.2	0	0	0.26	-741	0.257998	140.478	
Adip. tissue	27.8	59.8	11.4	0.7	0	0	0.1	0	0.1	0.1	0	0	0	0.95	-70	0.978975	43.3744	
Yellow marrow	23.1	64.4	11.5	0.7	0	0	0.1	0	0.1	0.1	0	0	0	0.98	-42	1.01303	42.5351	
Water	88.81	0	11.19	0	0	0	0	0	0	0	0	0	0	1	0	1	36.0823	
Breast	52.7	33.2	10.6	3	0	0.1	0.1	0	0.2	0.1	0	0	0	1.02	3	1.02898	37.8775	
Red marrow	43.9	41.4	10.5	3.4	0	0.1	0	0	0.2	0.2	0.2	0.1	0	1.03	14	1.04102	38.0317	
GI tract	75.1	11.5	10.6	2.2	0	0.1	0.1	0	0.1	0.2	0.1	0	0	1.03	23	1.02804	35.7108	
Lymph	83.2	4.1	10.8	1.1	0	0	0.3	0	0.1	0.4	0	0	0	1.03	28	1.02701	35.1365	
Testis	69.4	16.9	10.6	2.2	0	0.2	0.2	0	0.1	0.2	0.2	0	0	1.04	32	1.04104	35.7064	
Brain	71.2	14.5	10.7	2.2	0	0.4	0.2	0	0.2	0.3	0.3	0	0	1.04	37	1.04	35.3985	
Thyroid	74.5	11.9	10.4	2.4	0	0.1	0.2	0	0.1	0.2	0.1	0	0.1	1.05	40	1.04496	34.9107	
Muscle	71	14.3	10.2	3.4	0	0.2	0.1	0	0.3	0.1	0.4	0	0	1.05	42	1.04401	35.0533	
Kidney	72.4	13.2	10.3	3	0.1	0.2	0.2	0	0.2	0.2	0.2	0	0	1.05	43	1.04496	34.9957	
Lung defl.	74.9	10.5	10.3	3.1	0	0.2	0.2	0	0.3	0.3	0.2	0	0	1.05	44	1.04401	34.785	
Ovary	76.8	9.3	10.5	2.4	0	0.2	0.2	0	0.2	0.2	0.2	0	0	1.05	45	1.04601	34.7622	
Eye lens	64.6	19.5	9.6	5.7	0	0.1	0.1	0	0.3	0.1	0	0	0	1.07	50	1.06005	34.9361	
Liver	71.6	13.9	10.2	3	0	0.3	0.2	0	0.3	0.2	0.3	0	0	1.06	53	1.05396	34.6543	
Spleen	74.1	11.3	10.3	3.2	0	0.3	0.1	0	0.2	0.2	0.3	0	0	1.06	54	1.05396	34.5251	
Heart	74.5	11	10.2	3.3	0	0.1	0.1	0	0.2	0.3	0.2	0.1	0	1.06	55	1.053	34.5819	
Skin	64.5	20.4	10	4.2	0	0.1	0.2	0	0.2	0.3	0.1	0	0	1.09	75	1.08401	34.2994	
Cartilage	74.4	9.9	9.6	2.2	0	2.2	0.5	0	0.9	0.3	0	0	0	1.1	98	1.08097	32.4815	
Skel. spong.	36.7	40.4	8.5	2.8	7.4	3.4	0.1	0.1	0.2	0.2	0.1	0.1	0	1.18	260	1.15605	29.3503	
Skel. sacrum	43.8	30.2	7.4	3.7	9.8	4.5	0	0.1	0.2	0.1	0.1	0.1	0	1.29	413	1.23801	25.3857	
Vert. col. (D6,L3)	43.7	28.7	7	3.8	11.1	5.1	0	0.1	0.2	0.1	0.1	0.1	0	1.33	477	1.26696	24.0933	
Femur	36.8	34.5	7	2.8	12.9	5.5	0.1	0.1	0.2	0.1	0	0	0	1.33	499	1.26895	23.8711	
Ribs (2nd, 6th)	43.6	26.3	6.4	3.9	13.1	6	0.1	0.1	0.3	0.1	0.1	0	0	1.41	595	1.32907	22.0085	
Vert. col. (C4)	43.6	26.1	6.3	3.9	13.3	6.1	0.1	0.1	0.3	0.1	0.1	0.1	0	1.42	609	1.33693	21.7318	
Humerus	36.9	31.4	6	3.1	15.2	7	0.1	0.1	0.2	0	0	0	0	1.46	683	1.37006	20.8816	
Ribs (10th)	43.4	23.5	5.6	4	15.6	7.2	0.1	0.1	0.3	0.1	0.1	0	0	1.52	763	1.41299	19.6246	
Cranium	43.5	21.2	5	4	17.6	8.1	0.1	0.2	0.3	0	0	0	0	1.61	903	1.48007	17.9969	
Mandible	43.5	19.9	4.6	4.1	18.7	8.6	0.1	0.2	0.3	0	0	0	0	1.68	1006	1.53401	16.9642	
Cortical bone	43.5	15.5	3.4	4.2	22.5	10.3	0.1	0.2	0.3	0	0	0	0	1.92	1376	1.71398	14.0583	

the fit shown in their figure 9. The stopping power and radiation length values were calculated based on the elemental composition exploiting Geant4's functionality (Agostinelli et al. 2003). We used this table for some of the presented results to determine the regulariser weights (see section 2.10.1).

## 2.5 Notations

Throughout the paper, we will use the lower case variables  $h$  and  $r$  to refer to the HU and RSP values in the voxels, respectively. We use  $i, j, k$  as voxel indices and, by convention, associate  $k$  with the beam axis. With  $P$ , we refer to the proton radiography and with  $D$  to the pDRR. Both,  $P$  and  $D$ , share the same pixel indices denoted as  $i_p$  and  $j_p$ . It is sometimes convenient to combine the two indices  $i_p, j_p$  into a single index  $l$  so that, e.g.,  $P_l$  denotes the vector of pixels in the proton radiography. The  $N$  interval limits of the CT-RSP conversion curve are given as capital letter variables  $R_\alpha$  and  $H_\alpha$  with the Greek index enumerating the parametrisation intervals (see figure 1). We use the  $\vec{\cdot}$  symbol to indicate vector notation, e.g.  $\vec{R} = (R_0, R_1, \dots, R_N)$ .

## 2.6 Remapping of CT image

We exploit the piecewise linear character of the CT-RSP conversion curve to remap the CT numbers [in HU] in the CT image onto a set of coefficients. This allows us to state the optimisation as a linear minimisation problem. Let  $h$  be the HU value in one voxel of the patient CT image. Its conversion to  $r$  depends on a parameter tuple  $(H_\alpha, R_\alpha)$  if and only if  $h$  lies in the HU intervals left or right of  $H_\alpha$ , i.e.,  $h \in (H_{\alpha-1}, H_{\alpha+1})$ . Specifically,  $r$  is the linear barycentric interpolation between the parameter pair  $(R_\alpha, R_{\alpha+1})$  if  $h \in (H_\alpha, H_{\alpha+1})$  and between  $(R_{\alpha-1}, R_\alpha)$ , if  $h \in (H_{\alpha-1}, H_\alpha)$ .

This motivates the definition of a vector  $\vec{w}$  whose components  $w_\alpha$  are calculated in the following way:

Table 2: Lookup table according to Schneider et al. (2000) with additional columns containing RSP (at 200 MeV) and radiation length  $X_0$ .

O	C	H	N	Elemental fraction in %								$\rho$ [g/cm <sup>3</sup> ]	CT No. [HU]	RSP	$X_0$ [cm]
				Ca	P	Na	Mg	S	Cl	K	Ar				
23.2	0	0	75.5	0	0	0	0	0	0	0	1.3	0.000605	-1050	0.000537909	60521.5
74.9	10.5	10.3	3.1	0	0.2	0.2	0	0.3	0.3	0.2	0	0.052695	-950	0.0529346	693.293
19.8	68.1	11.6	0.2	0	0	0.1	0	0.1	0.1	0	0	0.907347	-120	0.934753	46.3824
30.8	56.7	11.3	0.9	0	0	0.1	0	0.1	0.1	0	0	0.943934	-82	0.967335	43.3455
41.1	45.8	11	1.5	0	0.1	0.1	0	0.2	0.2	0	0	0.97083	-52	0.989457	40.9452
50.9	35.6	10.8	2.2	0	0.1	0	0	0.2	0.2	0	0	0.997725	-22	1.01295	38.9291
57.8	28.4	10.6	2.6	0	0.1	0	0	0.2	0.2	0.1	0	1.02462	8	1.03652	37.2415
72.3	13.4	10.3	3	0	0.2	0.2	0	0.2	0.2	0.2	0	1.03056	19	1.0359	35.7234
62.2	20.7	9.4	6.2	0	0	0.6	0	0.6	0.3	0	0	1.09681	80	1.09214	34.0022
35.5	45.5	9.5	2.5	4.5	2.1	0.1	0	0.1	0.1	0.1	0	1.08746	120	1.0832	33.8555
36.3	42.3	8.9	2.7	6.4	3	0.1	0	0.1	0.1	0.1	0	1.13485	200	1.11908	31.2418
37.2	39.1	8.2	2.9	8.3	3.9	0.1	0	0.1	0.1	0.1	0	1.19409	300	1.16414	28.6198
38	36.1	7.6	3	10.1	4.7	0.1	0.1	0.2	0.1	0	0	1.25333	400	1.20957	26.3764
38.7	33.5	7.1	3.2	11.7	5.4	0.1	0.1	0.2	0	0	0	1.31257	500	1.25601	24.5068
39.4	31	6.6	3.3	13.2	6.1	0.1	0.1	0.2	0	0	0	1.37181	600	1.30145	22.8421
40	28.7	6.1	3.5	14.6	6.7	0.1	0.1	0.2	0	0	0	1.43104	700	1.34623	21.3871
40.5	26.5	5.6	3.6	15.9	7.3	0.1	0.2	0.3	0	0	0	1.49028	800	1.38998	20.0717
41.1	24.6	5.2	3.7	17	7.8	0.1	0.2	0.3	0	0	0	1.54952	900	1.43531	18.9582
41.6	22.7	4.9	3.8	18.1	8.3	0.1	0.2	0.3	0	0	0	1.60876	1000	1.4818	17.9454
42	21	4.5	3.9	19.2	8.8	0.1	0.2	0.3	0	0	0	1.668	1100	1.52574	17.0139
42.5	19.4	4.2	4	20.1	9.2	0.1	0.2	0.3	0	0	0	1.72724	1200	1.57143	16.2022
42.9	17.9	3.9	4.1	21	9.6	0.1	0.2	0.3	0	0	0	1.78648	1300	1.61658	15.4529
43.2	16.5	3.6	4.2	21.9	10	0.1	0.2	0.3	0	0	0	1.84572	1400	1.66118	14.7594
43.5	15.5	3.4	4.2	22.5	10.3	0.1	0.2	0.3	0	0	0	1.90496	1500	1.7082	14.1711

$$\begin{aligned}
w_\alpha(h) &\equiv \frac{h - H_{\alpha-1}}{H_\alpha - H_{\alpha-1}} \Omega_\alpha(h) + \frac{H_{\alpha+1} - h}{H_{\alpha+1} - H_\alpha} \Omega_{\alpha+1}(h) \\
w_0(h) &= \frac{H_1 - h}{H_1 - H_0} \Omega_{\alpha=1}(h) \\
w_N(h) &= \frac{h - H_{N-1}}{H_N - H_{N-1}} \Omega_{\alpha=N}(h)
\end{aligned} \tag{1}$$

with

$$\Omega_\alpha(h) = \begin{cases} 1 & \text{for } H_{\alpha-1} \leq h < H_\alpha \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

The value of  $w_\alpha$  thus quantifies how much  $r$  depends on the corresponding parameter  $R_\alpha$ . Note that most components of the vector  $\vec{w}$  are zero except for two (or one if  $h$  lies exactly on an interval boundary) and that the sum over all  $w_\alpha$  is always one. The RSP value  $r$  is obtained by contracting  $\vec{w}$  with the vector  $\vec{R}$  containing the parameters  $R_\alpha$  of the conversion curve:

$$r = \sum_{\alpha=0}^N w_\alpha R_\alpha = \vec{w} \cdot \vec{R}. \tag{3}$$

When applying the piecewise barycentric interpolation to all voxels in the CT image, the three dimensional HU matrix,  $h_{ijk}$ , is mapped onto a four dimensional coefficient matrix  $w_{ijk,\alpha}$ . In accordance with equation 3, the 3D RSP map is:

$$r_{ijk} = \sum_{\alpha=0}^N w_{ijk,\alpha} R_\alpha = \vec{w}_{ijk} \cdot \vec{R} \tag{4}$$

## 2.7 Proton digitally reconstructed radiography

The pDRR is obtained by projecting the RSP map along lines corresponding to the pencil beams used to acquire the proton radiography. We do not consider projection directions oblique to the CT grid in this work (see section 2.3). Denoting with  $\text{proj}_k$  the projection operation along axis

190  $k$ , the pDRR is calculated as

$$D_{i_p j_p} = \text{proj}_{k, i_p j_p}(r_{ijk}) = \text{proj}_{k, i_p j_p} \left( \sum_{\alpha=0}^N w_{ijk, \alpha} R_{\alpha} \right) = \sum_{\alpha=0}^N \text{proj}_{k, i_p j_p}(w_{ijk, \alpha}) R_{\alpha} = \sum_{\alpha=0}^N A_{i_p j_p, \alpha} R_{\alpha} \quad (5)$$

191 with

$$A_{i_p j_p, \alpha} = \text{proj}_{k, i_p j_p}(w_{ijk, \alpha}). \quad (6)$$

192 The projection operation is wrapped into the matrix  $A$  because it only acts on the coefficients  
 193  $w_{ijk}$ , but not on the parameters  $R_{\alpha}$ , and may therefore be interchanged with the contraction over  
 194 the  $\vec{R}$  vector. For an ideal pencil beam with FWHM=0 and neglecting scattering, the projection  
 195 is simply the sum along  $k$  multiplied by the voxel size  $s$ , i.e.,  $A_{i_p j_p, \alpha} = s \sum_k w_{i=i_p, j=j_p, k, \alpha}$ . A  
 196 more realistic projection operation must at least take into account the beam size and possibly  
 197 the effect of scattering. In this work, we investigated the following three projection models:  
 198 (1) Gaussian beam without scattering, (2) Gaussian beam with material and depth dependent  
 199 scattering, and (3) Monte Carlo based projection.

### 200 2.7.1 Gaussian beam without scattering

201 In this model, the beam is described by a Gaussian transverse profile characterised by the  
 202 FWHM, which is 5-10 mm for clinical proton beams at the energies used for proton imaging  
 203 ( $\approx 200$  MeV). The effect of scattering within the patient is neglected and the beam profile is  
 204 considered to be constant in depth and independent of the pencil beam position. Mathematically,  
 205 this corresponds to a convolution with a shift invariant Gaussian kernel  $G_{i_p - i, j_p - j}$  in the plane  
 206 orthogonal to the proton beam. The matrix  $A_{i_p j_p, \alpha}$  is then calculated as

$$A_{i_p j_p, \alpha} = s \sum_k \sum_{i, j} G_{i_p - i, j_p - j} w_{ijk, \alpha} = s \sum_{i, j} G_{i_p - i, j_p - j} \sum_k w_{ijk, \alpha}. \quad (7)$$

207 The convolution needs to be performed only once after summing along  $k$ , making the implemen-  
 208 tation computationally inexpensive.

### 209 2.7.2 Gaussian beam model with scattering

210 This more realistic projection model takes into account the effect of scattering which leads to  
 211 a gradual widening of the proton beam with depth. The most important contribution is due  
 212 to MCS, i.e., a succession of many small angle deflections as a proton propagates through the  
 213 medium. We take as depth dependent beam size a combination of the initial beam size and the  
 214 contribution due to MCS in the medium,

$$\sigma_{\text{beam}}(z_k) = \sqrt{\sigma_{\text{init}}^2 + \sigma_{\text{mcs}}^2(z_k)}, \quad (8)$$

215 because the Gaussian pencil beam can approximately be thought of as a weighted superposition  
 216 of many pin-like pencil beams, each of which undergoes MCS. We denote with  $z_k$  the depth of  
 217 the slices in the CT voxel grid along the beam axis.

218 We used a discretised version of the analytical formulae from (Lynch & Dahl 1991) and  
 219 (Gottschalk et al. 1993) to calculate  $\sigma_{\text{mcs}}^2$  as a function  $z_k$ :

$$\sigma_{\text{mcs}}^2(z_k) = \Omega_0^2 \left[ 1 + 0.038 \ln \left( \frac{z_k}{X_0} \right) \right] s \sum_{k'=0}^k \frac{z_{k'}^2}{\beta^2(z_{k'}) p^2(z_{k'}) X_0(z_{k'})}, \quad (9)$$

220 with  $\Omega_0 = 13.6$  MeV and  $s$  the voxel size. The radiation length  $X_0$  depends on the type  
 221 of material and parametrises the amount of MCS the protons suffer while traversing it. In  
 222 a heterogeneous medium like a patient, it depends on the three dimensional position within

the patient volume. The terms  $\beta$  and  $p$  are the proton's velocity relative to the speed of light,  $\beta = v/c$ , and its momentum, respectively. They relate to the proton energy as  $\beta p = (E^2 + 2E E_p)/(E + E_p)c$ , where  $E_p$  is the proton's rest mass, and decrease as a proton continuously slows down while traversing the patient.

Evaluating equation 9 requires some prior knowledge about the material composition within the patient which we obtained by converting the patient CT to an RSP and  $X_0$  map through the default tables implemented in *Fred*. Alternatively, one could perform a first optimisation of the CT-RSP curve using the Gaussian model without scattering, use the result to evaluate equation 9, and optimise the CT-RSP curve again, this time using the Gaussian beam model with scattering.

We make the following approximations to be able to sum along the beam axis  $z$  only in equation 9. For a given pencil beam, we determine the mean radiation length the protons encounter in each slice of the CT volume by calculating the weighted average value of  $X_0$  around the beam center  $(i_p, j_p)$  in the  $(i, j)$  plane at depth  $z_k$ . We use a two dimensional Gaussian weighting kernel with a standard deviation reflecting the beam spot size. This yields the depth dependent term  $X_0(z_k)$  in equation 9. Similarly, we calculate the water equivalent thickness  $z_{WET}$  that the protons belonging to a given pencil beam have traversed on average when penetrating the patient up to some geometric depth  $z_k$ . We then use the expression  $E(z_{WET}) = ((R - z_{WET})/\chi)^{1/\xi}$  from Bortfeld & Schlegel (1996) to calculate the average proton energy in depth  $z_k$ . In this formula,  $R$  refers to the proton range in water and  $\chi$  and  $\xi$  are fit constants.

We determine the approximate beam size for each pencil beam  $(i_p, j_p)$  and in all CT slices  $k$  as explained above to build a space variant Gaussian convolution kernel  $G_{i_p j_p; ij}(z_k)$  of width  $\sigma_{beam, i_p j_p}(z_k)$ . The matrix  $A_{i_p j_p, \alpha}$  is calculated as

$$A_{i_p j_p, \alpha} = s \sum_k \sum_{i, j} G_{i_p j_p; ij}(z_k) w_{ijk, \alpha}. \quad (10)$$

Fast Fourier transform based convolution algorithms cannot be used here because the kernel is shift variant. Instead, the multiplication and summation operations in equation 10 must be executed explicitly for all HU intervals  $\alpha$ . We therefore implemented equation 10 on an Nvidia Titan X GPU card using the OpenCL library. This helped lowering the computation time from several hours per projection down to about 15 sec.

### 2.7.3 Monte Carlo based projection

We used Monte Carlo to perform the projection in equation 5 to avoid the approximations inherent in the analytical beam model and to include the effect of nuclear interactions on the proton trajectories. To this end, we implemented a dedicated plugin to *Fred* which iterates through the following scheme: Initially, the patient CT is loaded and the components of the matrix  $\omega_{ijk, \alpha}$  are calculated. The pencil beams are defined and arranged in the same way as in the proton radiography simulation (see Sec. 2.3). The protons are then transported across the patient CT volume using the physics models of the MC code. After each transport step (before traversing into a new voxel the latest), the coefficient vector  $\vec{\omega}$  is extracted from the current voxel and multiplied by the step length projected onto the beam axis  $z$ . This information is cumulated along the entire trajectory and for all protons within one pencil beam. The result is the integrated coefficient vector  $\vec{\omega}$  associated with one pencil beam and repeating the procedure over the entire array of pencil beams yields the matrix  $A_{i_p j_p, \alpha}$ .

If the number of protons per pencil beam remained constant across the patient volume, the integrated coefficient vector would simply be normalised by this number. In reality, the number of protons fluctuates as a function of depth because (1) secondary protons are created through inelastic scattering or knock-on elastic collisions, (2) protons stop inside the patient because they lose a large fraction of their energy in a nuclear interaction, and (3) some protons are



scattered at large Rutherford-like angles and exit the patient laterally without fully traversing it. Our plugin accounts for these effects by applying a depth dependent normalisation and only includes protons which would have reached the detector. In the current implementation of *Fred*, the plugin can only be executed in a single thread on a CPU which makes the computation time much longer (order of hours) than for the analytical projection model.

## 2.8 Linear minimisation problem

Henceforth, we use a single index  $l$  for the pixels in the proton radiography (see section 2.5). We define the data attachment term of the cost function as

$$E_{\text{data}}(\vec{R}) = \sum_l \left( P_l - \sum_{\alpha=0}^N A_{l\alpha} R_\alpha \right)^2 = \|\vec{P} - A\vec{R}\|^2 \quad \text{with} \quad A_{l\alpha} = \text{proj}_k(w_{lk,\alpha}), \quad (11)$$

where  $\|\cdot\|^2$  refers to the  $\ell$  norm. We have introduced compact matrix/vector notation on the right hand side.

The optimal parameters of the CT-RSP conversion (denoted with the  $\hat{\cdot}$  symbol) curve are the components of the vector  $\hat{\vec{R}}$  which minimises the objective function, i.e.,

$$\hat{\vec{R}} = \text{argmin} \left\{ E(\vec{R}) \right\}. \quad (12)$$

Note that the HU intervals  $H_\alpha$  are implicitly encoded in the matrix  $A$  and not subject to optimisation, i.e. the parameters  $\hat{\vec{R}}$  are optimal for a given discretisation of the CT-RSP curve. Optimising  $H_\alpha$  and  $R_\alpha$  alongside each other would mean to optimise both  $A$  and  $\vec{R}$  in equation 11 - a much harder problem to solve, if solvable at all. In practice, the intervals  $H_\alpha$  can be chosen to yield a desired degree of granularity (e.g. smaller intervals in soft tissue region, fewer in bone region), or in accordance with a clinical calibration curve already in use in the facility. In this work, we used the same intervals as in the ground truth curve except for one dedicated test where different intervals were used in the optimization than in the proton radiography simulation (figure 4 in section 3.1).

## 2.9 Parameter correlation coefficients

The minimisation problem in equation 12 is mathematically well posed if matrix  $A$  has full rank, or, in practical terms, if there are at least as many pixels in the proton radiography as there are  $R_\alpha$  parameters to be optimised. Even if this criterion is met, the optimised CT-RSP curve would not be plausible from a physics point of view in many cases, as we will show in section 3.1. The reason is the following: The WET value encoded in a single pixel of the pDRR results from a summation of RSP values in many voxels along and around the proton beam (see section 2.7). Increasing the value of a parameter  $R_\alpha$  increases the RSP in some of the contributing voxels and consequently the average WET. At the same time, lowering the value of another parameter  $R_{\alpha'}$  tends to lower the WET value. Therefore, the  $R$  parameters are expected to be correlated. We also note that neighboring  $R_\alpha$  parameters are in any case anti-correlated because the CT-RSP curve interpolates linearly between them and is continuous by construction. The correlation coefficient of the parameters  $R_\alpha$  can be determined from the covariance matrix of the cost function as follows:

$$\rho_{\alpha\alpha'} = \frac{\text{Cov}_{\alpha\alpha'}}{\sqrt{\text{Cov}_{\alpha\alpha} \text{Cov}_{\alpha'\alpha'}}} \quad \text{with} \quad \text{Cov}_{\alpha\alpha'} = (A^T A)^{-1}_{\alpha\alpha'}. \quad (13)$$

The coefficient values always lie within the  $[-1, 1]$  interval and  $\rho_{\alpha\alpha} = 1$  by construction.

## 2.10 Regularised optimisation

As we will show in the results, correlation among the parameters  $R_\alpha$  would lead to unrealistic optimised CT-RSP curves. Therefore, we introduce a regularisation term in the cost function to control the variation of the CT-RSP curve. It penalises the intrinsic variation of the patient-specific CT-RSP curve. We define

$$E_{\text{reg}} = \sum_{\alpha=\alpha_0}^{\alpha_M} (d_{\alpha+1} - d_\alpha)^2, \quad (14)$$

with the discrete derivative  $d_\alpha$  of the curve calculated as

$$d_\alpha = \frac{R_\alpha - R_{\alpha-1}}{H_\alpha - H_{\alpha-1}}. \quad (15)$$

The regularisation may be applied to a subset of HU intervals only (see section 2.10.1) and  $\alpha_0$  and  $\alpha_M$  refer to the first and last of them, respectively.

Equation 14 can be written as matrix equation,

$$E_{\text{reg}} = \sum_{\alpha=\alpha_0}^{\alpha_M} (A_{\alpha'\alpha}^{\text{reg}} R_\alpha)^2 = \|A^{\text{reg}} \vec{R}\|^2, \quad (16)$$

where the matrix components are given as

$$A_{\alpha'\alpha}^{\text{reg}} = \frac{1}{H_{\alpha+1} - H_{\alpha-1}} (\delta_{\alpha',\alpha-1} - 2\delta_{\alpha',\alpha} + \delta_{\alpha',\alpha+1}), \quad \text{with } \delta_{\alpha',\alpha} = \begin{cases} 1 & \text{for } \alpha' = \alpha \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

### 2.10.1 Regularised cost function

The conversion curve is expected to be rather smooth and monotonic in the regions  $\text{HU} < -100$  and  $\text{HU} > 100$ , but should be allowed to fluctuate slightly in the region of soft tissue, i.e.,  $-100 < \text{HU} < 100$  (Schneider et al. 2000). For this reason, we applied the regularisation separately to the three intervals.

The regularised cost function is the sum of the data term (equation 11) and three regulariser terms (equation 16). It takes the form

$$E(\vec{R}) = E_{\text{data}} + \gamma_{\text{air}} E_{\text{reg,air}} + \gamma_{\text{soft}} E_{\text{reg,soft}} + \gamma_{\text{bone}} E_{\text{reg,bone}}. \quad (18)$$

We have used the extra subscripts “air”, “soft”, and “bone” as short hand notation to refer to the three distinct HU regions of the conversion curve subject to the regularisation, with the indices  $\alpha_0$  and  $\alpha_M$  in equation 16 chosen appropriately to confine the selected HU intervals. We used the python package *scipy.linalg.lstsq* to minimise the cost function, which is based on singular value decomposition.

The  $\gamma$  factors set the relative weight of the regularisers compared to the data term and their value should be chosen such that the optimised curve contains a plausible level of smoothness.

We adopted the following method to automatically select the weights  $\gamma$ : Starting with an initial guess, we ran the optimisation and then calculated the magnitude of each regulariser term once for the optimised CT-RSP curve and once for a reference curve. We increased the respective weight  $\gamma$ , if the regulariser term of the optimised curve was greater than that of the reference curve, and lowered it in the opposite case. This procedure was repeated iteratively until all regulariser terms of the optimised and reference curve corresponded to within a few percent. In this way, the optimised CT-RSP curve had a comparable amount of variation as the reference.

The reference curve should be chosen (or constructed) such that it includes a level of RSP variation which is thought to be necessary to reflect the tissue variability within one patient.

The relative weights need to be set for each optimisation because the data term to which they refer is linked to a specific proton radiography (and the choice of a region-of-interest - see section 2.12). Note also that the regulariser does not force the absolute RSP values of the optimised and reference curves to resemble each other, but only their degree of smoothness. Therefore, optimised RSP may still vary strongly among patients even in the case of a very smooth reference curve which suppresses RSP variation within one patient. In this work, we used the ground truth CT-RSP curve as reference to fix the regulariser weights, except for figure 4.

### 2.11 Assessment of range accuracy

The patient-specific CT-RSP calibration optimises the parameters  $R_\alpha$  by minimising the difference between predicted and measured average residual range for a field of pencil beams which fully penetrate the patient. From the treatment point of view, it is of interest to estimate how accurately the optimised curve predicts the range of lower energy protons stopping inside the patient.

To quantify the range accuracy which would be obtained with an optimised CT-RSP curve, we performed the following kind of Monte Carlo experiment. We defined a volume inside the patient's head and built an irradiation plan containing 10,000 pencil beams with randomised positions and beam energies. The positions were constrained by the lateral extension of the volume and the energy was chosen so that the pencil beams would penetrate the patient at least 20% of its total water equivalent thickness at that point and at most 80%. We scored the integrated depth dose profiles individually for all pencil beams. The same plan was simulated twice: once using the ground truth CT-RSP curve (see sections 2.2 and 2.3) to convert the patient CT to RSP and once using the optimised one. In this way, we obtained 10,000 pairs of integrated depth dose profiles. From each of them, we estimated the proton range by determining the water equivalent depth in which the distal edge decayed below 80 % of the peak value (Paganetti 2012a). The profiles were linearly interpolated to reach subvoxel resolution. We built histograms out of the 10,000 range differences and determined the mean value as estimate for the range accuracy.

We note that for certain pencil beams directed towards complex structures (e.g. nose and air cavities), the 80% distal fall-off technique may suffer from range mixing and therefore lead to additional variation of the range differences. Manual inspection of a random subsample of profiles led us to conclude that such cases are rare enough not to compromise the results in terms of average range accuracy in the Monte Carlo experiment.

### 2.12 Region of interest in the proton radiography

To perform the CT-RSP curve optimisation, it suffices to acquire a proton radiography of a region of interest of the patient instead of, e.g., the entire head to perform the patient-specific calibration as long as the image contains a sufficient amount of pixels. The region could for example be chosen such to reduce the dose given to healthy tissue. To investigate how much the optimised CT-RSP curve depends on such a choice, we simulated nine proton radiographies capturing differently positioned (but equally sized) rectangular areas (shown in figure 9) and calculated an optimised CT-RSP curve for each of them.

In practice, the region imaged by the proton radiography could be selected in accordance with the treatment volume or, instead, be placed so that the imaged volume does not correspond to the treated one, e.g. to avoid giving dose to an organ at risk distal to the treatment volume during image acquisition. We performed the Monte Carlo experiment (see section 2.11) to investigate whether this would have an impact on the range accuracy. Specifically, we selected a region of interest,  $\text{ROI}_{\text{exp}}$ , and defined the randomised irradiation plans so that the lateral coordinates of all pencil beams lay within it. We used the CT-RSP curve previously optimised based on another (or the same) region of interest proton radiography,  $\text{ROI}_{\text{opti}}$ , as look up table




Figure 3: Left: Optimised CT-RSP curves obtained without (blue) and with (red) regulariser applied. The black curve is the ground truth reference used to simulate the proton radiography. The two dashed vertical lines indicate the three separately regularised HU regions “air”, “soft tissue”, and “bone” (see section 2.10.1). Right: Relative difference between the DRRs calculated using the regularised and non-regularised CT-RSP curves, respectively.

in the Monte Carlo. We then determined the range accuracy as explained in section 2.11 for each pair ( $\text{ROI}_{\text{exp}}$ ,  $\text{ROI}_{\text{opti}}$ ).

## 3 RESULTS

### 3.1 Regularisation

In figure 3 left, we show a patient-specific CT-RSP with and without regularisation used in the optimisation process. The proton radiography was simulated with an 8 mm FWHM beam and we used a rectangular region of interest indicated as black box in the right panel containing 5500 pixels. The projection was performed using the Gaussian beam model with scattering (see section 2.7.2). The patient-specific calibration curve contained 33  $(H_\alpha, R_\alpha)$  pairs to be optimised. It is evident that the variation in the non-regularised curve is unacceptable, while it is kept at a plausible level when employing regularisation. The right panel shows the relative difference between two pDRR images: one calculated using the regularised CT-RSP curve and the other one using the non-regularised version. The two DRRs agree to better than 3% although the two conversion curves themselves differ by much more. This underlines that the non-regularised solution, despite its unrealistic variation, does indeed optimise the residual range similarly well as the regularised one.

The left panel of figure 4 shows optimised calibration curves for which the regulariser weights  $\gamma$  in equation 18 were determined based on two different reference CT-RSP relationships (see section 2.4). The regulariser terms have consequently enforced a degree of smoothness in the optimised curves similar to each reference curve. At the same time, both optimised curves approximately recovered the RSP values of the ground truth curve. This demonstrates that the regulariser does not affect absolute RSP, but only the RSP variability within the curve (see section 2.10.1).

The right panel shows optimised CT-RSP curves for which different HU intervals were used. The five HU thresholds in the yellow curve were  $[-1000, -100, 0, 100, 1376]$ . The  $\gamma$  weights were determined based on the dashed reference curve and consequently the optimised curves are smoother than the ground truth curve. It is interesting to note that the yellow curve does not exhibit any unrealistic variation despite the fact that no regularisation was applied.

The patient-specific calibration curves shown in figure 5 are based on proton radiographies

### 3. RESULTS

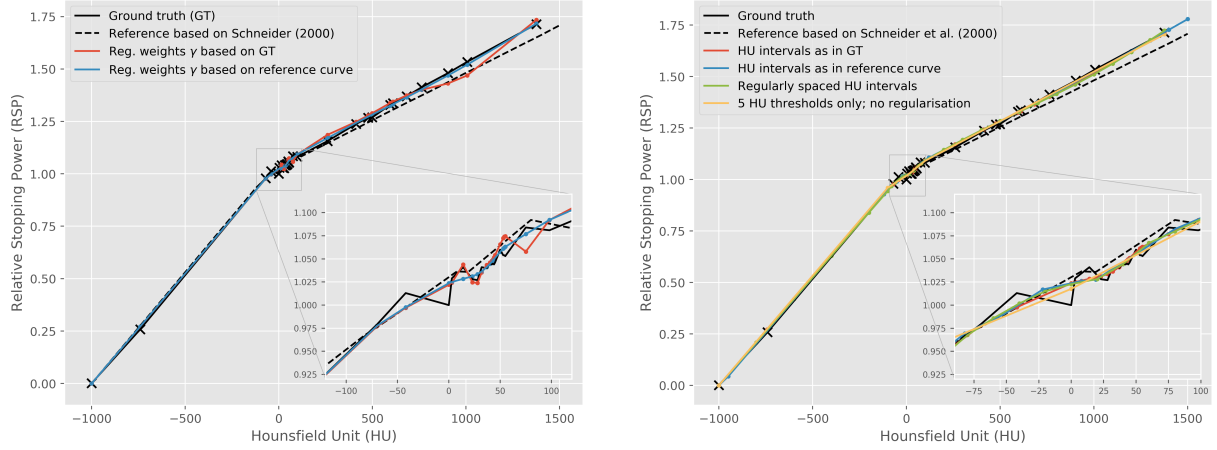


Figure 4: Left: Optimised CT-RSP curves using two different reference curves to determine the regulariser weights  $\gamma$ . Right: Optimised CT-RSP curves using different HU-intervals for the parameterisation. The regulariser weights  $\gamma$  were determined in accordance with the smoothness of the reference curve based on (Schneider et al. 2000) (see section 2.4). For the curve with five HU thresholds only, these were placed at HU=-1000, -100, 0, 100, 1376. No regularisation was applied to this curve.

simulated with different beam sizes, where “pin like” means FWHM=0 mm. We used the Gaussian beam model with scattering to perform the projection of the pDRR. The optimisation was regularised as explained in section 2.10. No regularisation was used for the green curve. The lower panel shows the relative difference in percent between the ground truth CT-RSP and the optimised ones. We have omitted the data point for  $HU = -1000$  because dividing by a very small ground truth RSP value (0.012) makes the relative difference large and meaningless.

Figure 6 illustrates the parameter correlation matrix  $\rho_{\alpha\alpha'}$  (equation 13) for three different beam sizes. We have used the Gaussian beam model with scattering (see section 2.7.2) to calculate the pDRR. The central diagonal contains values of 1 by construction, while the two diagonals off-set to the left and right are close to -1 because  $R_\alpha$  parameters of neighbouring HU intervals of the conversions curve tend to be anti-correlated. Non-zero correlation coefficients further away from the central diagonal occur because RSP information from many voxels is mixed during the projection operation (see section 2.9). Indeed, they are much more prominent in the left panel corresponding to an 8 mm FWHM beam than in the right panel, which refers to a pin-like beam.

In figure 7, we present optimised calibration curves obtained from proton radiographies simulated with different scattering processes activated. The lower panel shows the relative difference in percent between optimised and ground truth CT-RSP curves. “Full physics” includes all processes currently available in *Fred*, i.e., MCS and nuclear elastic and inelastic scattering; for “No nuclear”, only MCS was included; and “no nuclear, no MCS” corresponds to a pure continuous slowing down simulation in which protons travel on straight lines through the voxel geometry. Energy loss and straggling was left at default settings in all cases. The Gaussian beam model with scattering was used to calculate the pDRR except for “no nuclear, no MCS” where the Gaussian beam model without scattering was employed. The FWHM=8 mm curves were regularised (see section 2.10), while the pin-like curves were not.

Figure 8 left shows optimised CT-RSP curves obtained with the three projection models (see section 2.7). In all cases, the proton radiography was simulated with 8 mm FWHM pencil beams and the optimisation was regularised. The Monte Carlo projection (see section 2.7.3) was performed with 1000 protons per pencil beam. The right panel shows the relative difference between the pDRR obtained with the Gaussian projection model with scattering and the Monte Carlo projection.




Figure 5: CT-RSP curves obtained from proton radiographies simulated with different beam sizes. The lower panel shows the relative difference in percent between the optimised and the ground truth CT-RSP curve.




Figure 6: Visualisation of the parameter correlation matrix  $\rho_{\alpha\alpha'}$  (equation 13) for three different beam sizes.




Figure 7: CT-RSP curves obtained from proton radiographies simulated with different physics processes activated. The lower panel shows the relative difference in percent between the optimised and the ground truth CT-RSP curve.

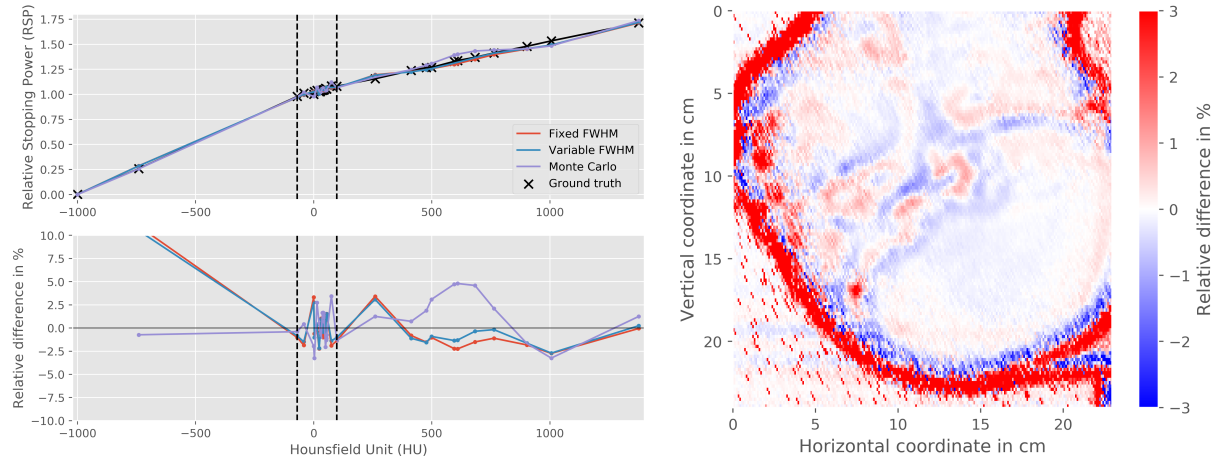


Figure 8: Left: CT-RSP curves obtained using three different projection models (see section 2.7). The lower panel shows the relative difference in percent between the optimised and the ground truth CT-RSP curve. Right: Relative difference between pDRR images obtained with the Gaussian projection model with scattering and the Monte Carlo based projection.

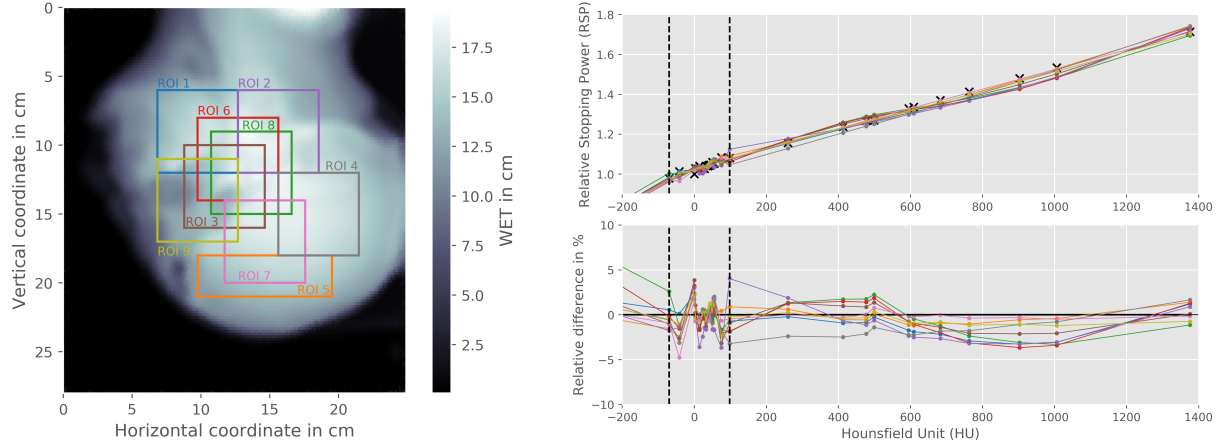


Figure 9: Left: regions of interest in the proton radiography used for the optimisation. Right: Optimised CT-RSP obtained using these regions of interest.

### 3.2 Range accuracy and region of interest

In the figure 9, we show optimised CT-RSP curves (right panel) using the different regions of interest in the proton radiography illustrated in the left panel. Colours in all panels match. The optimisation was performed using the Gaussian projection model with scatter (see section 2.7.2) and employing the regularisation separately in three HU regions (see section 2.10.1). The curves vary by a few percent among each other, but deviate no more than 4% from the ground truth curve. There appears to be some systematic under estimation in the range above 800 HU associated with bony tissue.

Figure 10 presents the results of the Monte Carlo range experiment explained in section 2.11. The upper left panel shows an example distribution of range differences obtained from 10,000 pencil beams. The blue distribution refers to the difference between predicted and true range of the pencil beams when using an optimised CT-RSP. The mean relative range difference quantifies the overall average range accuracy while the RMSE is indicative of the range precision. The yellow distribution was generated by repeating the Monte Carlo experiment twice, both times with the same ground truth curve, and it reflects the intrinsic range variations due to statistics in the depth dose profiles.

The upper right panel of figure 10 shows the mean relative range difference obtained when performing the Monte Carlo experiment in a region of interest,  $\text{ROI}_{\text{exp}}$ , while optimising the CT-RSP curve based on another,  $\text{ROI}_{\text{opti}}$  (see section 2.12). The diagonal elements in the graphic correspond to the case when the imaged area corresponds to the treated volume. On the other hand, for combinations such as  $\text{ROI}_{\text{exp}} = 1$  and  $\text{ROI}_{\text{opti}} = 4$  (see figure 9, left), there is no overlap between the imaged and the treated volume. The lower panels show the RMSE for the control experiment (left) and the optimised vs. ground truth runs (right). The fact that the RMSE tends to be slightly higher when the Monte Carlo experiment was performed in a relatively heterogeneous region (e.g. ROI 9) is in line with the remark in section 2.11 concerning the robustness of the 80% fall-off test. Overall, the RMSE associated with the optimised CT-RSP curves is less than 0.3% higher than the RMSE in the control run.

Figure 11 shows the mean relative range accuracy and the RMSE (as error bars) resulting from the Monte Carlo experiment using the optimised CT-RSP curves based on DRRs calculated with the three beam models (see section 2.7). We performed both, the optimisation and the Monte Carlo experiment, in the same nine regions of interest shown in the left panel of figure 9. The RMSE is only slightly larger than the RMSE of the control run (ground truth vs. ground truth; not shown), i.e. it is mainly due to Monte Carlo statistics and variation of the range retrieved by means of the 80% fall-off criterion (see section 2.11). Overall, the Gaussian beam



### 3. RESULTS

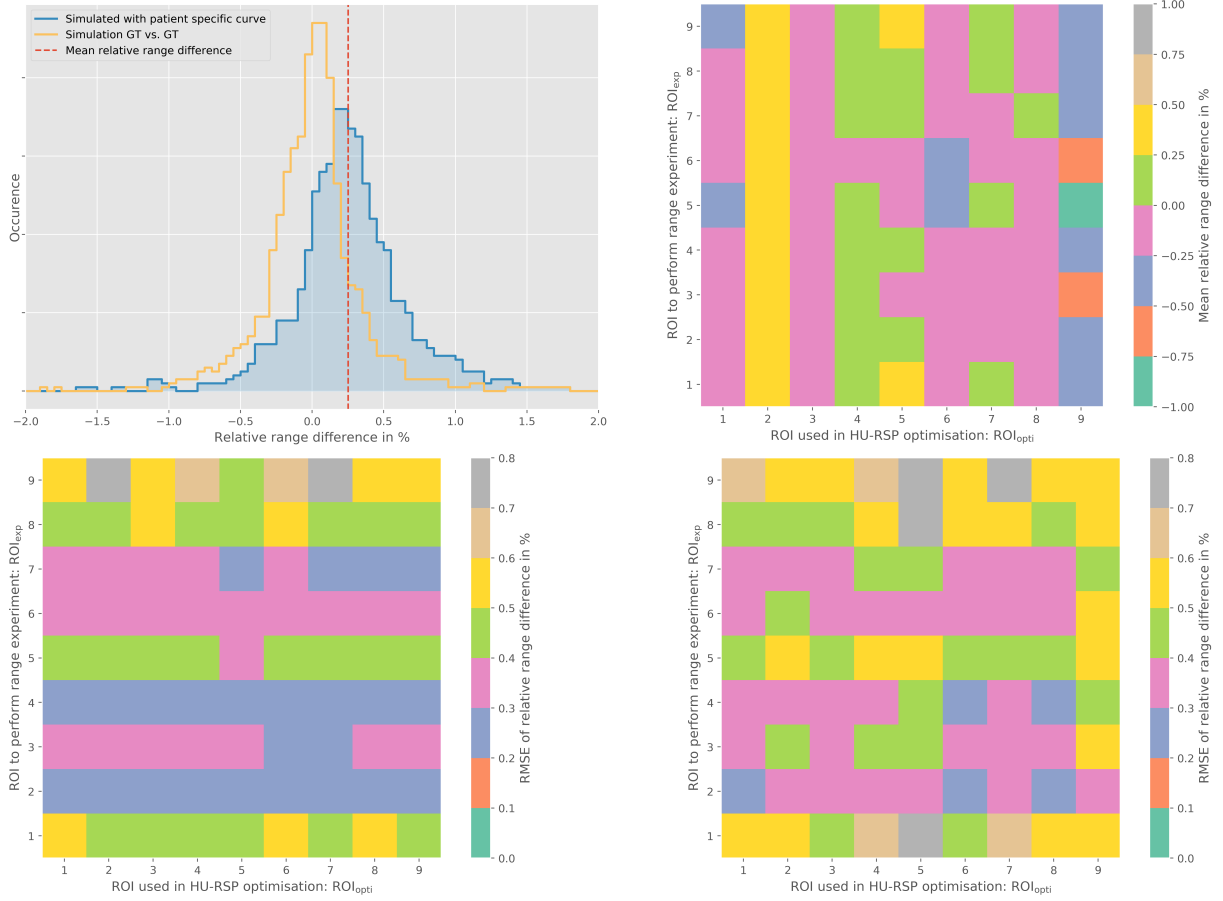


Figure 10: Upper left: example distribution of relative range differences obtained from the Monte Carlo experiment (see section 2.11). Upper right: Mean relative range difference between true and predicted range using different regions of interest for the optimisation and the Monte Carlo experiment (see section 2.11). Lower panels: RMSE in control experiment (left) and between true and predicted range (right).




Figure 11: Mean relative range accuracy obtained in the Monte Carlo experiment (see section 2.11). The CT-RSP curves were optimised for the nine regions of interest shown in figure 9 left and using DRRs calculated with three different beam models (see section 2.7).

model with fixed beam size seems to introduce a slight overestimation of the proton range. The other two beam models, on average, reproduce the ground truth range faithfully to within 0.2-0.5%. Whether or not the minimal positive bias observed with the variable FWHM model is statistically significant appears difficult to establish based on the nine data points.

## 4 DISCUSSION

In the presented work, we introduced methods for the patient-specific calibration of the CT-RSP conversion curve based on proton radiography, with a special regard to proton imaging set-ups which combine pencil beam scanning with a position-insensitive range telescope. Our results demonstrate that the optimisation of a piecewise linear conversion curve needs to be regularised because it otherwise leads to unrealistic variations in the optimised curve. The advantage of a quadratic norm is that the regulariser terms can be easily combined with the data attachment term keeping the optimisation computationally very feasible. With the proposed procedure to select the relative weights of the regulariser terms, the variation in the optimised curve corresponds to that observed in the stoichiometric curve used as reference.

Our results show that the main reason for the need to regularise the cost function is linked to the correlation among the parameters  $R_\alpha$  of the CT-RSP conversion curve, which in turn is strongly influenced by the beam size (figure 6). In a wider phantom than the one considered in this work, scattering becomes more important leading to increased correlation and potentially poorer optimisation results unless more strongly regularised. In case of a pin-like pencil beam, for example, the non-regularised optimisation yielded almost the same result as the regularised one (figure 5). This explains why Collins-Fekete et al. (2017), who investigated the patient-specific CT-RSP calibration based on proton radiographies acquired with single tracking set-ups, did not encounter similar problems: the integration in their case is performed along an estimate of individual proton trajectories with very thin uncertainty envelopes around them and the degree of parameter correlation is expected to be even smaller than with a pin-like pencil beam in our case. On the other hand, the set-up used by Doolan et al. (2015) detects many protons at the same time, as in our case, yet delivered in a passive field rather than by pencil beam scanning. Parameter correlation is therefore expected to be even stronger because the effective beam size is larger. The authors did not report any variation in the optimised CT-RSP curves, such as seen in figure 3, although they did not regularise the optimisation. The reason is probably that they overrode HU values to a single value in each organ making the phantom geometry much more homogeneous. They furthermore excluded heterogeneous regions prone to range mixing. We remark that such a region-wise homogenous RSP map might potentially over-simplify the patient geometry.

We note that the way of parametrising the CT-RSP curve has an impact on the correlation between parameters and thus on the need to regularise the optimisation. Figure 4 shows that no regularisation might be required when using relatively few HU intervals. On the other hand, the CT-RSP curve has then less granularity to resolve intra-patient tissue variability. In practice, a reasonable option would be to choose the HU intervals in accordance with the parametrisation of the clinical lookup table, possibly with some additional intervals for more granularity. The way of interpolating between HU intervals is also expected to affect the optimisation outcome. In our case, neighbouring  $R_\alpha$  values are intrinsically anti-correlated because the linear segments are assumed to form a continuous conversion curve. This is seen from the negative entries below and above the diagonal of the correlation matrix in figure 6. Alternative parametrisations include a piecewise linear interpolation without the constraint of continuity at the intervals boundaries or a piecewise constant conversion curve which assigns the same RSP value to all voxels whose HU values fall in the same interval. In the latter case, the  $R_\alpha$  parameters would be correlated only through the projection operation and consequently less variation would be expected in the optimised conversion curve even without regularisation. We remark that such a piecewise

constant conversion curve assigns the same RSP value to all voxels whose HU values fall in the same interval and is therefore similar to segmenting the CT image into homogeneous material regions as in (Doolan et al. 2015). Investigating the impact of different kinds of interpolation was beyond the scope of this work.

Among the three projection models used in our work, the Monte Carlo based one (section 2.7.3) takes the scattering processes into account most accurately through the physics models implemented in the code. In line with this, differences between the pDRR obtained with the Gaussian projection model with scattering and the Monte Carlo projection are mainly visible in heterogeneous regions and along the outer perimeter of the patient’s head (figure 8, right) where proton scattering has the greatest impact on the projected image. In the left panel of figure 8, the resulting optimised CT-RSP curve differs from the one obtained with the analytical beam model in its shape, but both lie within 3-5% of the ground truth curve. In terms of range accuracy (figure 11), our results suggest that the simple Gaussian beam model with a fixed beam size might lead to an overestimation of up to about 1%. The other two models reproduce the ground truth range expectation to 0.2-0.5% accuracy on average. Among the two, the Monte Carlo beam model seems to be slightly preferable over the variable FWHM Gaussian model, although we cannot deduce whether this observation is statistically significant. At the moment, the latter is much faster ( $< 1$  min) to calculate thanks to our GPU implementation compared to the former (order of hours) which currently runs on single CPU only, although these technical issues could be improved in the future.

Remapping the HU values  $h$  in the CT voxels to coefficient vectors  $\vec{w}$  (section 2.6) allowed us to separate the projection operation from the actual optimisation. This is particularly useful when using a cost function that requires iterative minimisation. We emphasise that the remapping can be used with proton radiographies acquired by any kind of set-up. For example, when using a single tracking proton scanner (Collins-Fekete et al. 2017), the projection operation to construct the pDRR (section 2.7) would be performed by integrating the  $w_{ijk,\alpha}$  matrices along the estimated most likely path of every detected proton.

In our Monte Carlo experiment, the range accuracy achievable with patient-specific conversion curves optimised using the methods presented here was well below 1%. Whether or not the imaged region-of-interest is chosen in accordance with the treated volume did not systematically impact the range accuracy. This statement only refers to the optimisation process itself and does not take into account possible local variability in the tissue composition.

We chose to perform our study entirely in Monte Carlo because it allowed us to define a ground truth CT-RSP conversion curve. This approach is idealised in one regard: In reality, no one-to-one correspondence exists between the HU and RSP of the patient tissue. Our ground truth scenario is therefore an approximation because we generate the RSP map through a one-to-one conversion process. Studying the impact of the tissue composition on range accuracy could be achieved by modelling the expected tissue variation (Yang et al. 2010, Möhler et al. 2016, Wohlfahrt et al. 2017) and perturbing the converted RSP value in each voxel. We emphasise that this was not the scope of this work. It should also be noted that the patient-specific CT-RSP calibration based on proton radiography cannot solve the issue of tissue variability within one patient because it intrinsically relies on a conversion curve. Only a full volumetric proton CT image would provide a voxel-wise RSP value and thereby circumvent the conversion process.

## 5 CONCLUSION

We developed a series of methods to perform a patient-specific calibration of the HU to RSP conversion curve based on proton radiography. We introduced a regulariser term in the cost function and showed that this is necessary to avoid unrealistic variation in the optimised curve. The main reason is the correlation between the parameters of the piecewise linearly interpolated CT-RSP curve and the effect is more pronounced with larger beam size. We performed Monte

Carlo experiments and obtained a range accuracy of better than 0.5%. Our results further indicate that from the point of view of the optimisation procedure, no systematic loss of range accuracy is probably to be expected if the imaged region of the patient does not correspond to the treatment volume. This work underlines that it is possible to perform a patient-specific calibration of the CT-RSP conversion curve on the basis of a single proton radiographic projection acquired with a detector typically available in a treatment facility.

## Acknowledgements

The work of Nils Krah was supported by funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 753370. The work of Ilaria Rinaldi was supported partially by a research fellowship “Visiting Professor Sapienza” from the University of Rome. This work was performed within the framework the project DIC20161236452 of the “Fondation pour la Recherche Médicale” (FRM), of the SIRIC LYriCAN Grant INCa.INSERM.DGOS.12563 and of the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007) operated by the ANR. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan X Pascal GPU used for this research.

## References

- Agostinelli, S., Allison, J., Amako, K., Apostolakis, J., Araujo, H., Arce, P., Asai, M., Axen, D., Banerjee, S., Barrand, G., Behner, F., Bellagamba, L., Boudreau, J., Broglia, L., Brunengo, A., Burkhardt, H., Chauvie, S., Chuma, J., Chytrcek, R., Cooperman, G., Cosmo, G., Degt'yarenko, P., Dell’Acqua, A., Depaola, G., Dietrich, D., Enami, R., Feliciello, A., Ferguson, C., Fesefeldt, H., Folger, G., Foppiano, F., Forti, A., Garelli, S., Giani, S., Giannitrapani, R., Gibin, D., Gomez Cadenas, J. J., Gonzalez, I., Gracia Abril, G., Greeniaus, G., Greiner, W., Grichine, V., Grossheim, A., Guatelli, S., Gumplinger, P., Hamatsu, R., Hashimoto, K., Hasui, H., Heikkinen, A., Howard, A., Ivanchenko, V., Johnson, A., Jones, F. W., Kallenbach, J., Kanaya, N., Kawabata, M., Kawabata, Y., Kawaguti, M., Kelner, S., Kent, P., Kimura, A., Kodama, T., Kokoulin, R., Kossov, M., Kurashige, H., Lamanna, E., Lampen, T., Lara, V., Lefebvre, V., Lei, F., Liendl, M., Lockman, W., Longo, F., Magni, S., Maire, M., Medernach, E., Minamimoto, K., Mora de Freitas, P., Morita, Y., Murakami, K., Nagamatsu, M., Nartallo, R., Nieminen, P., Nishimura, T., Ohtsubo, K., Okamura, M., O’Neale, S., Oohata, Y., Paech, K., Perl, J., Pfeiffer, A., Pia, M. G., Ranjard, F., Rybin, A., Sadilov, S., di Salvo, E., Santin, G., Sasaki, T., Savvas, N., Sawada, Y., Scherer, S., Sei, S., Sirotenko, V., Smith, D., Starkov, N., Stoecker, H., Sulkimo, J., Takahata, M., Tanaka, S., Tcherniaev, E., Safai Tehrani, E., Tropeano, M., Truscott, P., Uno, H., Urban, L., Urban, P., Verderi, M., Walkden, A., Wander, W., Weber, H., Wellisch, J. P., Wenaus, T., Williams, D. C., Wright, D., Yamada, T., Yoshida, H. & Zschesche, D. (2003), ‘GEANT4 - A simulation toolkit’, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**(3), 250–303.
- Allison, J., Amako, K., Apostolakis, J., Arce, P., Asai, M., Aso, T., Bagli, E., Bagulya, A., Banerjee, S., Barrand, G., Beck, B., Bogdanov, A., Brandt, D., Brown, J., Burkhardt, H., Canal, P., Cano-Ott, D., Chauvie, S., Cho, K., Cirrone, G., Cooperman, G., Cortés-Giraldo, M., Cosmo, G., Cuttone, G., Depaola, G., Desorgher, L., Dong, X., Dotti, A., Elvira, V., Folger, G., Francis, Z., Galoyan, A., Garnier, L., Gayer, M., Genser, K., Grichine, V., Guatelli, S., Guèye, P., Gumplinger, P., Howard, A., Hivnáčová, I., Hwang, S., Incerti, S., Ivanchenko, A., Ivanchenko, V., Jones, F., Jun, S., Kaitaniemi, P., Karakatsanis, N., Karamitros, M., Kelsey, M., Kimura, A., Koi, T., Kurashige, H., Lechner, A., Lee, S., Longo,

- F., Maire, M., Mancusi, D., Mantero, A., Mendoza, E., Morgan, B., Murakami, K., Nikitina, T., Pandola, L., Paprocki, P., Perl, J., Petrović, I., Pia, M., Pokorski, W., Quesada, J., Raine, M., Reis, M., Ribon, A., Ristić Fira, A., Romano, F., Russo, G., Santin, G., Sasaki, T., Sawkey, D., Shin, J., Strakovsky, I., Taborda, A., Tanaka, S., Tomé, B., Toshito, T., Tran, H., Truscott, P., Urban, L., Uzhinsky, V., Verbeke, J., Verderi, M., Wendt, B., Wenzel, H., Wright, D., Wright, D., Yamashita, T., Yarba, J. & Yoshida, H. (2016), ‘Recent developments in G eant 4’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **835**, 186–225.  
**URL:** <http://www.sciencedirect.com/science/article/pii/S0168900216306957>  
<https://linkinghub.elsevier.com/retrieve/pii/S0168900216306957>
- Arbor, N., Dauvergne, D., Dedes, G., Létang, J. M., Parodi, K., Quiñones, C. T., Testa, E. & Rit, S. (2015), ‘Monte Carlo comparison of x-ray and proton CT for range calculations of proton therapy beams’, *Physics in Medicine and Biology* **60**(19), 7585.  
**URL:** <http://stacks.iop.org/0031-9155/60/i=19/a=7585>
- Böhlen, T., Cerutti, F., Chin, M., Fassò, A., Ferrari, A., Ortega, P., Mairani, A., Sala, P., Smirnov, G. & Vlachoudis, V. (2014), ‘The FLUKA Code: Developments and Challenges for High Energy and Medical Applications’, *Nuclear Data Sheets* **120**, 211–214.  
**URL:** <http://adsabs.harvard.edu/abs/2014NDS...120..211B>
- Bortfeld, T. & Schlegel, W. (1996), ‘An analytical approximation of depth - dose distributions for therapeutic proton beams’, *Physics in Medicine and Biology* **41**(8), 1331–1339.  
**URL:** <http://stacks.iop.org/0031-9155/41/i=8/a=006>
- Collins-Fekete, C.-A., Brousmiche, S., Hansen, D. C., Beaulieu, L. & Seco, J. (2017), ‘Pre-treatment patient-specific stopping power by combining list-mode proton radiography and x-ray CT’, *Physics in Medicine & Biology* **62**(17), 6836–6852.  
**URL:** <http://stacks.iop.org/0031-9155/62/i=17/a=6836>
- Doolan, P. J., Testa, M., Sharp, G., Bentefour, E. H., Royle, G. & Lu, H.-M. (2015), ‘Patient-specific stopping power calibration for proton therapy planning based on single-detector proton radiography’, *Physics in Medicine and Biology* **60**(5), 1901–1917.  
**URL:** <http://stacks.iop.org/0031-9155/60/i=5/a=1901>
- Durante, M. & Paganetti, H. (2016), ‘Nuclear physics in particle therapy: a review’, *Reports on Progress in Physics* **79**(9), 096702.  
**URL:** <http://stacks.iop.org/0034-4885/79/i=9/a=096702>
- Farace, P., Righetto, R. & Meijers, A. (2016), ‘Pencil beam proton radiography using a multilayer ionization chamber’, *Physics in medicine and biology* **61**(11), 4078–87.  
**URL:** <http://www.ncbi.nlm.nih.gov/pubmed/27164479>
- Ferrari, A., Sala, P. R., Fassò, A. & Ranft, J. (2005), *FLUKA: A multi-particle transport code (program version 2005)*, Technical Report CERN-2005-10, INFN/TC 05/11, SLAC-R-773, Geneva.
- Garbacz, M., Battistoni, G., Durante, M., Gajewski, J., Krah, N., Patera, V., Rinaldi, I., Schiavi, A., Scifoni, E., Skrzypek, A., Tommasino, F. & Rucinski, A. (2019), Proton Therapy Treatment Plan Verification in CCB Krakow Using Fred Monte Carlo TPS Tool, in L. Lhotska, L. Sukupova, I. Lacković & G. S. Ibbott, eds, ‘World Congress on Medical Physics and Biomedical Engineering 2018’, Springer Singapore, Singapore, pp. 783–787.

- Gottschalk, B., Koehler, A., Schneider, R., Sisterson, J. & Wagner, M. (1993), ‘Multiple Coulomb scattering of 160 MeV protons’, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **74**(4), 467–490.  
**URL:** <http://www.sciencedirect.com/science/article/pii/0168583X9395944Z>
- Jäkel, O., Jacob, C., Schardt, D., Karger, C. P. & Hartmann, G. H. (2001), ‘Relation between carbon ion ranges and x-ray CT numbers’, *Medical Physics* **28**(4), 701–703.  
**URL:** <http://doi.wiley.com/10.1118/1.1357455>
- Johnson, R. P. (2018), ‘Review of medical radiography and tomography with proton beams’, *Reports on Progress in Physics* **81**(1), 016701.  
**URL:** <http://stacks.iop.org/0034-4885/81/i=1/a=016701>
- Koehler, A. M. (1968), ‘Proton Radiography’, *Science* **160**(3825), 303 LP – 304.  
**URL:** <http://science.sciencemag.org/content/160/3825/303.abstract>
- Krah, N., Khellaf, F., Létang, J. M., Rit, S. & Rinaldi, I. (2018), ‘A comprehensive theoretical comparison of proton imaging set-ups in terms of spatial resolution’, *Physics in Medicine & Biology* **63**(13), 135013.  
**URL:** <http://stacks.iop.org/0031-9155/63/i=13/a=135013>
- Lynch, G. R. & Dahl, O. I. (1991), ‘Approximations to multiple Coulomb scattering’, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **58**(1), 6–10.  
**URL:** <http://linkinghub.elsevier.com/retrieve/pii/0168583X9195671Y>
- Möhler, C., Wohlfahrt, P., Richter, C. & Greulich, S. (2016), ‘Range prediction for tissue mixtures based on dual-energy CT’, *Physics in Medicine and Biology* **61**(11), N268–N275.  
**URL:** <http://stacks.iop.org/0031-9155/61/i=11/a=N268>
- Paganetti, H. (2012a), *Proton Therapy Physics*, Taylor & Francis Group, Boca Raton.
- Paganetti, H. (2012b), ‘Range uncertainties in proton therapy and the role of Monte Carlo simulations’, *Physics in Medicine and Biology* **57**(11), R99–R117.  
**URL:** <http://stacks.iop.org/0031-9155/57/i=11/a=R99>
- Parodi, K. (2014), ‘Heavy ion radiography and tomography’, *Physica Medica* **30**(5), 539–543.  
**URL:** <http://dx.doi.org/10.1016/j.ejmp.2014.02.004>
- Poludniowski, G., Allinson, N. M. & Evans, P. M. (2015), ‘Proton radiography and tomography with application to proton therapy’, *The British Journal of Radiology* **88**(1053), 20150134.  
**URL:** <http://www.birpublications.org/doi/10.1259/bjr.20150134>
- Rinaldi, I., Brons, S., Gordon, J., Panse, R., Voss, B., Jäkel, O. & Parodi, K. (2013), ‘Experimental characterization of a prototype detector system for carbon ion radiography and tomography’, *Physics in Medicine and Biology* **58**(3), 413–427.  
**URL:** <http://stacks.iop.org/0031-9155/58/i=3/a=413>
- Rinaldi, I., Brons, S., Jäkel, O., Voss, B. & Parodi, K. (2014), ‘Experimental investigations on carbon ion scanning radiography using a range telescope.’, *Physics in medicine and biology* **59**(12), 3041–57.  
**URL:** <http://stacks.iop.org/0031-9155/59/i=12/a=3041>
- Schiavi, A., Senzacqua, M., Pioli, S., Mairani, A., Magro, G., Molinelli, S., Ciocca, M., Battistoni, G. & Patera, V. (2017), ‘Fred: A GPU-accelerated fast-Monte Carlo code for rapid treatment plan recalculation in ion beam therapy’, *Physics in Medicine and Biology* **62**(18), 7482–7504.  
**URL:** <https://doi.org/10.1088/1361-6560/aa8134>

- Schneider, U. & Pedroni, E. (1995), ‘Proton radiography as a tool for quality control in proton therapy’, *Medical Physics* **22**(4), 353–363.  
**URL:** <http://link.aip.org/link/?MPH/22/353/1>
- Schneider, U., Pedroni, E. & Lomax, A. (1996), ‘The calibration of CT Hounsfield units for radiotherapy treatment planning’, *Physics in Medicine and Biology* **41**(1), 111.  
**URL:** <http://stacks.iop.org/0031-9155/41/i=1/a=009>
- Schneider, U., Pemler, P., Besserer, J., Pedroni, E., Lomax, A. & Kaser-Hotz, B. (2005), ‘Patient specific optimization of the relation between CT-Hounsfield units and proton stopping power with proton radiography’, *Medical Physics* **32**(1), 195–199.  
**URL:** <http://link.aip.org/link/?MPH/32/195/1>
- Schneider, W., Bortfeld, T. & Schlegel, W. (2000), ‘Correlation between CT numbers and tissue parameters needed for Monte Carlo simulations of clinical dose distributions.’, *Physics in medicine and biology* **45**(2), 459–78.  
**URL:** <http://www.ncbi.nlm.nih.gov/pubmed/10701515>
- Schulte, R., Bashkirov, V., Tianfang Li, Zhengrong Liang, Mueller, K., Heimann, J., Johnson, L., Keeney, B., Sadrozinski, H.-W., Seiden, A., Williams, D., Lan Zhang, Zhang Li, Peggs, S., Satogata, T. & Woody, C. (2004), ‘Conceptual design of a proton computed tomography system for applications in proton radiation therapy’, *IEEE Transactions on Nuclear Science* **51**(3), 866–872.  
**URL:** <http://ieeexplore.ieee.org/document/1311983/>
- Schulte, R. W., Penfold, S. N., Tafas, J. T. & Schubert, K. E. (2008), ‘A maximum likelihood proton path formalism for application in proton computed tomography’, *Medical Physics* **35**(11), 4849.  
**URL:** <http://scitation.aip.org/content/aapm/journal/medphys/35/11/10.1118/1.2986139>
- Steward, V. W. & Koehler, A. M. (1973), ‘Proton Beam Radiography in Tumor Detection’, *Science* **179**(4076), 913–914.  
**URL:** <http://www.sciencemag.org/cgi/doi/10.1126/science.179.4076.913>
- Taasti, V. T., Bäumer, C., Dahlgren, C. V., Deisher, A. J., Ellerbrock, M., Free, J., Gora, J., Kozera, A., Lomax, A. J., De Marzi, L., Molinelli, S., Kevin Teo, B.-K., Wohlfahrt, P., Petersen, J. B., Muren, L. P., Hansen, D. C. & Richter, C. (2018), ‘Inter-centre variability of CT-based stopping-power prediction in particle therapy: Survey-based evaluation’, *Physics and Imaging in Radiation Oncology* **6**, 25–30.  
**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S2405631618300216>
- Takada, Y., Kondo, K., Marume, T., Nagayoshi, K., Okada, I. & Takikawa, K. (1988), ‘Proton computed tomography with a 250 MeV pulsed beam’, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **273**(1), 410–422.  
**URL:** <https://www.sciencedirect.com/science/article/pii/0168900288908443>
- West, D. & Sherwood, A. (1972), ‘Radiography with 160 MeV protons’, *Nature* **239**, 157–159.  
**URL:** <http://www.nature.com/nature/journal/v239/n5368/abs/239157b0.html>
- Williams, D. C. (2004), ‘The most likely path of an energetic charged particle through a uniform medium’, *Physics in Medicine and Biology* **49**(13), 2899–2911.  
**URL:** <http://stacks.iop.org/0031-9155/49/i=13/a=010>

Wohlfahrt, P., Möhler, C., Richter, C. & Greilich, S. (2018), ‘Evaluation of Stopping-Power Prediction by Dual- and Single-Energy Computed Tomography in an Anthropomorphic Ground-Truth Phantom’, *International Journal of Radiation Oncology\*Biology\*Physics* **100**(1), 244–253.

**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0360301617338920>

Wohlfahrt, P., Möhler, C., Stützer, K., Greilich, S. & Richter, C. (2017), ‘Dual-energy CT based proton range prediction in head and pelvic tumor patients’, *Radiotherapy and Oncology* **125**(3), 526–533.

**URL:** <https://linkinghub.elsevier.com/retrieve/pii/S0167814017326282>

Yang, M., Virshup, G., Clayton, J., Zhu, X. R., Mohan, R. & Dong, L. (2010), ‘Theoretical variance analysis of single- and dual-energy computed tomography methods for calculating proton stopping power ratios of biological tissues’, *Physics in Medicine and Biology* **55**(5), 1343–1362.

**URL:** <http://stacks.iop.org/0031-9155/55/i=5/a=006?key=crossref.dfc557bb88c6d3a5127ef1dd48a5f62f>