



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Quantitative electron density CT imaging for radiotherapy planning

Citation for published version:

Mason, JH, Perelli, A, Nailon, WH & Davies, ME 2017, Quantitative electron density CT imaging for radiotherapy planning. in *Medical Image Understanding and Analysis - 21st Annual Conference, MIUA 2017, Proceedings*. vol. 723, Communications in Computer and Information Science, vol. 723, Springer-Verlag, pp. 297-308, 21st Annual Conference on Medical Image Understanding and Analysis, MIUA 2017, Edinburgh, United Kingdom, 11/07/17. https://doi.org/10.1007/978-3-319-60964-5_26

Digital Object Identifier (DOI):

[10.1007/978-3-319-60964-5_26](https://doi.org/10.1007/978-3-319-60964-5_26)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Medical Image Understanding and Analysis - 21st Annual Conference, MIUA 2017, Proceedings

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Quantitative Electron Density CT Imaging for Radiotherapy Planning

Jonathan Mason¹, Alessandro Perelli¹, William Nailon², and Mike Davies¹

¹ Institute for Digital Communications, University of Edinburgh, EH9 3JL, UK,
`{j.mason},{a.perelli},{mike.davies}@ed.ac.uk`,

² Oncology Physics Department, Edinburgh Cancer Centre,
Western General Hospital, Edinburgh, EH4 2XU, UK,
`bill.nailon@luht.scot.nhs.uk`

Abstract Computed tomography (CT) is the imaging modality used to calculate the deposit of dose in radiotherapy planning, where the physical interactions are modelled based upon the electron density, which can be calculated from CT images. Traditionally this is a three step process: linearising the raw x-ray measurements and correcting for beam-hardening and scatter; inverting the system with analytic or iterative reconstruction algorithms into linear attenuation coefficient; then applying a non-linear calibration into electron density. In this work, we propose a new method for statistically inferring a quantitative image of electron density directly from the raw CT measurements, with no pre- or post-processing necessary, and able to cope with both beam-hardening from a single polyenergetic source and additive scatter. We evaluate this concept with cone-beam CT (CBCT) imaging for bladder cancer, where we demonstrate significantly higher electron density accuracy than other quantitative approaches. We also show through simulated photon and proton beam calculation, that our method may facilitate superior dose estimation, especially with regions containing bony structures.

Keywords: computed tomography, quantitative imaging, statistical reconstruction, proton therapy

1 Introduction

Computed tomography (CT) imaging facilitates radiation therapy planning by providing a quantitative map of electro-magnetic attenuation from a low x-ray energy source. With this, one can approximate how the higher energy treatment beams will interact with the patient and distribute dose throughout tumours and sensitive organs. For photon and proton therapy, dose calculation methods usually rely upon electron density [1], since it allows the interaction processes at the high energies to be modelled accurately. However, the attenuation with CT imaging energies is not solely determined by quantity of electrons, but their atomic environment, which induces whether the photon will interact through photoelectric effect, scattering or other [2]. When a specimen contains significantly different materials, such as soft-tissue and bone, the attenuation will

be both spatially and spectrally non-linear, so calculating a consistent energy independent electron density is not straight forward.

With traditional approaches to CT imaging, the mapping from raw x-ray measurement to electron density is performed in three distinct steps. Firstly, the measurements are converted to linearised projections, which requires correcting for scatter, taking the logarithm, and calibrating from the polyenergetic to approximate monoenergetic source—known as beam hardening correction. Common assumptions for this conversion are either that the specimen has the spectral properties of water, or a mix of bone and water [3], the latter requiring a preliminary reconstruction and segmentation of the bony structures.

After the preprocessing from measurement to projections comes the ‘reconstruction algorithm’ [4], that attempts to invert the system into a spatial distribution of linear attenuation coefficient. These are broadly classed into ‘analytic methods’, where a closed form approximate inverse is applied to the projections, or ‘iterative methods’, where noise modelling and regularisation can be incorporated to infer a more accurate reconstruction through optimisation. Due to its preprocessing, the resulting images are approximations of a monoenergetic source attenuation.

Finally, the reconstruction is non-linearly converted into electron density through calibration [1]. In general, a three component linear fit is accurate for a biological specimen. At each of the three stages in this reconstruction work-flow, approximations are made that inevitably introduce artefacts that will propagate through the mapping, which will become increasingly pronounced as the dose is lowered. If this processes could be replaced by a direct inference with an accurate measurement model, then it is likely that these errors could be mitigated.

To date, there are several techniques that go some of the way to this direct inference. In [5], the authors propose a statistical algorithm for inferring the mass density from polyenergetic measurements and additive noise. Here, given a prior segmentation of the specimen into distinct material classes, such as water and bone, the attenuation within each region is modelled as a product of mass density and energy dependent mass attenuation coefficient of the corresponding class. They then presented an alternative model in [6], where the material classification is itself a function of mass density, and so no prior segmentation is necessary. A pre-dating method [7] attempts to model the photoelectric and scattering cross-sections, determining the polyenergetic attenuation, as a function of monoenergetic attenuation.

In both of the methods [6][7], there is an explicit assumption that for biological tissues, the polyenergetic attenuation can be modelled as an energy dependent piecewise linear function of either mass density or monoenergetic attenuation coefficient. Since the calibration from either of these into electron density [1] is also a piecewise linear calibration, we suggest it is reasonable to assume that these may be combined into a single mapping, which is what we set out to do in Section 3.1.

From the new relation in Section 3.1, we then couple this with an accurate CT model incorporating spectral properties of the source, and additive scatter

or noise. In Section 3.2, we then show how this new model can be used for quantitative reconstruction, and how through properties of the model design, efficient algorithms can be realised that may include many smooth and non-smooth regularisation functions, such as generalized Gaussian Markov random field (GGMRF) [8], total variation (TV) [9] or wavelet sparsity [10].

As a proof of concept, we evaluate the accuracy of our method in the challenging geometry of cone-beam CT (CBCT), with its high scatter [11] and poor sampling [12], for radiation therapy planning of the bladder region. We demonstrate the superior quantitative accuracy of our approach in both image and dose, when compared to standard three-step reconstructions, and the model in [5].

2 Background

2.1 X-ray Attenuation

The mechanism that allows various regions in a heterogeneous specimen to be differentiated is their degree of x-ray attenuation. For biological tissues irradiated with a diagnostic x-ray source, the significant phenomena contributing to the attenuation of incident radiation are photoelectric and scattering effects—consisting of Compton and coherent scatter. The combined attenuation strength of a given element is quantified with the linear attenuation coefficient, defined as

$$\mu = \frac{\rho(\sigma_{\text{p.e.}} + \sigma_{\text{incoh.}} + \sigma_{\text{coh.}})}{uA}, \quad (1)$$

where ρ is the mass density, $\sigma_{\text{p.e.}}$, $\sigma_{\text{incoh.}}$ and $\sigma_{\text{coh.}}$ represent the interaction cross sections of photoelectric, incoherent (Compton) and coherent (Rayleigh and Thompson) effects, A is the atomic mass of the element, and u is a constant—the unified atomic mass unit. For a material, its attenuation coefficient is found by a sum of its constituent elemental coefficients, weighted by their mass fraction.

Since the strength of each attenuation effect varies with energy, we find that μ is dependent on a material’s mass density, elemental composition and the incident x-ray source spectrum.

2.2 Probabilistic Measurement Model

In CT, one is able to observe a specimen’s attenuation through the radiation intensity after transmission. The magnitude of this is often approximated using the Beer–Lambert law, given for a monoenergetic beam as

$$I_o = I_i \exp\left(-\int_{\ell} \mu(\ell) d\ell\right),$$

where I_i is the incident intensity, ℓ is the line-of-sight path of the beam through specimen, I_o is the output intensity one is able to measure. Since in practice, μ

is energy dependent and we have a polyenergetic source, the output intensity in this case is

$$I_o = \int_{\xi} I_i(\xi) \exp \left(- \int_{\ell} \mu(\ell, \xi) d\ell \right) d\xi,$$

where ξ is energy.

For a finite number of photons, the measured intensity will be probabilistic with an approximate Poisson distribution [13]. If we also move the attenuation, measurements and energy spectrum into a discretised setting, we can write the measurement process as

$$y_i \sim \text{Poisson} \left\{ \sum_{j=1}^{N_{\xi}} b_i(\xi_j) \exp(-[\Phi \mu(\xi_j)]_i) + s_i \right\} \text{ for } i = 1, \dots, N_{\text{ray}}, \quad (2)$$

where N_{ray} is the number of CT measurements, N_{ξ} is the number of energy bins, $\mathbf{b}(\xi) \in \mathbb{R}^{N_{\text{ray}}}$ is a vector of incident intensities, $\mu(\xi) \in \mathbb{R}^{N_{\text{vox}}}$ is the vector of attenuation coefficients with N_{vox} the number of voxels, $\Phi \in \mathbb{R}^{N_{\text{ray}} \times N_{\text{vox}}}$ is the system matrix describing the paths from source through specimen onto each detector, and $\mathbf{s} \in \mathbb{R}^{N_{\text{ray}}}$ is expectation of scatter or other background noise reaching the detector.

2.3 Standard Reconstruction Approaches

Inferring μ from measurements with the model in (2) is not straightforward due to its the non-linearity and energy dependence. Instead, a standard approach is to assume a monoenergetic model and attempt to linearise the system. In a simple form, this linearised model can be expressed as

$$\mathbf{p} = \Phi \mu + \mathbf{n}, \quad (3)$$

where $\mathbf{n} \in \mathbb{R}^{N_{\text{ray}}}$ is noise, $\mathbf{p} \in \mathbb{R}^{N_{\text{ray}}}$ is the linearised projection, calculated by

$$p_i = \log \left(\frac{\bar{b}_i}{y_i - s_i} \right) \text{ for } i = 1, \dots, N_{\text{ray}},$$

where $\bar{b}_i = \sum_{j=1}^{N_{\xi}} b_i(\xi_j)$ is the total incident flux. In practice, \mathbf{p} is calibrated to approximate the monoenergetic equivalent before attempting to invert the system in (3).

One class of reconstruction methods—known as ‘analytic’ [4]—attempt to find a closed form expression for the inverse of (3). These normally are classed as ‘filtered back-projection’ (FBP) or ‘back-projection filtration’ (BPF). For the cone-beam geometry, the popular Feldkamp-Davis-Kress (FDK) is FBP. These methods are usually very fast, but perform poorly in high noise or limited measurement scenarios.

Another class of reconstruction algorithm is known as ‘iterative’ [4], which usually try to find an optimal μ by some metric. A popular version of iterative

reconstruction for (3) is penalised weighted-least-squares (PWLS), where $\boldsymbol{\mu}$ is found as [5]

$$\hat{\boldsymbol{\mu}} = \underset{\boldsymbol{\mu}}{\operatorname{argmin}} (\boldsymbol{\Phi}\boldsymbol{\mu} - \boldsymbol{p})^T \boldsymbol{W} (\boldsymbol{\Phi}\boldsymbol{\mu} - \boldsymbol{p}) + \lambda R(\boldsymbol{\mu}), \quad (4)$$

where $\boldsymbol{W} \in \mathbb{R}^{N_{\text{ray}} \times N_{\text{ray}}}$ is a diagonal weighting matrix with entries $w_{ii} = (y_i - s_i)^2 / y_i$, $R(\boldsymbol{\mu})$ is a regularisation function to promote desirable structure in $\boldsymbol{\mu}$, and λ is usually a scalar constant trade-off between data fit and regularisation. The weights are used here to approximate the Poisson noise model in (2).

In order to calculate dose from reconstructions in $\boldsymbol{\mu}$, whether the linearised system is used within FBP and PWLS, or a more accurate maximum-likelihood polyenergetic model is used [5][6][7], conversion to electron density is necessary. In [1], this is done through a piecewise linear fit derived from the chemical properties of many biological tissues.

3 Method

3.1 Electron Density to Attenuation

The essence of our method is to use the model in (2) for reconstruction directly into electron density, ρ_e . This involves a single non-linear energy dependent fitting from ρ_e to attenuation, which depends on the existence of a simple relation. To investigate the possibility of such a mapping, we calculated the relative electron density—normalised to that of water—for 52 material classes from the International Commission on Radiological Protection (ICRP) report 89 [14]. This is plotted against the relative attenuation—normalised to maximum attenuation—in Figure 1.

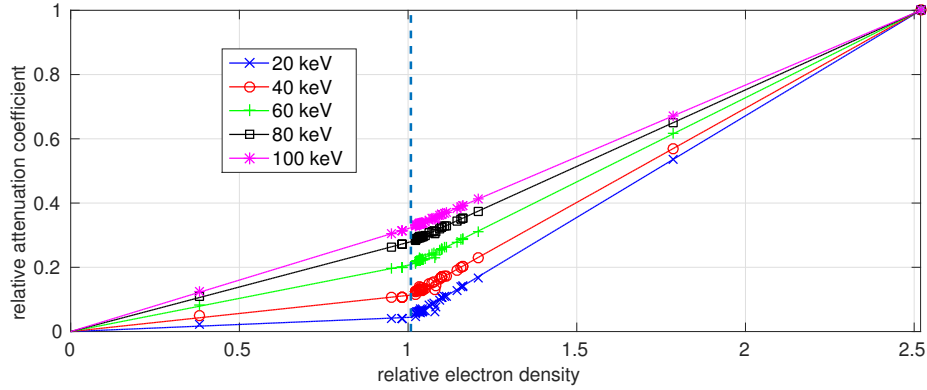


Figure 1: Plot of relative attenuation coefficient against electron density for ICRP tissues over a range in energy. Two linear fits are made to this data, with the transition between them shown by the dashed line.

By observing the data in Figure 1, it becomes apparent that there is indeed a simple well fitting model that can map from electron density to attenuation. We have superimposed two linear fits around a constant point, and imposed constraints that the first fit should pass through the origin and both fits cross at a consistent ‘knee’ position, which was calculated to maximise the accuracy of overall fit—1.0083 in this case.

To make the realisation of a practical algorithm easier, we make use of a single continuously differentiable function to approximate the piecewise linear nature from the two fits shown in Figure 1. For this we use the generalised logistic function

$$f(x) = \frac{1}{1 + \exp(-k(x - x_0))},$$

where k is the steepness of its transition and x_0 is the central point, which we set to our knee. The continuous mapping from electron density to attenuation coefficient can then be written as

$$\hat{\mu}(\rho_e, \xi) = [1 - f(\rho_e)]\alpha(\xi)\rho_e + f(\rho_e)[\beta(\xi)\rho_e + \gamma(\xi)], \quad (5)$$

where α is the gradient of first fit, and β, γ are equation of the second fit; we use the constraint $\gamma = x_0(\alpha - \beta)$ for continuity. In practice, k is set to be very large to give a sharp transition between fits.

3.2 Direct Quantitative Electron Density Reconstruction

Our quantitative method combines (5) with (2) and infers ρ_e from the model through optimisation. Although this is a complicated task explicitly, we will show how through our choice of fitting functions, simplifications can yield an efficient optimisation object, and we point towards simple algorithms for solving it.

Firstly, combining (5) with (2) results in the relation

$$\sum_{j=1}^{N_\xi} b_i(\xi_j) \exp(-[\Phi\boldsymbol{\mu}(\xi_j)]_i) \approx \sum_{j=1}^{N_\xi} b_i(\xi_j) \exp(-[\Phi\hat{\boldsymbol{\mu}}(\boldsymbol{\rho}_e, \xi_j)]_i) \text{ for } i = 1, \dots, N_{\text{ray}}.$$

If we introduce a function $\psi(\cdot, \cdot)$ to simplify notation as

$$\psi_i(\boldsymbol{\rho}_e, \xi) \equiv b_i(\xi) \exp(-[\Phi\hat{\boldsymbol{\mu}}(\boldsymbol{\rho}_e, \xi)]_i) \text{ for } i = 1, \dots, N_{\text{ray}},$$

we can write the negative log-likelihood (NLL) for the Poisson model as

$$\text{NLL}(\boldsymbol{\rho}_e; \mathbf{y}) = \sum_{i=1}^{N_{\text{ray}}} \sum_{j=1}^{N_\xi} \psi_i(\boldsymbol{\rho}_e, \xi_j) + s_i - y_i \log \left(\sum_{j=1}^{N_\xi} \psi_i(\boldsymbol{\rho}_e, \xi_j) + s_i \right). \quad (6)$$

Reconstruction of the electron density map can be performed by finding a $\boldsymbol{\rho}_e$ that minimises (6). We will look at gradient descent methods, for which we

require an expression for the derivative of NLL. If we simplify notation with the following

$$\mathbf{d}(\boldsymbol{\rho}_e) = \mathbf{y} \oslash \left(\sum_{j=1}^{N_\xi} \psi(\boldsymbol{\rho}_e, \xi_j) + \mathbf{s} \right) - \mathbf{1},$$

where \oslash represents component-wise division. An expression for the derivative is then

$$\begin{aligned} \frac{\partial \text{NLL}(\boldsymbol{\rho}_e; \mathbf{y})}{\partial \boldsymbol{\rho}_e} = & (\mathbf{1} - \mathbf{f}(\boldsymbol{\rho}_e)) \odot \mathbf{u}_\alpha(\boldsymbol{\rho}_e) \odot \boldsymbol{\Phi}^T \left[\sum_{j=1}^{N_\xi} \alpha(\xi_j) \psi(\boldsymbol{\rho}_e, \xi_j) \odot \mathbf{d}(\boldsymbol{\rho}_e) \right] \\ & + \mathbf{f}(\boldsymbol{\rho}_e) \odot \mathbf{u}_\beta(\boldsymbol{\rho}_e) \odot \boldsymbol{\Phi}^T \left[\sum_{j=1}^{N_\xi} \beta(\xi_j) \psi(\boldsymbol{\rho}_e, \xi_j) \odot \mathbf{d}(\boldsymbol{\rho}_e) \right] \\ & + \mathbf{u}_\gamma(\boldsymbol{\rho}_e) \odot \boldsymbol{\Phi}^T \left[\sum_{j=1}^{N_\xi} \gamma(\xi_j) \psi(\boldsymbol{\rho}_e, \xi_j) \odot \mathbf{d}(\boldsymbol{\rho}_e) \right], \end{aligned} \quad (7)$$

where $\boldsymbol{\Phi}^T$ represents a transpose of the system matrix or ‘back-projection’, and \odot is component-wise multiplication. $\mathbf{u}_\alpha(\cdot)$, $\mathbf{u}_\beta(\cdot)$ and $\mathbf{u}_\gamma(\cdot)$ are factors from the derivative of (5) with respect to $\boldsymbol{\rho}_e$. If $\mathbf{f}(\cdot)$ were instead independent on $\boldsymbol{\rho}_e$, then we would have $\mathbf{u}_\alpha(\cdot) = \mathbf{1}$, $\mathbf{u}_\beta(\cdot) = \mathbf{1}$ and $\mathbf{u}_\gamma(\cdot) = \mathbf{0}$, which would cancel the third term in (7) for faster computation, and we have noted this to have negligible difference from the exact form for very high k , as is desirable.

With the necessary tools, we can write quantitative reconstruction as

$$\hat{\boldsymbol{\rho}}_e = \underset{\boldsymbol{\rho}_e \in \mathcal{C}}{\text{argmin}} \text{NLL}(\boldsymbol{\rho}_e; \mathbf{y}) + \lambda R(\boldsymbol{\rho}_e), \quad (8)$$

where $R(\cdot)$ is some regularisation function, and \mathcal{C} is a set of box constraints on $\boldsymbol{\rho}$ so that $0 \leq \rho_i \leq \zeta$ for $i = 1, \dots, N_{\text{vox}}$, where ζ is the maximum allowable electron density value. In the experimental section, we use the TV regularisation function for $R(\cdot)$ [9] to promote a piece-wise constant image.

We note that the objective function in (8) is both non-linear and nonconvex, so its minimisation should be treated carefully, though there exist several gradient descent algorithms that may be invoked to minimise it, such as VMFB [15] or paraboloid surrogates [16] if only smooth $R(\cdot)$ is used.

4 Experimentation

4.1 Data

The data we used was derived from the ‘Adult Reference Computational Phantoms’ [17]—a segmented structure with defined biological tissue and elemental composition—from which we isolated the bladder region. A single slice from this data is shown in Figure 2a, where the voxel intensities simply encode the material type. For

example, 30 is muscle tissue, 50 is fat, 10 the femoral head, 52 the rectum, 53 urine and 1 is air. To simulate a planning protocol, we also synthesised a prior image from this data by non-rigidly deforming several of the soft-tissue regions, and this is shown in Figure 2d. The electron densities of the oracle and planning images are shown in Figures 2b and 2e respectively.

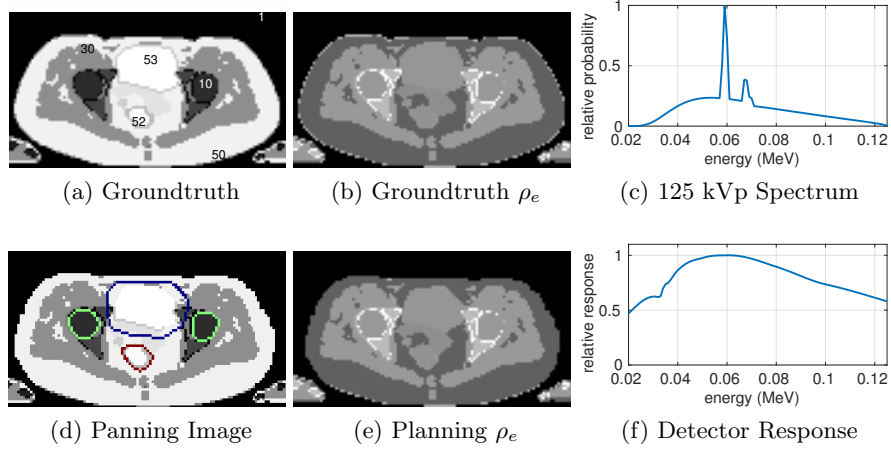


Figure 2: Experimental data used: (a) is the oracle test image as material index; (b) is the oracle electron density with grey scale $[0.8, 1.2]$; (c) is the 125 kVp source spectrum used; (d) is the simulated planning image with contours shown for the PTV, GI tract and femoral heads; (e) is the planning image electron density with grey scale $[0.8, 1.2]$; and (f) is the detector response function used.

To generate CBCT measurements, we used the Monte-Carlo engine, Gate [18], with a total of 2×10^{10} photons over 160 projection angles. To try and closely approximate a real clinical CBCT scanner, we derived the filtered energy spectrum and detector response, shown in Figures 2c and 2f from a Varian TrueBeam™ system (Varian Medical Systems, Palo Alto CA, USA). We also matched the geometry, source profile from offset detector bow-tie filter to its ‘half-fan’ mode, and included a model of its trans-axial focused scatter collimation grid.

We used a number of different methods to map these measurements into electron density, where we utilised CBCT projection and back-projection operators for Φ and Φ^T from the Michigan Image Reconstruction Toolbox [19]. From relative electron density (normalised to water), we then calculated dose distributions using the matRad Toolbox (German Cancer Research Centre) for photons and protons. In both cases, we optimised for delivering a dose of 64 Gy (J/kg) into a planning target volume (PTV) for bladder cancer [20], where we added a margin of ~ 1 cm around the planning image bladder—this is shown

by the contour (blue) in Figure 2d, and applied identical beams to each other method under test and the oracle image.

To assess the various methods, we calculated the electron density accuracy as the error euclidean norm to the oracle, and also the mean and maximum absolute deviations in dose estimation within the PTV, femoral heads and GI tract, the latter of which are sensitive to excessive dose and should be carefully monitored.

4.2 Methods Under Test

- Plan: using the planning image directly, with its oracle bone registration and matching.
- FDK: ‘analytic’ FBP reconstruction for CBCT [21] using prior Monte-Carlo scatter correction [22] and water-based beam hardening compensation [3] as preprocessing, with Hann windowing to suppress noise amplification, weighting for offset detector [23], and non-linear calibration to electron density [1] as postprocessing.
- PWLS: ‘statistical iterative’ reconstruction with (4), with same pre- and postprocessing as FDK, and Total Variation (TV) regularisation—an equivalent objective model as in [24].
- PolySIR: polyenergetic reconstruction model from [5] with oracle bone segmentation, prior Monte-Carlo scatter expectation estimation [22], TV regularisation, and electron density calibration [1] as postprocessing.
- ρ_e CT: our proposed quantitative approach with prior Monte-Carlo scatter expectation estimation [22], TV regularisation, and no postprocessing.

In all cases of iterative reconstruction—PWLS, PolySIR and ρ_e CT—the regularisation constant for TV, λ , was individually numerically optimised to give maximum soft tissue accuracy, which ensured fairness throughout the different data-fidelity terms.

4.3 Results

Results are summarised in Table 1, and illustrated visually in Figures 3 and 4.


From the numerical results, it is evident that our proposed method has a significant numerical advantage in calculating accurate electron density, with its error norm 35% lower than the next best tested. This gain appears to translate into average dose calculation, where it is the best performing across the board. There are a couple of instances where it is outperformed in the worst case however, in Table 2, which we predict is due to the less smooth soft-tissue structures in the optimal TV regularisation in ρ_e CT than PWLS. Indeed, if we manually increase the regularisation parameter λ , this maximum error does drop below PWLS, but we did not include this result to maintain fairness and consistency. Additionally, from the substantially better performance in the fem. heads and GI tract for proton, we also think this trade-off is acceptable. It may be that invoking varying regularisation for the different tissue classes could be key to pushing our performance even further.

Table 1: Quantitative results 1: norm ρ_e error and mean dose error into planning target volume (PTV), femoral heads (fem. heads) and gastrointestinal tract (GI tract)

scheme	ρ_e error norm	mean photon error (Gy)			mean proton error (Gy)		
		PTV	fem. heads	GI tract	PTV	fem. heads	GI tract
Plan	56.3	0.587	0.587	0.320	0.321	0.438	0.0976
FDK	135	7.52	3.06	2.79	0.417	3.96	2.90
PWLS	30.7	0.288	0.198	0.229	0.102	0.156	0.141
PolySIR	30.0	0.321	0.209	0.227	0.0862	0.131	0.144
ρ_e CT	19.5	0.275	0.182	0.209	0.0809	0.0696	0.0683

Table 2: Quantitative results 2: worst case dose error—maximum absolute difference

scheme	max. photon error (Gy)			max. proton error (Gy)		
	PTV	fem. heads	GI tract	PTV	fem. heads	GI tract
Plan	48.9	47.1	25.0	39.4	45.8	33.5
FDK	10.2	15.6	15.7	39.5	61.9	61.8
PWLS	3.51	7.15	9.87	9.04	11.5	12.0
PolySIR	2.78	5.92	8.48	24.6	24.4	17.7
ρ_e CT	2.59	5.84	8.86	11.1	5.50	5.03

Figure 3: Visual illustration of photon dose results: oracle dose maps have colour scale $[0.2, 64]$; dose differences are absolute and have colour scale $[0.2, 5]$; and all electron densities are shown with grey scale $[0, 1.8]$.

5 Conclusion

In conclusion, we have successfully introduced a method for inferring electron density directly from raw CT measurements, with additive scatter and a single




Figure 4: Visual illustration of proton dose results, with same colour scales as Figure 3.

polyenergetic source. From our preliminary experiments, it significantly outperforms existing approaches, and appears to be an accurate quantitative tool, even in the severely low-dose and high scatter setting under test. There are certainly many avenues for further investigation, such as a relation to the underlying physical phenomena and dual energy systems, and more advanced regularisation strategies. We also predict it could be readily extended to model metal objects and implants also.

6 Acknowledgements

The authors would like to thank the Maxwell Advanced Technology Fund, EPSRC DTP studentship funds and ERC project: C-SENSE (ERC-ADG-2015-694888) for supporting this work.

References

1. U Schneider, E Pedroni, and A Lomax. The calibration of CT Hounsfield units for radiotherapy treatment planning. *Phys. Med. Biol.*, 41(1):111–124, 1996.
2. T S Curry, J E Dowdey, and R C Murry. *Christensen's Physics of Diagnostic Radiology*. 1990.
3. P M Joseph and R D Spital. A method for correcting bone induced artifacts in computed tomography scanners., 1978.
4. J A Fessler. Fundamentals of CT Reconstruction in 2D and 3D. In *Compr. Biomed. Phys.*, pages 263–295. Elsevier, 2014.
5. I A Elbakri and J A Fessler. Statistical image reconstruction for polyenergetic X-ray computed tomography. *IEEE Trans. Med. Imaging*, 21(2):89–99, feb 2002.

6. I A Elbakri and J A Fessler. Segmentation-free statistical image reconstruction for polyenergetic x-ray computed tomography with experimental validation. *Phys. Med. Biol.*, 48(15):2453–2477, aug 2003.
7. B De Man, J Nuyts, P Dupont, G Marchal, and P Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Trans. Med. Imaging*, 20(10):999–1008, 2001.
8. C Bouman and K Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. Image Process.*, 2(3):296–310, jul 1993.
9. L I Rudin, S Osher, and E Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D Nonlinear Phenom.*, 60(1-4):259–268, nov 1992.
10. I Daubechies, M Fornasier, and I Loris. Accelerated Projected Gradient Method for Linear Inverse Problems with Sparsity Constraints. *J. Fourier Anal. Appl.*, 14(5-6):764–792, dec 2008.
11. J H Siewerdsen and D A Jaffray. Cone-beam computed tomography with a flat-panel imager: Magnitude and effects of x-ray scatter. *Med. Phys.*, 28(2):220, 2001.
12. H K Tuy. An inversion formula for cone-beam reconstruction. *SIAM J. Appl. Math.*, 43(3):546–552, 1983.
13. Z Chang, R Zhang, J-B Thibault, K Sauer, and C Bouman. Statistical x-ray computed tomography imaging from photon-starved measurements. *SPIE Comput. Imaging*, 9020:90200G, 2014.
14. ICRP Publication 89. Basic Anatomical and Physiological Data for Use in Radiological Protection Reference Values. *Ann. ICRP*, 32 (3-4), 2002.
15. E Chouzenoux, J C Pesquet, and A Repetti. Variable Metric Forward-Backward Algorithm for Minimizing the Sum of a Differentiable Function and a Convex Function. *J. Optim. Theory Appl.*, 162(1):107–132, 2014.
16. H Erdogan and J A Fessler. Monotonic algorithms for transmission tomography. *IEEE Trans. Med. Imaging*, 18(9):801–814, 1999.
17. ICRP Publication 110. Adult Reference Computational Phantoms. *Ann. ICRP*, 39 (2), 2009.
18. S Jan, D Benoit, E Becheva, T Carlier, F Cassol, P Descourt, T Frisson, L Grevillot, L Guigues, L Maigne, C Morel, Y Perrot, N Rehfeld, D Sarrut, D R Schaart, S Stute, U Pietrzyk, D Visvikis, N Zahra, and I Buvat. GATE V6: a major enhancement of the GATE simulation platform enabling modelling of CT and radiotherapy. *Phys. Med. Biol.*, 56(4):881–901, 2011.
19. J A Fessler. Image Reconstruction Toolbox.
20. A Barrett, S Morris, J Dobbs, and T Roques. *Practical Radiotherapy Planning Fourth Edition*. 2009.
21. L A Feldkamp, L C Davis, and J W Kress. Practical cone-beam algorithm. *J. Opt. Soc. Am. A*, 1(6):612, 1984.
22. Y Xu, T Bai, H Yan, L Ouyang, A Pompos, J Wang, L Zhou, S B Jiang, and X Jia. A practical cone-beam CT scatter correction method with optimized Monte Carlo simulations for image-guided radiation therapy. *Phys. Med. Biol.*, 60(9):3567–3587, may 2015.
23. G Wang. X-ray micro-CT with a displaced detector array. *Med. Phys.*, 29(7):1634, 2002.
24. K Choi, J Wang, L Zhu, T-S Suh, S Boyd, and L Xing. Compressed sensing based cone-beam computed tomography reconstruction with a first-order methoda). *Med. Phys.*, 37(9):5113–5125, aug 2010.