One-step iterative reconstruction approach based on eigentissue decomposition for spectral photon-counting computed tomography

Mikaël Simard[a,*] **and Hugo Bouchard**[a,b,c]

[a]Université de Montréal, Département de physique, Montréal, Québec, Canada
[b]Centre de recherche du Centre hospitalier de l'Université de Montréal, Montréal, Québec, Canada
[c]Centre hospitalier de l'Université de Montréal (CHUM), Département de radio-oncologie, Montréal, Québec, Canada

Abstract

Purpose: We propose a one-step tissue characterization method for spectral photon-counting computed tomography (SPCCT) using eigentissue decomposition (ETD), tailored for highly accurate human tissue characterization in radiotherapy.

Methods: The approach combines a Poisson likelihood, a spatial prior, and a quantitative prior constraining eigentissue fractions based on expected values for tabulated tissues. There are two regularization parameters: α for the quantitative prior, and β for the spatial prior. The approach is validated in a realistic simulation environment for SPCCT. The impact of α and β is evaluated on a virtual phantom. The framework is tested on a virtual patient and compared with two sinogram-based two-step methods [using respectively filtered backprojection (FBP) and an iterative method for the second step] and a post-reconstruction approach with the same quantitative prior. All methods use ETD.

Results: Optimal performance with respect to bias or RMSE is achieved with different combinations of α and β on the cylindrical phantom. Evaluated in tissues of the virtual patient, the one-step framework outperforms two-step and post-reconstruction approaches to quantify proton-stopping power (SPR). The mean absolute bias on the SPR is 0.6% (two-step FBP), 0.6% (two-step iterative), 0.6% (post-reconstruction), and 0.2% (one-step optimized for low bias). Following the same order, the RMSE on the SPR is 13.3%, 2.5%, 3.2%, and 1.5%.

Conclusions: Accurate and precise characterization with ETD can be achieved with noisy SPCCT data without the need to rely on post-reconstruction methods. The one-step framework is more accurate and precise than two-step methods for human tissue characterization.

 [DOI: 10.1117/1.JMI.9.4.044003]

Keywords: spectral photon-counting computed tomography; tissue characterization; radiotherapy; one-step reconstruction; quantitative imaging; eigentissue decomposition.

Paper 22077GR received Apr. 10, 2022; accepted for publication Jul. 1, 2022; published online Jul. 27, 2022.

1 Introduction

Personalized radiation therapy requires estimates of tissue physical properties for the calculation of treatment dose. The common method for obtaining such properties in a clinical environment is single-energy computed tomography (SECT), where a single CT number is used to infer physical parameters of interest for a given radiotherapy modality. For instance, in megavoltage photon therapy and proton therapy, the relative electron density (ρ_e) and the relative proton stopping power (SPR) are respectively required. SECT struggles to provide accurate physical properties,[1]

*Address all correspondence to Mikaël Simard, mikael.simard@umontreal.ca

2329-4302/2022/$28.00 © 2022 SPIE

with errors of up to 3.5% on the SPR.[2] This intrinsically limits the accuracy achievable with treatment plans. In the last decade, multi-energy CT (MECT) was explored to improve the accuracy of tissue characterization for radiotherapy. For instance, dual-energy CT (DECT) can provide a more accurate estimation of ρ_e and SPR for radiotherapy using information from two CT numbers acquired at different energies.[1,3] More recently, another MECT modality, spectral photon-counting computed tomography (SPCCT), which uses photon counting detectors to provide up to four energy-resolved measurements, has been shown to provide further benefits compared to DECT for tissue characterization and reducing proton range uncertainties.[4–6]

In those studies, the benefits of MECT were obtained using post-reconstruction methods, where tissue characterization is performed on reconstructed CT number maps. The obtained benefits typically require the use of bespoke methods for tissue characterization with MECT. For instance, eigentissue decomposition (ETD) was introduced[7] as an optimized decomposition basis to maximize the reconstruction accuracy of human tissue elemental composition with MECT. Furthermore, to deal with noisy MECT data, ETD was reformulated into a maximum *a posteriori* (MAP) framework[8] where eigentissue fractions are regularized to lie close to expected values of tabulated human tissues. This regularization strategy is referred to as quantitative regularization. Combining ETD and quantitative regularization for MECT data results in efficient denoising performance and allows material decomposition with more than two basis materials for human tissues, which has been shown to be mandatory to maximize quantitative accuracy of estimated radiotherapy-related parameters.[6–9]

For post-reconstruction characterization, the achievable accuracy of estimated physical parameters is limited by reconstruction artifacts such as beam hardening. As a solution, two-step sinogram-based tissue characterization methods were used to produce physical parameter maps directly from sinogram data.[10,11] The two-step framework fits physical parameters sinograms from sinogram data, then reconstructs physical parameters maps using either a filtered backprojection (FBP) or an iterative method. While useful to mitigate artifacts, two-step methods are susceptible to noise due to the instability (ill-conditioning) of the decomposition step, resulting in material projections corrupted by a high amount of correlated noise.[12,13] Furthermore, the two-step framework cannot fully take advantage of the ETD parametrization, as only two eigentissue sinograms can be solved during the first step due to the similarity between eigentissue basis functions and the lack of regularization in sinogram space.[11] This overall limits the intrinsic accuracy attainable with the two-step method for SPCCT; e.g., the benefits of using SPCCT against DECT for accurate tissue characterization with ETD were lost due to the use of two eigentissues.[11]

One solution is to implement a one-step tissue characterization framework, where reconstruction is formulated as a large-scale optimization problem and physical properties are estimated directly from sinogram data. One-step methods were introduced in the last decade, first for DECT[14,15] and more recently for SPCCT,[16,17] with a focus on K-edge imaging. As constraints are formulated in the image domain, the approach is free from the above-mentioned limitations of the two-step framework and may allow more accurate and precise tissue characterization.

Therefore, the main objective of this study is to develop a novel one-step reconstruction framework with the ETD parametrization, to recapture the benefits of using SPCCT for human tissue characterization, as obtained with post-reconstruction ETD.[5] More specifically, we propose a unified framework that incorporates ETD, spatial regularization as well as the quantitative prior[8] into a one-step reconstruction approach. The aim is to maximize the quantitative accuracy and precision of reconstructed physical parameters of human tissues relevant for radiotherapy without the limitations of post-reconstruction methods. Each regularizer contributes toward increased performance. The framework is validated in a realistic simulation environment for SPCCT with four energy bins on a cylindrical phantom and a virtual patient. The performance of the algorithm with various levels of quantitative and spatial regularization is first evaluated. Then, using three optimized sets of regularization parameters, it is compared to (1) two implementations of the two-step framework using the ETD parametrization, and (2) the post-reconstruction MAP implementation of ETD.[8] As such, a second aim of this work is to report the differences between the four tissue characterization schemes for the task of extracting highly accurate radiotherapy-related parameters from SPCCT data. This was only reported previously for DECT and a restricted characterization methods, each with its own parametrization.[18]

By using the same parametrization of the attenuation coefficient for all methods, we can provide insight into the intrinsic accuracy of all tissue characterization strategies.

2 Theory

2.1 *Eigentissue Decomposition*

2.1.1 *Principles of eigentissue decomposition*

Let j denote the index of a pixel in the image to reconstruct, and J the total number of pixels. Let $(\delta\epsilon, \ldots, N_\epsilon\delta\epsilon)$ be a uniform energy sampling, where N_ϵ and $\delta\epsilon$ are the number of energies sampled and the energy resolution, respectively; ϵ is the index referring to photon energies. The energy-dependent x-ray attenuation coefficient of a tissue located at the j'th pixel and for the ϵ^{th} energy, $\mu_{j\epsilon}$, is expressed as the product of the electron density relative to water of the tissue at the j'th pixel, $\rho_{e,j}$, and a quantity proportional to its energy-dependent electronic cross section, $g_{j\epsilon}$

$$\mu_{j\epsilon} = \rho_{e,j} g_{j\epsilon}, \tag{1}$$

where $g_{j\epsilon}$ is, by definition,[11] equal to

$$g_{j\epsilon} = N_A \left(\frac{Z}{A}\right)_w \rho_w \sigma_{j\epsilon}, \tag{2}$$

where $\left(\frac{Z}{A}\right)_w$ is the ratio of atomic to mass number of water, ρ_w is the mass density of water, N_A is Avogadro's number, and $\sigma_{j\epsilon}$ is the energy-dependent electronic cross section of the tissue, which includes all relevant physical effects in the diagnostic energy range, and can be calculated with the law of mixtures considering the elemental cross sections $\sigma_{m,\epsilon}$. The attenuation coefficient of Eq. (1) can also be expressed directly with the law of mixtures as a sum of its elemental contributions

$$\mu_{j\epsilon} = \sum_{m=1}^{M} x_{jm} g_{m\epsilon}, \tag{3}$$

where M denotes the number of elements m which compose the tissue, x_{jm} is the partial electron density of the m'th element in the tissue located at the j'th pixel, and $g_{m\epsilon}$ is proportional to the energy-dependent electronic cross section of the m'th element, and is obtained from Eq. (2) using the electronic cross section of the corresponding element, $\sigma_{m\epsilon}$ (i.e., $\sigma_{j\epsilon}$ is replaced by $\sigma_{m\epsilon}$). For brevity, $g_{m\epsilon}$ is now referred to as an electronic cross section. By construction, $\sum_{m=1}^{M} x_{jm} = \rho_{e,j}$. For a given pixel, the M values $x_{j1}, \ldots, x_{jM}$ thus span the electron density and the elemental composition.[7] Full tissue characterization implies the extraction of all x_{jm} from the N energy-resolved measurements with MECT. For human tissues, typically $M \geq N$, and the elemental basis cannot be used for tissue characterization with MECT.

This is solved by representing human tissues as a sum of $K \leq N$ basis materials, such as water and bone,[19,20] water, lipids and proteins,[21] or water and iodine for contrast-enhanced imaging.[22,23] However, such decomposition bases cannot reproduce the full variability of the elemental composition for human tissues.[7] This can lead, for human tissues, to biased estimates of ρ_e and SPR,[11,24] and indicates that such bases are not optimal to characterize human tissues. Lalonde and Bouchard[7] performed principal component analysis (PCA) on the set of elemental partial electron densities from a database of human tissues[25] to extract a more optimal decomposition basis, named the eigentissue basis, to represent the full range of variation in human tissue composition. Using the $K \leq N$ most meaningful eigentissues for decomposition, Eq. (3) is made compatible with MECT by using the model

$$\hat{\mu}_{j\epsilon} = \sum_{k=1}^{M} y_{jk} g_{k\epsilon}^{\mathrm{ET}} \approx y_0 g_{0\epsilon}^{\mathrm{ET}} + \sum_{k=1}^{K} y_{jk} g_{k\epsilon}^{\mathrm{ET}}, \tag{4}$$

where y_{jk} represents the partial electron density of the k'th eigentissue in the j'th pixel, and $g^{\mathrm{ET}}_{k\epsilon}$ is its corresponding energy-dependent electronic cross section. The residual terms y_0 and $g^{\mathrm{ET}}_{0\epsilon}$ represent the contribution of the $(K+1)$'th to the M'th eigentissues, which are assumed constant and equal, for every tissue, to the average over all tissues from the reference database. This explains the lack of spatial dependency for y_0. More details can be found in previous studies.[7,11] The electron density is estimated as $\rho_{e,j} \approx y_0 + \sum_{k=1}^{K} y_{jk}$. Equation (4) is the image-based formulation of ETD, and the K parameters to estimate at each image pixel j is a vector of partial electron densities $\mathbf{y}_j = (y_{j1} \dots y_{jK})^{\mathrm{T}}$. Using the transformation matrix between the elemental and eigentissue partial electron densities obtained with PCA, the y_{jk} can be mapped back to the elemental composition of the sample.[7]

2.1.2 *Adaptation to sinogram space*

Let the index i represent a ray projected at angle θ and detected at location ξ. There are I sinogram pixels (I = number of projections $N_p\times$ number of detectors N_d). Let also $\psi_{n\epsilon}$ denote the normalized effective x-ray spectrum, which includes contributions of the source and detector response, where n is an index for a specific combination of source and/or detector response; there are N different source/detector response combinations. The detected x-ray photon counts from the i'th ray coming from the effective spectrum $\psi_{n\epsilon}$, with no object (blank scan), is N^0_{in}, while the detected x-ray photon counts after propagating through an object is N_{in}. The transmission of the i'th ray for the n'th spectrum, t_{in}, is the ratio $t_{in} = \frac{N_{in}}{N^0_{in}}$. The transmission model, $\hat{t}_{in}$, is the polyenergetic forward model

$$\hat{t}_{in} = \sum_{\epsilon=1}^{N_\epsilon} \psi_{n\epsilon} \mathrm{soft}\exp\left(-\sum_{j=1}^{J} a_{ij}\mu_{j\epsilon}\right), \tag{5}$$

where a_{ij} are elements of the projection matrix $\mathbf{A}$, of size $I \times J$, which provides the path length of the i'th ray through the j'th pixel of the geometry. The soft exponential function, used instead of the conventional exponential function, is

$$\mathrm{softexp}(r) = \begin{cases} e^{-r} & \text{if } r \geq 0, \\ 1 - r & \text{if } r < 0. \end{cases} \tag{6}$$

This dampens the effect of possible negative attenuation obtained in early iterations, which often leads to convergence issues, as proposed by Sidky et al.[26] Inserting the ETD parametrization of $\mu_{j\epsilon}$ from Eq. (4) gives

$$\hat{t}_{in}(y_{jk}) = \sum_{\epsilon=1}^{N_\epsilon} \psi_{n\epsilon}\, \mathrm{softexp}\left(g^{\mathrm{ET}}_{0\epsilon} Y_{i0} + \sum_{k=1}^{K} g^{\mathrm{ET}}_{k\epsilon} \sum_{j=1}^{J} a_{ij} y_{jk}\right). \tag{7}$$

As y_0 is constant, we defined $Y_{i0} = y_0 \sum_{j=1}^{J} a_{ij}$, which is the Radon transform of the residual eigentissue, assuming that all image pixels have the residual y_0.

2.2 *A Bayesian Framework for Image Reconstruction*

Let $\mathbf{y_k} = (y_{1k} \dots y_{Jk})^{\mathrm{T}}$ be the $J \times 1$ vector that contains all pixels of the image for the k'th eigentissue fraction. All $\mathbf{y_k}$ are regrouped in the $JK \times 1$ vector of joint partial electron densities $\mathbf{y} = (\mathbf{y_1^T} \dots \mathbf{y_K^T})^{\mathrm{T}}$. Similar vectors are also defined for the transmission measurements. $\mathbf{t_n} = (t_{1n} \dots t_{In})^{\mathrm{T}}$ is the $I \times 1$ transmission vector which includes all I transmission measurements for the n'th overall spectrum, and the $IN \times 1$ joint transmission vector is $\mathbf{t} = (\mathbf{t_1^T} \dots \mathbf{t_N^T})^{\mathrm{T}}$. The MAP estimator of the joint partial electron density vector, $\hat{\mathbf{y}}$, is

$$\hat{\mathbf{y}} = \arg\max_{\mathbf{y}} P(\mathbf{y}|\mathbf{t}) = \arg\max_{\mathbf{y}} \left[\frac{P(\mathbf{t}|\mathbf{y})P(\mathbf{y})}{P(\mathbf{t})}\right]. \tag{8}$$

$P(\mathbf{y}|\mathbf{t})$ is the posterior probability of the image containing physical properties defined in $\mathbf{y}$ knowing that the transmission values $\mathbf{t}$ are measured. The right-hand side is obtained using Bayes' theorem, where $P(\mathbf{t}|\mathbf{y})$ is the likelihood and $P(\mathbf{y})$ is the prior probability of $\mathbf{y}$. It is more convenient to minimize the negative logarithm of the posterior

$$\hat{\mathbf{y}} = \arg\min_{\mathbf{y}} [\ell_L(\mathbf{y}) + \ell_P(\mathbf{y})] = \arg\min_{\mathbf{y}} L(\mathbf{y}), \tag{9}$$

where $\ell_L(\mathbf{y}) = -\ln(P(\mathbf{t}|\mathbf{y}))$ is the negative log-likelihood and $\ell_P(\mathbf{y}) = -\ln(P(\mathbf{y}))$ is the negative log-prior. $L(\mathbf{y})$ is the total cost function to minimize. The likelihood and prior terms are defined in Secs. 2.2.1 and 2.2.2.

2.2.1 *Likelihood*

Poisson statistics are used to describe energy-resolved photon counts on a detector pixel, t_{in}. Each count is assumed to follow a Poisson distribution and is independent from the other events. This is realistic for spectral photon counting CT, as (1) the arrival of Poisson-distributed photons is directly counted on the detector and (2) the low electronic noise does not importantly affect the counting statistics. The measurement joint probability (the likelihood) is

$$P(\mathbf{t}|\mathbf{y}) = \prod_{n=1}^{N}\prod_{i=1}^{I} P(t_{in}|\mathbf{y}) = \prod_{n=1}^{N}\prod_{i=1}^{I} \exp(-\hat{t}_{in}(\mathbf{y}))(\hat{t}_{in}(\mathbf{y}))^{t_{in}} \frac{1}{t_{in}!}, \tag{10}$$

where $\hat{t}_{in}(\mathbf{y})$ is the model [Eq. (7)] of the transmission. The corresponding negative log-likelihood is

$$\ell_L(\mathbf{y}) = \sum_{n=1}^{N}\sum_{i=1}^{I} (\hat{t}_{in}(\mathbf{y}) - t_{in}\ \ln(\hat{t}_{in}(\mathbf{y}))). \tag{11}$$

2.2.2 *Prior*

For the prior $P(\mathbf{y})$, by construction, eigentissue partial electron densities are expected to vary within a specific range for human tissues,[8] and fitted values can be constrained accordingly. This constraint increases the robustness of image-based methods with noisy MECT data.[5,9,27] As regularization is performed pixel-wise on physical parameters, it is referred to as the quantitative prior. Spatial regularization[28] on maps of $\mathbf{y}$ is also used in this work and is referred to as the spatial prior. The total prior is the product of quantitative and spatial priors.

Quantitative prior. Deriving eigentissues using PCA on a database of P tabulated human tissues[25] produces a set of known eigentissue partial electron densities for each tissue in the reference database, each represented by a $K \times 1$ vector $\mathbf{y}_{\text{ref},p}$, $p \in [1,\ldots,P]$, with elements $y_{\text{ref},pk}$, $k \in [1,\ldots,K]$. Therefore, there is prior knowledge available on the distribution of the y_{jk}, if one assumes that the scanned object is made of human tissues of similar composition to those of the reference database. It is expected that, for the j'th pixel, the set of the K eigentissue partial electron densities $\mathbf{y}_j = (\, y_{j1} \quad \ldots \quad y_{jK}\,)^{\mathrm{T}}$ will lie in proximity to the partial electron densities of one tissue in the reference database, $\mathbf{y}_{\text{ref},p}$. This proximity is modeled with a kernel density estimator (KDE) using a Gaussian kernel

$$P_Q(\mathbf{y}_j) = \sum_{p=1}^{P} \frac{w_{jp}}{\sqrt{(2\pi)^K |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{y}_j - \mathbf{y}_{\text{ref},p})^{\mathrm{T}} \boldsymbol{\Sigma}^{-1} (\mathbf{y}_j - \mathbf{y}_{\text{ref},p})\right). \tag{12}$$

The Gaussian kernel accounts for variability in the elemental composition and electronic density of the patient's tissues with respect to the reference database's tissues. The pixel-

dependent weight w_{jp} of each Gaussian is defined later. The $K \times K$ matrix $\mathbf{\Sigma}$ is the kernel bandwidth. By construction, eigentissues are uncorrelated within reference tissues. As proposed by Lalonde et al.,[8] Σ is the diagonal and defined as

$$\mathbf{\Sigma} = \alpha \cdot \mathbf{\Sigma}_{\mathbf{y}} = \alpha \cdot \begin{pmatrix} \sigma_{y_1}^2 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \sigma_{y_K}^2 \end{pmatrix}, \tag{13}$$

where $\mathbf{\Sigma}_{\mathbf{y}}$ is a $K \times K$ matrix with elements σ_{y_k} representing the standard deviation in the k'th eigentissue over all tissues of the database. To provide flexibility in cases of high noise or larger compositional variation, a regularization parameter α is introduced. Considering Eq. (13), the probability density function $P_Q(\mathbf{y}_j)$ reduces to

$$P_Q(\mathbf{y}_j) = \frac{1}{\sqrt{(2\pi\alpha)^K \prod_{k=1}^{K} \sigma_{y_k}^2}} \sum_{p=1}^{P} w_{jp} \exp\left(-\frac{d_{jp}^{(y)}}{2\alpha}\right), \tag{14}$$

where $d_{jp}^{(y)}$ represents the squared generalized distance between pixel $\mathbf{y}_j$ and the p'th reference tissue $\mathbf{y}_{\text{ref},p}$

$$d_{jp}^{(y)} = \sum_{k=1}^{K} \frac{(y_{jk} - y_{\text{ref},pk})^2}{\sigma_{y_k}^2}. \tag{15}$$

For this work, we develop a prior assuming that each tissue is equiprobable, such that optimization is not biased towards a particular tissue. Note that this may not hold true for specific anatomical sites; in that case, we recommend to tailor the reference database and prior probabilities to reflect the site being studied. For a given $\mathbf{y}_j$, a simple formulation of w_{jp} leading to equal probability $P_Q(\mathbf{y}_j)$ for each reference tissue is

$$w_{jp} = \begin{cases} 1 & \text{if } p = \arg\min_p (d_{jp}^{(y)}); \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

Equation (16) is not twice differentiable; a smooth approximation to Eq. (16) is proposed, using the transform

$$\pi_{jp} = (d_{jp}^{(y)} + \varepsilon)^{-\nu}, \tag{17}$$

which amplifies the differences between distances $d_{jp}^{(y)}$ of reference tissues, if ν is large (in this work, $\nu = 60$). The constant ε is introduced to ensure no division by zero occurs if $\mathbf{y}_j$ coincides with a reference tissue; it is fixed to $\varepsilon = 10^{-5}$. It should be noted that varying ν and ε within a reasonable range has no visible impact on the quality of reconstructed physical parameters. π_{jp} is normalized by its sum over all reference tissues to obtain a final weight w_{jp} between 0 and 1

$$w_{jp} = \frac{(d_{jp}^{(y)} + \varepsilon)^{-\nu}}{\sum_{p'=1}^{P} (d_{jp'}^{(y)} + \varepsilon)^{-\nu}}. \tag{18}$$

Equation (18) provides a smooth approximation to Eq. (16). The joint quantitative prior for the entire image is

$$P_Q(\mathbf{y}) = P_Q(\mathbf{y}_1, \dots, \mathbf{y}_j) = \prod_{j=1}^{J} P_Q(\mathbf{y}_j) = \left((2\pi\alpha)^K \prod_{k=1}^{K} \sigma_{y_k}^2\right)^{-J/2} \prod_{j=1}^{J} \sum_{p=1}^{P} w_{jp} \exp\left(-\frac{d_{jp}^{(y)}}{2\alpha}\right). \tag{19}$$

Removing constant terms independent of y_{jk}, the negative logarithm of the quantitative prior on the eigentissue partial electron densities, $\ell_{P_Q}(\mathbf{y})$ is

$$\ell_{P_Q}(\mathbf{y}) = -\sum_{j=1}^{J} \ln\left(\sum_{p=1}^{P} w_{jp} \exp\left(-\frac{1}{2\alpha}\sum_{k=1}^{K}\frac{(y_{jk}-y_{\mathrm{ref},pk})^2}{\sigma_{y_k}^2}\right)\right). \tag{20}$$

Considering two eigentissues ($K = 2$) and $\alpha = 1000$, Fig. 1 shows $\ell_{P_Q}(\mathbf{y})$.

Spatial prior. For image reconstruction purposes, the negative logarithm of the spatial prior, $\ell_{P_S}(\mathbf{y})$, typically takes the form of a penalty function that acts on intensity differences between nearest neighbor pixels[29–31]

$$\ell_{P_S}(\mathbf{y}) = \sum_{k=1}^{K}\beta_k \sum_{j=1}^{J}\sum_{\eta\in\mathcal{N}_j}\omega_{j\eta}\,\phi(y_{jk}-y_{\eta k},\delta_k), \tag{21}$$

where each β_k and δ_k are user-defined regularization parameters which control the strength of the spatial prior for the k'th eigentissue fraction, while $\phi(y_{jk}-y_{\eta k},\delta_k)$ is a user-defined potential function that penalizes, for each of the K materials, differences between the j'th pixel and the set of pixels defined in its neighborhood $\mathcal{N}_j$, noted with the index η. $\omega_{j\eta}$ weights the importance of the interaction between the j'th pixel and its η'th neighbor; it is set as the inverse of the Euclidian distance between the pixels. The four direct neighbors are given a weight of 1, while the four diagonal neighbors are given a weight of $\frac{1}{\sqrt{2}}$. The desired spatial distribution of $\mathbf{y}$ is encoded into the potential function $\phi(y_{jk}-y_{\eta k},\delta_k)$. ϕ should encourage smoothness across neighboring pixels within the same organ/tissue, while accommodating sharp changes in $\mathbf{y}$ due to boundaries between tissues. Many edge-preserving potential functions have been proposed in the literature[28]; in this work, the Huber potential[32] is used for all material channels

$$\phi(y_{jk}-y_{\eta k},\delta_k) = \begin{cases} \frac{(y_{jk}-y_{\eta k})^2}{2}, & |y_{jk}-y_{\eta k}| \le \delta_k; \\ \delta_k|y_{jk}-y_{\eta k}| - \frac{\delta_k^2}{2}, & |y_{jk}-y_{\eta k}| > \delta_k. \end{cases} \tag{22}$$

Finite differences between nearest neighbors are encoded in a $J \times J$ finite difference matrix $\mathbf{D}$. For the set of all pixels corresponding to the k'th eigentissue partial electron density, $\mathbf{y_k}$, the $J \times 1$ vector of finite differences for the image $\mathbf{y_k}$ is calculated as the product $\mathbf{Dy_k}$. Equation (21) is then formulated as

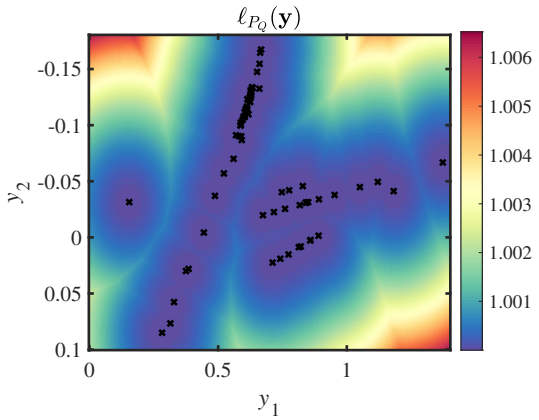

Fig. 1 Negative logarithm of the quantitative prior for $K = 2$ and $\alpha = 1000$. Points illustrated with black crosses are the reference tissues from White and Woodard's database.[25] The cost function is constructed such that all reference tissues have equal probability, while regions of the search space far from reference tissues are penalized.

$$\ell_{P_S}(\mathbf{y}) = \sum_{k=1}^{K} \beta_k \sum_{j=1}^{J} \phi([\mathbf{D}\mathbf{y}_\mathbf{k}]_j, \delta_k), \tag{23}$$

where the sum on neighboring pixels η and the weights $\omega_{j\eta}$ are encoded within $\mathbf{D}$. The notation $[\mathbf{D}\mathbf{y}_\mathbf{k}]_j$ denotes the j'th row (or element) of the vector $\mathbf{D}\mathbf{y}_\mathbf{k}$. An explicit form is $[\mathbf{D}\mathbf{y}_\mathbf{k}]_j = \sum_{c=1}^{J} D_{jc} y_{ck}$. Furthermore, each β_k is defined to produce a similar level of spatial regularization across eigentissue fractions, similar to the quantitative prior

$$\beta_k = \frac{\beta}{\sigma_{y_k}^2}, \tag{24}$$

where β is a global regularization parameter that controls the strength of the spatial regularization. The $\sigma_{y_k}^2$ was defined in Eq. (13). The parameters δ_k define, in the Huber function, the transition between uniform regions of interest and edges between tissues. Similar, they are set as proportional to the expected variability for each eigentissue fraction

$$\delta_k = \delta \sigma_{y_k}. \tag{25}$$

The regularization parameter becomes a unique δ. Equations (24) and (25) allow to reduce the number of free parameters in the spatial prior from $2K$ to 2.

2.2.3 *Cost function*

The cost function $L(\mathbf{y})$ is the negative log of the posterior, obtained by combining the negative log-likelihood of Eq. (11) with the negative log-prior [Eqs. (20) and (23)], assuming independence of the prior terms

$$L(\mathbf{y}) = \ell_L(\mathbf{y}) + \ell_P(\mathbf{y}) = \ell_L(\mathbf{y}) + \ell_{P_Q}(\mathbf{y}) + \ell_{P_S}(\mathbf{y}). \tag{26}$$

2.3 *Optimization Algorithm*

A nonlinear conjugate gradient (CG) algorithm is used in this work, similar to the work of Cai et al.[14] Pseudocode for the optimization algorithm is given in Appendix A. For iteration κ, the update direction $\mathbf{d}^{[\kappa]}$ is guided by the Polak–Ribière scalar β^{PR}, while the update step $\gamma^{[\kappa]}$ is obtained through an approximative line search in the update direction $\mathbf{d}^{[\kappa]}$. The performance of nonlinear CG has been shown to produce similar outcomes compared to other existing algorithms.[28] The CG method finds a solution $\hat{\mathbf{y}}$ to Eq. (9) through the following iterative scheme:

$$\mathbf{y}^{[\kappa+1]} = \mathbf{y}^{[\kappa]} + \gamma^{[\kappa]} \mathbf{d}^{[\kappa]}. \tag{27}$$

Calculation of $\mathbf{d}^{[\kappa]}$ and $\gamma^{[\kappa]}$ at each iteration are detailed in Secs. 2.3.1 and 2.3.2, respectively.

2.3.1 *Update direction* $\mathbf{d}^{[\kappa]}$

The update direction for the first iteration is the steepest descent $\mathbf{d}^{[\kappa=0]} = -\nabla_L(\mathbf{y}^{[\kappa=0]})$, where $\nabla_L(\mathbf{y})$ is the gradient of $L(\mathbf{y})$. Subsequent update directions are calculated via Gram–Schmidt conjugation of the residuals, with the approach of Polak and Ribière[33]

$$\mathbf{d}^{[\kappa+1]} = -\nabla_L(\mathbf{y}^{[\kappa+1]}) + \beta_{\mathrm{PR}}^{[\kappa+1]} \mathbf{d}^{[\kappa]}, \tag{28}$$

where $\beta_{\mathrm{PR}}^{[\kappa+1]}$ is the Polak–Ribière scalar, equal to

$$\beta_{\mathrm{PR}}^{[\kappa+1]} = \max\left\{ \frac{(\nabla_L(\mathbf{y}^{[\kappa+1]}))^{\mathrm{T}} \nabla_L(\mathbf{y}^{[\kappa+1]}) - (\nabla_L(\mathbf{y}^{[\kappa+1]}))^{\mathrm{T}} \nabla_L(\mathbf{y}^{[\kappa]})}{(\nabla_L(\mathbf{y}^{[\kappa]}))^{\mathrm{T}} \nabla_L(\mathbf{y}^{[\kappa]})}, 0 \right\}. \tag{29}$$

2.3.2 *Line search for γ*

To obtain the step size $\gamma^{[k]}$, a line search is performed in the direction of the update, $\mathbf{y}^{[\kappa]} + \gamma^{[\kappa]}\mathbf{d}^{[\kappa]}$, to find the $\gamma^{[k]}$ that minimizes $L(\mathbf{y})$ along that direction. The optimal $\gamma^{[k]}$ is such that the following directional derivative is zero:

$$\frac{\mathrm{d}}{\mathrm{d}\gamma^{[k]}}[L(\mathbf{y}^{[\kappa]} + \gamma^{[\kappa]}\mathbf{d}^{[\kappa]})] = (\nabla_L(\mathbf{y}^{[\kappa]} + \gamma^{[\kappa]}\mathbf{d}^{[\kappa]}))^{\mathrm{T}}\mathbf{d}^{[\kappa]} = 0. \tag{30}$$

By expanding the cost function $L(\mathbf{y})$ into a second-order Taylor series[14] and inserting it into Eq. (30), a closed form expression for $\gamma^{[\kappa]}$ is found

$$\gamma^{[\kappa]} = -\frac{(\nabla_L(\mathbf{y}^{[\kappa]}))^{\mathrm{T}}\mathbf{d}^{[\kappa]}}{(\mathbf{d}^{[\kappa]})^{\mathrm{T}}\mathbf{H}(\mathbf{y}^{[\kappa]})\mathbf{d}^{[\kappa]}}, \tag{31}$$

where $\mathbf{H}(\mathbf{y})$ is the Hessian of $L(\mathbf{y})$. Equation (31) results from the approximation that the cost function $L(\mathbf{y})$ is a parabola; for a nonquadratic cost function such as the one of Eq. (26), Eq. (31) is not exact, and can provide an inexact line search. To ensure that the update step provides a descent step,[14,28] the cost function at $\mathbf{y}^{[\kappa]} + \gamma^{[\kappa]}\mathbf{d}^{[\kappa]}$ is calculated; if $L(\mathbf{y}^{[\kappa]} + \gamma^{[\kappa]}\mathbf{d}^{[\kappa]}) < L(\mathbf{y}^{[\kappa]})$, then the update step $\gamma^{[\kappa]}$ is conserved. If the cost function increases, $\gamma^{[\kappa]}$ is divided by two, and the cost function is re-evaluated. This procedure is performed up to $\zeta_{\max}$ times if the cost function does not decrease. After $\zeta_{\max}$ unsuccessful divisions, it is assumed that the search direction is erroneous, and $\gamma^{[\kappa]}$ is reset to the negative of its original value for the current iteration [Eq. (31)]. In our experiments, we found that this procedure is always helpful to reach convergence. However, when reset occurs in late iterations, it may be due to convergence instead of an inaccurate line search; in that case, early stopping may be considered. The algorithm requires closed-form expressions for $\nabla_L(\mathbf{y})$, and $\mathbf{d}^{\mathrm{T}}\mathbf{H}(\mathbf{y})\mathbf{d}$, which are reported in Appendix B.

3 Methodology

3.1 *Simulation Framework*

3.1.1 *Phantom geometry and spectral information*

The framework is implemented in MATLAB R2021a (MathWorks, Natick, Massachusetts, United States), and uses two different geometries. First, a 16-plug cylindrical phantom is generated, where each insert contains a tissue from White and Woodard's database of human tissues.[25] The name and ρ_e of the inserts, the layout of the phantom as well as the ρ_e map are respectively given in Figs. 2(a), 2(c), and 2(e). Phantom size is set to 512×512 pixels to generate projection data [Eq. (5)], while the image is reconstructed on a 256×256 grid. Reconstructed pixel size is always 1×1 mm^2.

The virtual patient is similar to one introduced in previous work,[9] although no contrast agent is used in the current study. The virtual patient is generated from a DECT scan of a pelvis region, with a semiautomatic assignation of tissues and densities. The density is scaled with the HU map, as detailed in Ref. 9. There are 16 materials assigned to the virtual patient, and elemental composition, as well as mass density variability, are introduced following the scheme of Yang et al.[34] Details on the geometry are provided in Figs. 2(b), 2(d), and 2(f). Similar to the cylindrical phantom, the size of the virtual patient is set to 460×784 pixels to produce sinograms and is reconstructed on a 230×392 grid. To generate SPCCT data, a 140 kVp source spectrum from the Siemens SOMATOM Definition Flash provided by the manufacturer (Siemens Healthineers, Forchheim, Germany) is used. Four energy bins ($N = 4$) are generated with a realistic photon counting detector response model,[35] which includes spectral distortion (K-escape, charge sharing). Energy thresholds are set to 60, 72, and 91 keV, which correspond to optimized thresholds to minimize overlap between energy bins, following the global search procedure detailed in precedent work.[11] The four resulting effective spectra $\psi_{n\epsilon}$ are shown in figure 3b of that study.[11]

(a)

Plug #	Tissue name	ρ_e
1	Liver #2	1.050
2	Prostate	1.033
3	Blood whole	1.050
4	Adipose tissue #3	0.933
5	Spleen	1.051
6	Red marrow	1.023
7	Humerus total bone	1.389
8	Skeletal muscle #1	1.039
9	Skin #1	1.078
10	Ribs (second and sixth)	1.347
11	Sacrum (female)	1.330
12	Brain (gray matter)	1.035
13	Mammary gland #1	0.987
14	Kidney #1	1.040
15	Sacrum (male)	1.244
16	Testis	1.034
17	Skeletal muscle #3	1.040
18	Air	0.011

(b)

Plug #	Tissue name	ρ_e
	Water	1.000
	Vertebral column (whole)	1.279
	Small intestine wall	1.024
	Skin #2	1.078
	Ribs (10^{th})	1.441
	Red marrow	1.023
	Pancreas	1.034
	Skeletal muscle #2	1.040
	Liver #2	1.050
	Kidney #2	1.041
	Kidney #1	1.040
	Cartilage	1.083
	Blood whole #1	1.050
	Adipose tissue # 3	0.933
	Adipose tissue # 1	0.966
	Air	0.011

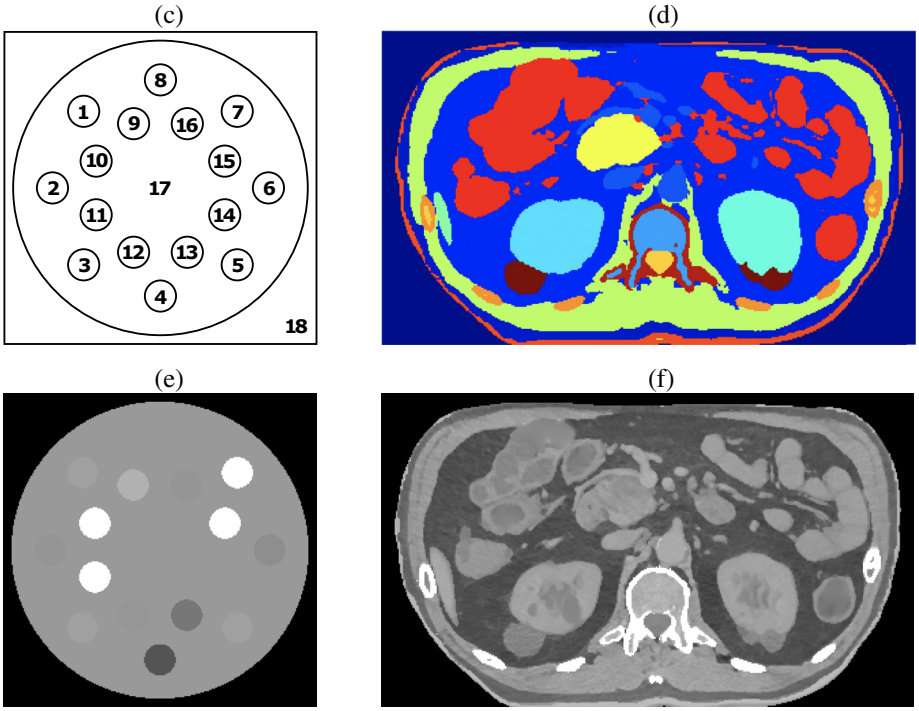

Fig. 2 Two geometries used in this work. (a) and (b) respectively list the inserts used in the cylindrical phantom and the virtual patient, along with their electron density. Figures (c) and (d) illustrate the layouts of the plugs/tissues within the geometries, while (e) and (f) illustrate ρ_e maps (dynamic range between 0.8 and 1.2).

3.1.2 *Sinogram generation*

Sinogram data [Eq. (5)] is generated by considering a weighted sum of monoenergetic projections from 1 to 140 keV, at $\delta\epsilon = 1$ keV intervals,[9,24] and is obtained individually for each spectrum $\psi_{n\epsilon}$. The strip-integral model of the MIRT toolbox[36] is used to generate the projection matrix $\mathbf{A}$. Elemental photon cross sections used to generate $\mu_{j\epsilon}$ are obtained from the XCOM database.[37] Sinogram data are generated using $N_p = 720$ projections and $N_d = 729$ (cylindrical phantom) or 911 (virtual patient) detectors, each of 1 mm size. Poisson noise is added to sinogram data N_{in} to obtain noisy measurements. Sinograms are generated without scattering. Approximately 10^6 photons per projection angle are distributed between the four energy bins, following the relative fluence of $\psi_{n\epsilon}$ in each bin. This results in a noise level of $\approx$100 HU per energy bin, which roughly matches experimental SPCCT noise levels reported in the literature, between 80 and 112 HU.[4,6]

3.2 *Practical Implementation of the Reconstruction Algorithm*

A total of $P = 70$ tissues from White and Woodard's database[25] (all 25 bone and 45 soft tissues including all but thyroid) are used to (1) generate eigentissues[7] and (2) create the KDE in the quantitative prior [Eq. (14)]. For regularization parameters, δ is fixed to 0.02 based on the previous work of Cai et al.[14] with the Huber prior and our own experiments. For the spatial prior, δ has a limited impact on the quantitative accuracy of the method compared to β, hence its fixed value. The parameters α and β are optimized on the cylindrical phantom, as described in Sec. 3.3. For the optimization algorithm, following the nomenclature of the pseudocode in Appendix A, the maximum number of iterations $\kappa_{\max}$ is 2500, while $\zeta_{\max}$ for the line search is set to 10, similar to Mory et al.[28] In all experiments, setting $\kappa_{\max} = 2500$ resulted in a cost function varying by $<10^{-5}\%$ between iterations, ensuring that convergence is reached.

Finally, the optimization algorithm requires an initial guess $\mathbf{y}^{[\kappa=0]}$. For this study, to focus solely on the impact of α and β, the initial guess is obtained from ground truth ρ_e data instead of estimated values. Due to the nonconvexity of the quantitative prior, all-zero initialization is not recommended. Each pixel in the image to reconstruct is first assigned to one of five classes following its ρ_e. Pixels with $\rho_e < 0.1$ are identified as air, $0.1 \leq \rho_e < 0.8$ are lung, $0.8 \leq \rho_e < 1$ fat tissues, $1 \leq \rho_e < 1.11$ soft tissues, and $\rho_e \geq 1.11$ bone. For each of the five classes c, an initial set of eigentissue partial electron densities $\mathbf{y}_c = (y_{c1} \ldots y_{cK})^{\mathrm{T}}$ is assigned, resulting in an initial image guess $\mathbf{y}^{[\kappa=0]}$ that can take, for one eigentissue channel k, at most five different values. The initial $\mathbf{y}_c$ of lung is taken, using the nomenclature of the quantitative prior, as $\mathbf{y}_{\mathrm{ref},p_{\mathrm{lung}}}$, with p_{lung} corresponding to the index, in the reference database, of the deflated lung. Similar, for fat-based tissues, the initial $\mathbf{y}_c$ is taken as the average of the $\mathbf{y}_{\mathrm{ref},p}$ corresponding to low-density fat-like tissues in the reference database (i.e., adipose tissue 3, adipose tissue 2, adipose tissue 1, yellow marrow, mammary gland 1). For the soft tissue and bone classes, a similar procedure is performed by averaging the $\mathbf{y}_{\mathrm{ref},p}$ of all the remaining soft and bone tissues in the reference database, respectively. For air, as it is not present in the database that generates eigentissues, there is no value of $\mathbf{y}_{\mathrm{ref},p}$ corresponding to air. To generate one, a global search is performed over a range of $\mathbf{y}_c$ to find a $\mathbf{y}_{\mathrm{air}}$ that produces minimal x-ray attenuation, while matching the density of air. Practically, considering Eq. (4), a $\mathbf{y}_{\mathrm{air}}$ is found such that $\sum_{k=1}^{K} y_{\mathrm{air},k} g_{k\epsilon}^{\mathrm{ET}} - y_0 g_{0\epsilon}^{\mathrm{ET}} = 0$ under the constraint that $\mathbf{J}_{1\times K}\,\mathbf{y}_{\mathrm{air}} = \rho_{e,\mathrm{air}}$, where $\mathbf{J}_{1\times K}$ denotes the $1 \times K$ array of ones. Finally, it should be noted that using the post-reconstruction solution as an initial guess produces similar results to those obtained with the above initialization. It is recommended to use an initial guess reasonably close to the solution to maximize performance.

3.3 *Comparative Study*

3.3.1 *Cylindrical phantom*

The cylindrical phantom is reconstructed for a range of α and β. A qualitative comparison showing parametric maps as well as relative error maps on relevant physical properties as a function of selected values of α and β is first shown in Figs. 3 and 4. Then, various quantitative metrics are reported in Fig. 5 to further compare reconstructions and help selecting optimal regularization parameters. The quantitative metrics are defined below.

Metrics are defined for a physical parameter θ, where $\theta \in \{\mathrm{SPR}, \rho_e\}$. The SPR is obtained with the Bragg additivity rule using the elemental fractions obtained from eigentissue fractions for the I-values. First, let $\hat{\theta}_j$ represent the reconstructed parameter at pixel j. The relative error (in %) on θ for pixel j is $\delta_j^{(\theta)} = 100 \times (\hat{\theta}_j - \theta_j)\theta_j^{-1}$, where θ_j is the ground truth value. For parameter θ and insert q in the cylindrical phantom, the mean relative error (bias) is obtained by averaging over pixels of the insert: $B_{\theta,q} = \mathcal{N}_q^{-1} \sum_{j\in V_q} \delta_j^{(\theta)}$. Similar, $\mathrm{RMSE}_{\theta,q} = \sqrt{\mathcal{N}_q^{-1} \sum_{j\in V_q} (\delta_j^{(\theta)})^2}$. Here, V_q is a region of interest associated with insert q (the inner 90% of all pixels in the insert) and $\mathcal{N}_q$ is the number of pixels in V_q. The standard deviation of θ in insert q is $\sigma_{\theta,q} = \sqrt{\mathrm{RMSE}_{\theta,q}^2 - B_{\theta,q}^2}$.

To quantify the overall accuracy and precision of parameter θ, the mean absolute bias, RMSE and standard deviation averaged over all Q plugs are reported; e.g., the mean absolute bias is $|B_\theta| = \frac{1}{Q}\sum_{q=1}^{Q}|B_{\theta,q}|$. The other metrics are labeled RMSE_θ and σ_θ.

While the plug-averaged metrics represent the overall quantitative performance, they may not adequately factor in any form of spatial distortion or visual artifacts that can be introduced in the reconstruction process, especially with strong regularization. Another comparison metric that takes into account the overall spatial distribution of errors is thus proposed. The metric is based on the water equivalent path length (WEPL) used in proton therapy. The WEPL is defined, along a given path ℓ, as $\text{WEPL}_\ell = \sum_{j\in\ell}\text{SPR}_j$, where j denotes a pixel on the path ℓ. The relative error (in %) on WEPL_ℓ is $\delta_\ell^{\text{WEPL}} = 100\times(\widehat{\text{WEPL}}_\ell - \text{WEPL}_\ell)\text{WEPL}_\ell^{-1}$. The reported metric is the RMSE on $\delta_\ell^{\text{WEPL}}$ calculated over L various paths ℓ. Similar to Vilches-Freixas et al.,[18] WEPL_ℓ are generated by taking the Radon transform on ground truth and estimated SPR maps for angles distributed from 0 to 179 deg, with increments of 1 deg. Paths with $\text{WEPL}_\ell = 0$ are removed from the analysis, and the SPR of all pixels outside the phantom geometry are set to 0. Practically, this results in $L = 43353$ WEPL values for the cylindrical phantom, which essentially covers all possible beam paths traversing the phantom.

The optimal set of regularization parameters may depend on the subsequent task to perform (e.g., quantitative imaging for diagnostic purposes or treatment planning). Three sets of optimized regularization parameters are extracted by looking at the combinations shown in Figs. 3 and 4: the ones minimizing (1) the absolute bias $|B_{\text{SPR}}|$, (2) the root-mean-square error RMSE_{SPR}, and (3) the RMSE on the WEPL, calculated over the L paths defined previously.

3.3.2 *Virtual patient*

The virtual phantom is reconstructed with the three sets of optimized regularization parameters, as suggested in Sec. 3.3.1. The three resulting images obtained with the one-step framework are compared with three methods. The first, referred to as two-step (FBP), is the sinogram-based two-step ETD decomposition on all tissues,[11] with $K = 2$ eigentissues. The second, named two-step (iterative), uses the same first step as the previous method, but reconstructs the y_{jk} using an adaptation of the iterative framework of Schirra et al.,[38] a penalized weighted least squares algorithm based on separable paraboloidal surrogates. All-zero initialization was used, and regularization parameters β_k and δ_k (one per eigentissue channel) were optimized to minimize SPR bias, as presented in Sec. 3.3.1. Optimized parameters were $\delta_k = 0.02\sigma_k$ and $\beta_k = 0.0005/\sigma_k^2$. For completeness, the update step is

$$y_{jk}^{[\kappa+1]} = y_{jk}^{[\kappa]} + \frac{\sum_{i=1}^{I} a_{ij}\dot{h}_{ik}(Y_{ik}) - \beta_k'\phi'([\mathbf{D}\mathbf{y}_\mathbf{k}]_j,\delta_k)}{-\sum_{i=1}^{I}\ddot{h}_{ik}a_{ij}a_i + \beta_k'\phi'([\mathbf{D}\mathbf{y}_\mathbf{k}]_j,\delta_k)}, \tag{32}$$

where $Y_{ik} = \sum_{j=1}^{J} a_{ij}y_{jk}$ is the forward projection of eigentissue fractions y_{jk}. $\dot{h}_{ik}(Y_{ik}) = \frac{Y_{ik}^0 - Y_{ik}}{\sigma_{ik}^2}$ and $\ddot{h}_{jk} = \frac{-1}{\sigma_{ik}^2}$ are respectively the first and second derivatives of the marginal log-likelihood function, obtained after computing the diagonal elements of the inverse of the Fisher information matrix, σ_{ik}^2, once the first step of the two-step method is complete and returns eigentissue sinograms Y_{ik}^0; see Ref. 38 for details. Also, $a_i \equiv \sum_{j=1}^{J} a_{ij}$, and $\beta_k' = -\beta_k\sum_{i=1}^{I} a_{ij}\ddot{h}_{ik}a_i$.

The third comparison method is the post-reconstruction ETD with quantitative regularization ($\alpha = 0.4$).[8] For the post-reconstruction approach, the beam hardening correction algorithm of Kijewski and Bjärngard[39] is used, and the Z-space calibration is performed on a simulated Gammex phantom, as detailed in Lalonde et al.[5] A similar qualitative and quantitative analysis to the cylindrical phantom is performed and reported in Figs. 6 and 7.

4 Results

4.1 *Impact of α and β*

Figure 3 illustrates the resulting image quality for ρ_e using a selected set of regularization parameters α and β. Note that quantitative regularization increases as α gets smaller; the case $\alpha = \infty$ corresponds to no quantitative regularization. Spatial regularization increases as β gets larger, with $\beta = 0$ corresponding to no spatial regularization. To complement Fig. 3, corresponding absolute error maps on ρ_e are also presented in Fig. 4.

Figures 3 and 4 illustrate that α and β must be carefully tuned to maximize precision and accuracy, i.e., minimize the uncertainties and biases. To further investigate the quantitative performance with respect to α and β and select optimal regularization parameters, Fig. 5 shows the performance for various metrics related to the SPR, as this physical parameter naturally combines ρ_e and the elemental composition.

Figure 5 generally shows that a set of parameters (α, β) that produce low bias lead to larger standard deviations. The parameters that minimize $|B_{\mathrm{SPR}}|$ are $(\hat{\alpha}_{\mathrm{B}}, \hat{\beta}_{\mathrm{B}}) = (15000, 0.0001)$, and result in a performance of $|B_{\mathrm{SPR}}| = 0.19\%$ and $\mathrm{RMSE}_{\mathrm{SPR}} = 1.32\%$. The regularization values that minimize $\mathrm{RMSE}_{\mathrm{SPR}}$ are $(\hat{\alpha}_{\mathrm{R}}, \hat{\beta}_{\mathrm{R}}) = (2000, 0.001)$, which provide a performance of $|B_{\mathrm{SPR}}| = 0.40\%$ and $\mathrm{RMSE}_{\mathrm{SPR}} = 0.48\%$. Finally, the α and β that minimize the RMSE on the WEPL are $(\hat{\alpha}_{\mathrm{W}}, \hat{\beta}_{\mathrm{W}}) = (5000, 0.0005)$, and result in $|B_{\mathrm{SPR}}| = 0.35\%$ and $\mathrm{RMSE}_{\mathrm{SPR}} = 0.55\%$. The set of parameters $(\hat{\alpha}_{\mathrm{W}}, \hat{\beta}_{\mathrm{W}})$ thus corresponds to a tradeoff between the minimal bias and minimal RMSE cases.

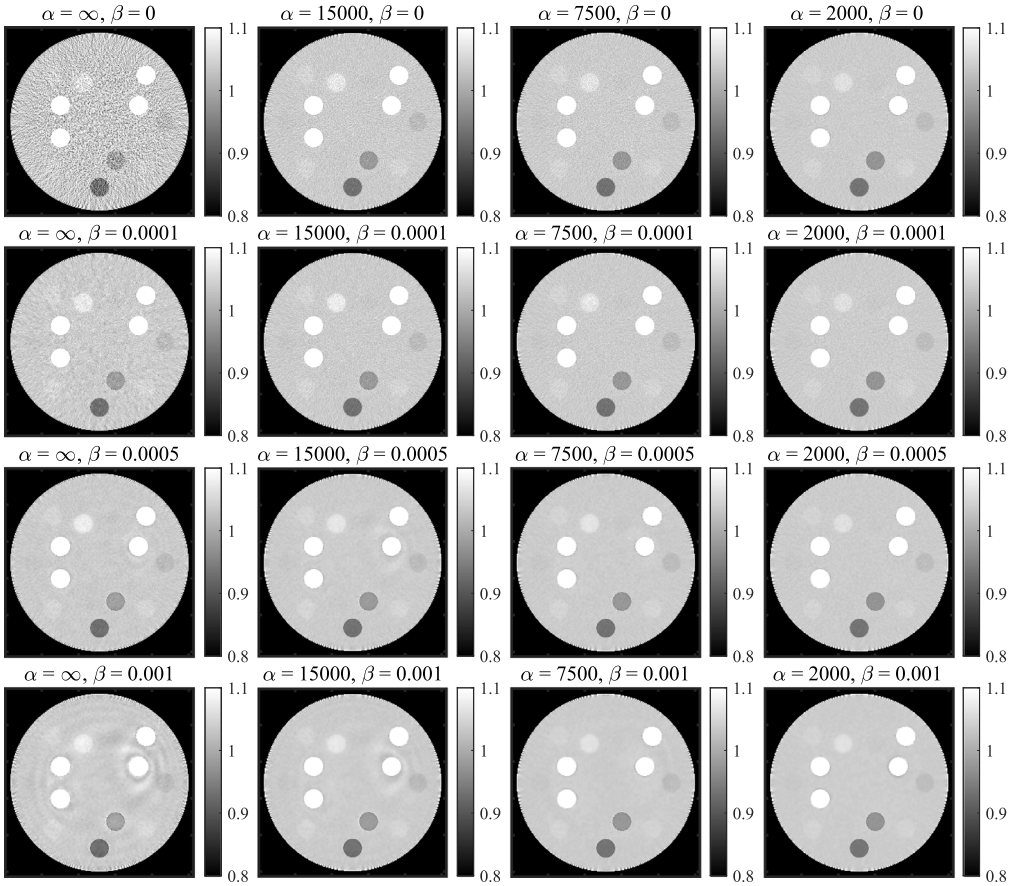

Fig. 3 Electron density maps for various combinations of regularization parameters α and β. Spatial variation increases from top to bottom, while quantitative regularization increases from left to right.

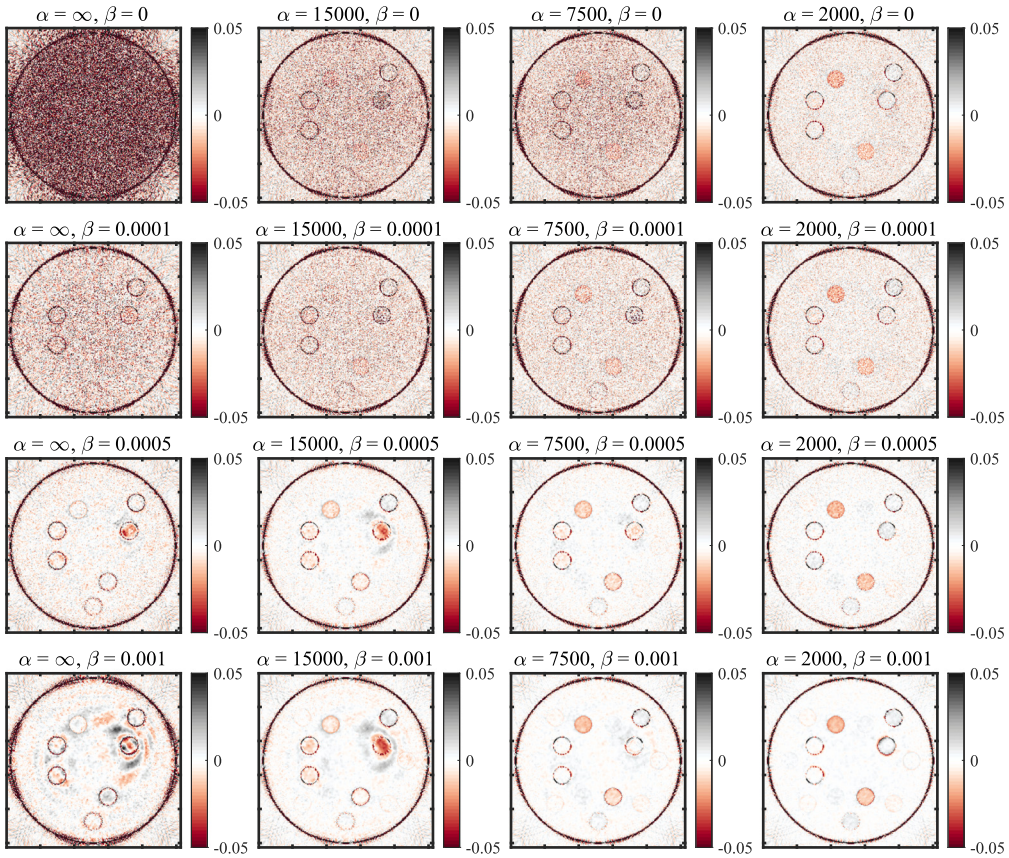

Fig. 4 Maps of absolute error on the electron density using the same set of regularization parameters as presented in Fig. 3.

4.2 *Performance on the Virtual Patient*

The one-step reconstruction algorithm's performance is compared with the two-step frameworks (using FBP[11] and an adapted iterative scheme[38] for the second step) and the regularized post-reconstruction[8] frameworks. All approaches use eigentissues as a material basis. In the following figures and the discussion, the one-step method which uses optimized regularization parameters $(\hat{\alpha}_{\mathrm{B}}, \hat{\beta}_{\mathrm{B}})$, $(\hat{\alpha}_{\mathrm{R}}, \hat{\beta}_{\mathrm{R}})$, and $(\hat{\alpha}_{\mathrm{W}}, \hat{\beta}_{\mathrm{W}})$ are respectively referred to as the minimal bias, minimal RMSE, and WEPL-optimized cases. Figure 6 illustrates the relative visual performance of the different ETD frameworks to estimate the SPR.

The quantitative performance of all implementations of ETD for the virtual patient is shown in Fig. 7, where the bias and RMSE on the SPR for 6 selected tissues in the virtual patient are reported for all methods. In addition, Table 1 summarizes the results of Fig. 7 by reporting the average values of both metrics over all tissues.

5 Discussion

The proposed one-step framework combines spatial and quantitative priors to provide superior performance compared to two-step and post-reconstruction methods with ETD. Figures 3–5 illustrate that both priors have a significant impact on the performance of the algorithm; using only one prior leads to suboptimal performance. While the quantitative prior generally increases precision, it also limits unnatural features in the images,[27] which often appear with strong spatial regularization. The case with $(\alpha, \beta) = (\infty, 0.001)$ leads to visible artifacts in the ρ_e map (Fig. 4), especially around the sacrum (male) insert. Incorporating quantitative regularization, for instance with $(\alpha, \beta) = (7500, 0.001)$, almost entirely eliminates non-physical features.

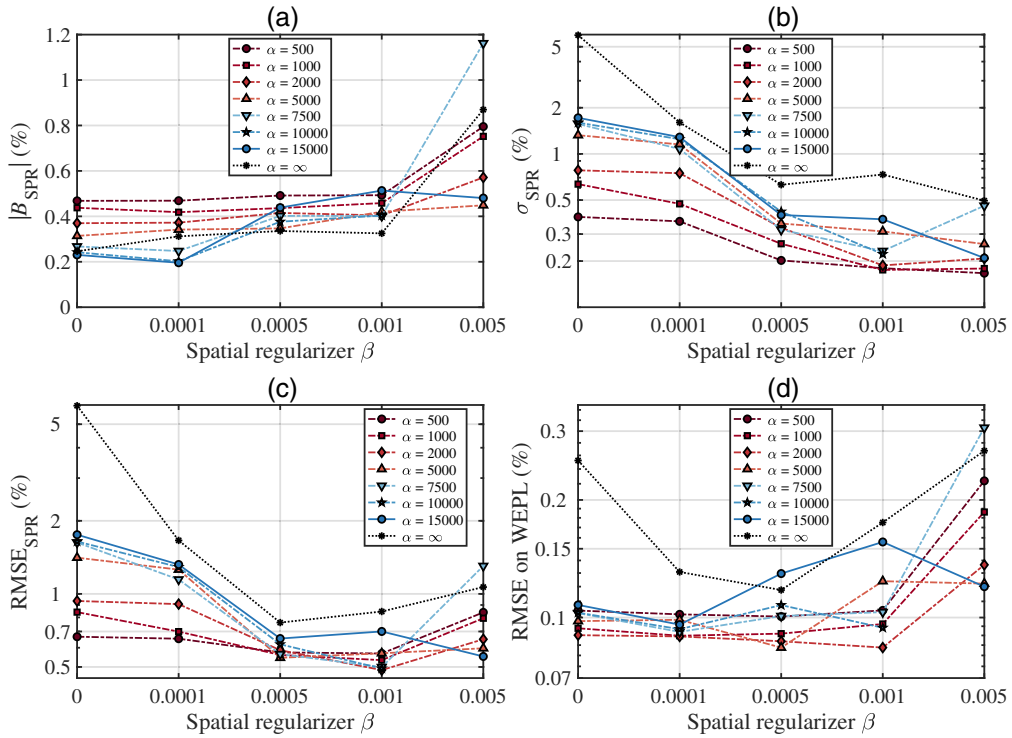

Fig. 5 Quantitative performance of the one-step iterative reconstruction algorithm to estimate the SPR on the cylindrical phantom for a range of regularization parameters α and β. (a) Through (c) respectively show the plug-averaged absolute bias, standard deviation and RMSE on the SPR in %, while (d) shows the RMSE on the WEPL calculated from SPR maps and multiple paths. Note that $\text{RMSE}_{\text{SPR}} = \sqrt{B^2_{\text{SPR}} + \sigma^2_{\text{SPR}}}$.

This study suggests that one-step reconstruction frameworks can benefit from quantitative regularization, in addition to the commonly used spatial regularization.

A visual assessment of Fig. 4 illustrates how the minimal bias dataset, with $(\hat{\alpha}_{\text{B}}, \hat{\beta}_{\text{B}}) = (15000, 0.0001)$, results in important denoising without significantly introducing systematic errors in plugs. While the minimal RMSE dataset $[(\hat{\alpha}_{\text{R}}, \hat{\beta}_{\text{R}}) = (2000, 0.001)]$ leads to stronger denoising, it produces larger biases in inserts, and results in poor performance around the sacrum (male) insert. Although the average RMSE per plug can be an adequate metric, it is apparent that it does not take into account potential image artifacts similar to those around the sacrum (male) insert, hence the usefulness of another metric that considers the entire spatial distribution of errors, such as the RMSE on the WEPL, to obtain optimal α and β. Such metric is particularly relevant, for instance, in proton therapy beam range calculation.

Figures 6 and 7 illustrate the performance of the virtual patient. For precision, the one-step framework is less noisy than its post-reconstruction and two-step counterparts. This is confirmed visually in Fig. 6, as well as quantitatively in Fig. 7. Indeed, as shown in Table 1, the tissue-averaged RMSE on the SPR is largely reduced with the use of the one-step framework with any set of optimized regularization parameters. For the two-step FBP framework, the large RMSE of 13.27% is due to the use of the FBP and no regularization, which leads to strong amplification of sinogram noise, which is equivalent to ≈100 HU in the image domain. The improved precision compared to the two-step (iterative) framework is likely due to the inclusion of the quantitative prior, and the removal of the first step (material decomposition in sinogram space), which is not regularized. For post-reconstruction, the relatively large RMSE of 3.24% is explained by the use of a quantitative prior only; there is no form of spatial regularization in the post-reconstruction implementation of ETD. As noted in Figs. 3–5, setting $\beta = 0$ and adding quantitative regularization α can only provide a limited reduction in RMSE compared to using both regularizers.

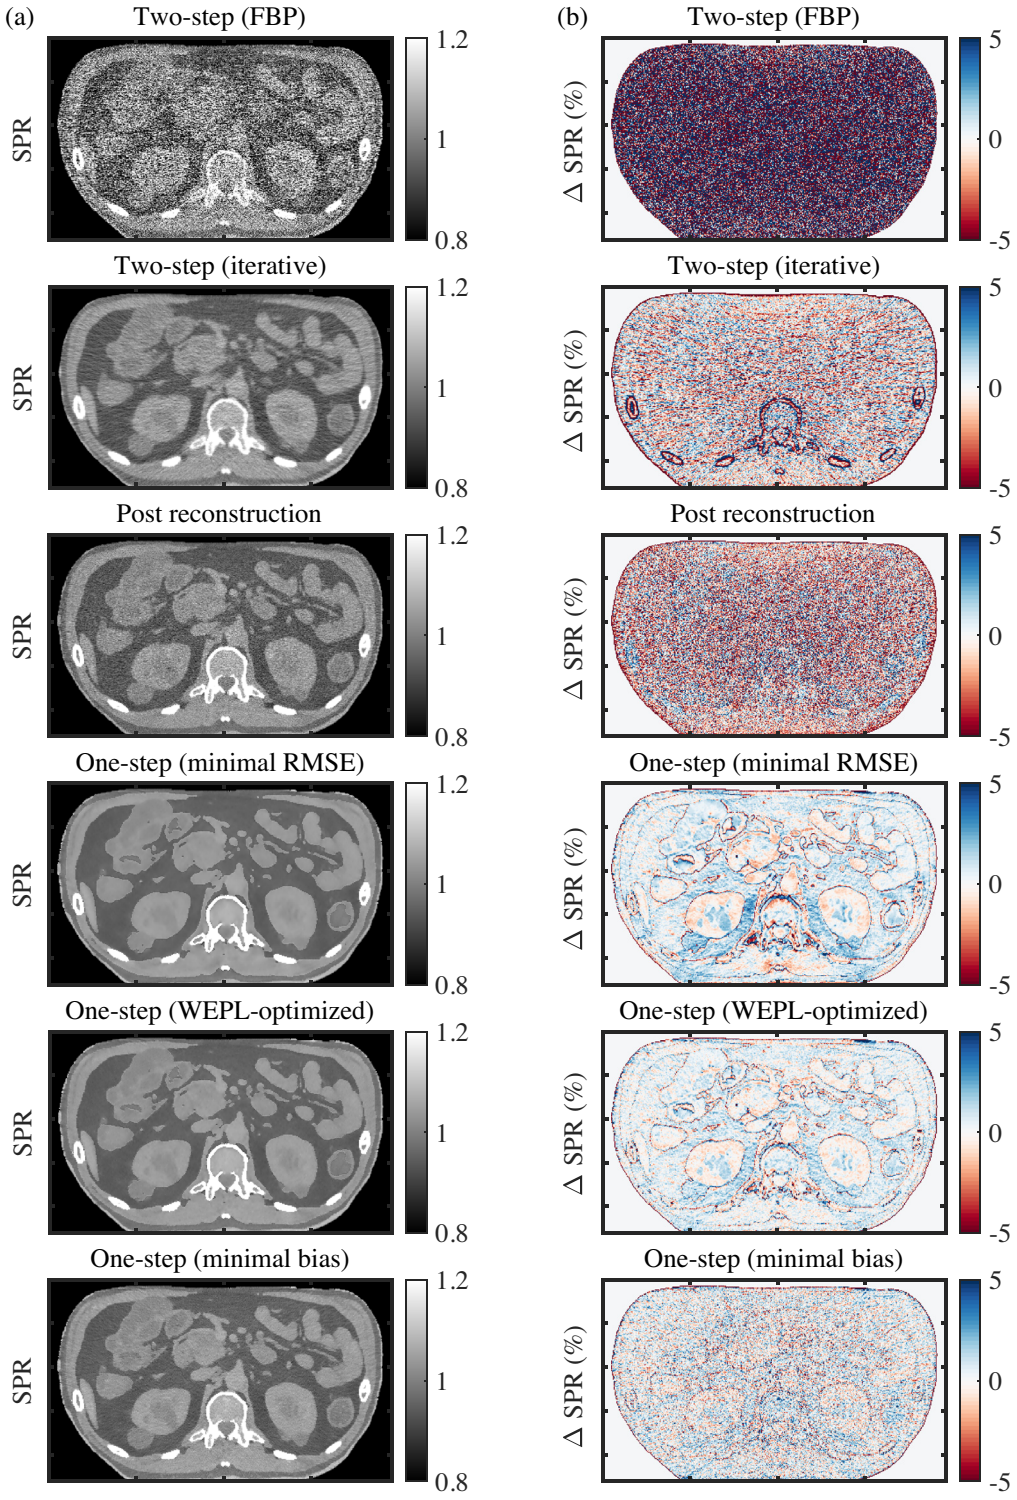

Fig. 6 Performance of various implementations of ETD to estimate the SPR on the virtual patient. (a) Shows SPR maps, while (b) shows relative SPR error maps.

Furthermore, Figs. 6 and 7 and Table 1 also show that the one-step method outperforms other ETD frameworks to minimize the bias on the SPR. First, the tissue-averaged absolute bias (Table 1) is reduced from a range of 0.56% to 0.63% with the two-step framework (FBP or iterative) to a range of 0.24% to 0.43% with the one-step framework using different

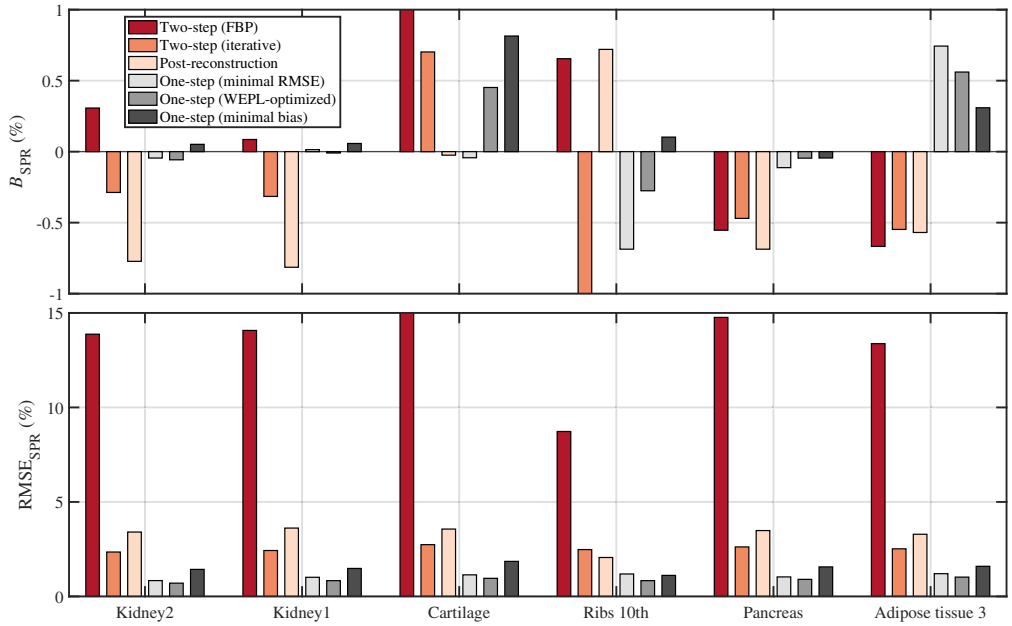

Fig. 7 Quantitative performance of various implementations of ETD on selected tissues of the virtual patient. Top illustrates the RMSE on the SPR, while the bottom plot shows the bias.

Table 1 Absolute bias and RMSE on the SPR for all ETD frameworks tested on the virtual patient. Reported values correspond to the average over the 6 tissues shown in Fig. 7.

Framework	Absolute bias (%)	RMSE (%)
Two-step (FBP)	0.56	13.27
Two-step (iterative)	0.63	2.53
Post-reconstruction	0.60	3.24
One-step (minimal bias)	0.24	1.49
One-step (WEPL-optimized)	0.31	0.88
One-step (minimal RMSE)	0.43	1.09

regularization parameters. This reduction is mostly attributed to the use of four eigentissues (one-step) instead of two eigentissues (two-step) to parametrize the attenuation coefficient, which reduces the truncation error that appears when estimating the elemental composition of the eigentissues.[7] Comparing the one-step framework with post-reconstruction, it might be expected that post-reconstruction would be less biased than the one-step framework, as different eigentissues are derived for soft tissues and bone in the post-reconstruction framework, which leads to a theoretically more accurate parametrization for individual tissues. We attribute the worse performance of post-reconstruction ETD (average absolute bias of 0.60%) to beam hardening artifacts.

Qualitative and quantitative results also allow us to formulate recommendations for the selection of regularization parameters α and β for the one-step framework. First, the minimal RMSE case [with $(\hat{\alpha}_R, \hat{\beta}_R) = (2000, 0.001)$] produces highly corrupted SPR images of the virtual patient (Fig. 6), especially along soft tissue and bone interfaces between the kidneys. For this set of regularization parameters, the amplitude of such systematic errors is not as important on the cylindrical phantom. This might be explained by the much more complex geometry to reconstruct for the virtual patient. Furthermore, based on Table 1, the performance of the one-step

method (minimal RMSE case) in terms of RMSE on the SPR is worse than the WEPL-optimized case for the virtual patient. This shows that the optimization of α and β performed on a specific geometry (cylindrical phantom) may not produce optimal parameters on a patient geometry. We conclude that the regularization parameters optimized such that the RMSE in selected regions of interest is minimal results in an unacceptable level of regularization, and poor image quality.

Fortunately, considering the SPR maps of Fig. 6, the minimal bias case, with $(\hat{\alpha}_{\mathrm{B}}, \hat{\beta}_{\mathrm{B}}) = (15000, 0.0001)$, does not produce any significant systematic errors comparable with the minimal RMSE and WEPL-optimized cases. This case produces the minimal bias on selected tissues of the virtual patient, results in natural-looking SPR maps and still largely reduces the RMSE on the SPR compared with post-reconstruction, from 3.24% to 1.49%. We thus recommend to use the current one-step framework with regularization parameters optimized to minimize the bias.

It should be noted that there have been previous implementations of quantitative priors for iterative reconstruction in CT. Regularization has been proposed directly on attenuation data,[40] or implemented through a conservation of volumes constraints,[15] which are not directly compatible with ETD. More recently, Nakada et al.[41] have introduced a quantitative prior for tissue type assignation for spectral CT. While their cost function also used a Gaussian kernel centered around reference tissues, it differs from the proposed prior as it is based on a Compton scattering and photoelectric effect parametrization, which is not optimal to reproduce radiotherapy-related quantities such as the proton stopping power.[24] Furthermore, their quantitative prior uses three regularization parameters, a covariance matrix sampled from the scanned object, and a restricted set of possible tissues. This may require additional fine-tuning compared to the single parameter (α) used in this work for the quantitative prior.

This study has some limitations. First, the conclusions and optimized regularization parameters are valid for this simulated imaging context, noise level, detector response, and selected energy bins; results may vary otherwise. As noted previously, the optimization procedure for α and β was also performed on a specific cylindrical phantom and used on the virtual patient; due to differences in materials and geometry, the results obtained for the virtual patient may not be optimal. Second, 2D reconstruction is used in the current study, but further improvements could be obtained with the use of 3D regularization. Third, scattering was not simulated nor included in the model proposed in this work. It is expected that the overall performance of all methods investigated in this work may be reduced with the inclusion of scattering. In addition, the quantitative prior heavily relies on the work of White and Woodard[25] to generate reference human tissue centroids in the eigentissue space. The method is validated using the same reference tissues, and may not be ideal for realistic data, if real tissue composition differs from the reference database. Future studies on tissue stoichiometry for patients could be useful to improve the prior. Finally, this study does not explicitly compare the performance of ETD against other parametrizations of the attenuation coefficient. However, previous studies have demonstrated that ETD provides more accurate tissue characterization compared to material basis decomposition for post-reconstruction[7,8] and two-step sinogram-based[11] frameworks. It is therefore expected that the same benefits would translate to the one-step framework.

Furthermore, additional steps are required before a full experimental evaluation of this method can be carried out. For instance, ψ_{ne}, which includes the source spectrum and the detector response, must be evaluated through a careful calibration procedure. Scattering should also be incorporated into the model. This study illustrates the potential of the proposed method, and experimental performance may be reduced. Also, the initial guess is derived using ground truth values of ρ_e; to maximize the quality of the initial guess with experimental data, SPCCT bins could be merged to produce a low-noise SECT dataset, from which ρ_e can be inferred, similar to the work of Yu et al.[42] Future work will be focused on an experimental implementation of the method presented herein with a SPCCT scanner, and an adaptation to contrast-enhanced imaging with ETD,[9] which has been left out due to the complexity of incorporating the conservation of volumes constraint in the model.

6 Conclusion

This study reports a one-step reconstruction algorithm to extract highly precise and accurate radiotherapy-related parameters for human tissues with SPCCT, although it is compatible with

DECT. The method uses eigentissues, a PCA-derived decomposition basis tailored for human tissues, as well as two separate prior terms, which provide optimal performance when combined. A quantitative prior, controlled by regularization parameter α, encourages eigentissue fractions to lie in proximity of expected values for human tissues. The spatial prior, with parameter β, encourages smooth intensity variations between eigentissue fractions of neighboring pixels. With proper tuning of α and β, we show that the one-step framework provides increased accuracy and precision to estimate ρ_e and SPR compared to a state-of-the-art post-reconstruction implementation of ETD,[8] as well as a finely tuned two-step framework with iterative reconstruction, although at the cost of increased computational complexity. Based on the results for the virtual patient, we suggest that optimal parameters α and β should be selected to minimize biases in selected regions of interest. The one-step framework is a key component in helping SPCCT reach its full potential for tissue characterization and quantitative imaging.

7 Appendix A: Pseudocode for the Optimization Algorithm

The set of variables to optimize is $\mathbf{y}$, the starting point $\mathbf{y}^{[\kappa=0]}$, and the gradient and Hessian are $\nabla_L(\mathbf{y})$ and $\mathbf{H}(\mathbf{y})$. CG iterations are noted with κ, and the maximum number of CG iterations is $\kappa_{\max}$. For the line search, the index ζ is used to count the number of times that the update step is halved, up to a maximum of $\zeta_{\max}$ times.

Algorithm 1 Nonlinear conjugate gradient for one-step ETD.

$\kappa \leftarrow 0$
$\mathbf{y} \leftarrow \mathbf{y}^{[\kappa=0]}$
$\mathbf{r} \leftarrow -\nabla_L(\mathbf{y})$
$\mathbf{d} \leftarrow \mathbf{r}$
$\delta \leftarrow \mathbf{r}^T\mathbf{d}$
while $\kappa < \kappa_{\max}$ **do**
  $L_0 \leftarrow L(\mathbf{y})$
  $\gamma \leftarrow \frac{\mathbf{r}^T\mathbf{d}}{\mathbf{d}^T\mathbf{H}(\mathbf{y})\mathbf{d}}$
  $\gamma_0 \leftarrow \gamma$
  $\zeta \leftarrow 0$
  while $\zeta < \zeta_{\max} + 1$ and $L_0 > L(\mathbf{y} + \gamma\mathbf{d})$ **do**
    if $\zeta < \zeta_{\max}$
      $\gamma \leftarrow \gamma/2$
    else if $\zeta = \zeta_{\max}$
      $\gamma \leftarrow -\gamma_0$
    $\zeta \leftarrow \zeta + 1$
  $\mathbf{y} \leftarrow \mathbf{y} + \gamma\mathbf{d}$
  $\mathbf{r} \leftarrow -\nabla_L(\mathbf{y})$
  $\delta_p \leftarrow \delta$
  $\delta_m \leftarrow \mathbf{r}^T\mathbf{d}$
  $\delta \leftarrow \mathbf{r}^T\mathbf{r}$
  $\beta_{\text{PR}} = \max\left(\frac{\delta - \delta_m}{\delta_p}, 0\right)$
  $\mathbf{d} \leftarrow \mathbf{r} + \beta_{\text{PR}}\mathbf{d}$
  $\kappa \leftarrow \kappa + 1$

8 Appendix B: Closed form Expressions for $\nabla_L(\mathbf{y})$ and $\mathbf{d}^{\mathsf{T}}\mathbf{H}(\mathbf{y})\mathbf{d}$

For ease of reimplementation, final closed form expressions for derivatives of the cost function are provided. For reference, some intermediate expressions used are listed in Table 2.

8.1 *Gradient of the Cost Function*

The gradient is divided into likelihood, quantitative prior and spatial prior contributions, per Eq. (26)

$$\nabla_L(\mathbf{y}) \equiv \frac{\partial L(\mathbf{y})}{\partial y_{jk}} = \frac{\partial \ell_L(\mathbf{y})}{\partial y_{jk}} + \frac{\partial \ell_{P_Q}(\mathbf{y})}{\partial y_{jk}} + \frac{\partial \ell_{P_S}(\mathbf{y})}{\partial y_{jk}}. \tag{33}$$

The gradient of the negative log likelihood can be conveniently expressed as K backprojection operations

$$\frac{\partial \ell_L(\mathbf{y})}{\partial y_{jk}} = \sum_{i=1}^{I} \mathcal{B}_{ik} a_{ij}, \tag{34}$$

where

$$\mathcal{B}_{ik} = \sum_{\epsilon=1}^{N_e} g_{k\epsilon} F_{i\epsilon} \left(\sum_{n=1}^{N} \psi_{n\epsilon} \left(\frac{t_{in}}{\hat{t}_{in}} - 1 \right) \right). \tag{35}$$

The gradient of the negative log prior (quantitative) is

$$\frac{\partial \ell_{P_Q}(\mathbf{y})}{\partial y_{jk}} = \frac{1}{\sigma_{y_k}^2 \mathcal{P}_j} \sum_{p=1}^{P} \left(2\nu K_{jp} \left(\mathcal{D}_{jkp} - w_{jp} \left(\sum_{p'=1}^{P} \mathcal{D}_{jkp'} \right) \right) + \frac{1}{\alpha} w_{jp} K_{jp} \Delta_{jkp} \right), \tag{36}$$

while the spatial contribution to the negative log prior is given as

$$\frac{\partial \ell_{P,S}(\mathbf{y})}{\partial \mathbf{y_k}} = \beta \sigma_{y_k}^2 \sum_{r=1}^{J} D_{jr}^{\mathrm{T}} \phi' \left(\sum_{c=1}^{J} D_{rc} y_{ck}, \delta_k \right) = \beta \sigma_{y_k}^2 \mathbf{D}^{\mathrm{T}} \phi'(\mathbf{D y_k}), \tag{37}$$

where $\phi'(\cdot)$ denotes the derivative of the Huber potential function.

8.2 *Denominator of the Step Size*

Similar to Eq. (33), the denominator of the step size from Eq. (31) is divided into contributions from the likelihood, quantitative and spatial priors

Table 2 Intermediate expressions for the calculation of the gradient and Hessian.

Parameter	Formula
Attenuation factor	$\mathcal{A}_{i\epsilon} = g_{0\epsilon}^{\mathrm{ET}} Y_{i0} + \sum_{k=1}^{K} g_{k\epsilon}^{\mathrm{ET}} \sum_{j=1}^{J} a_{ij} y_{jk}$
Soft exponential factor	$F_{i\epsilon} = \begin{cases} \exp(-\mathcal{A}_{i\epsilon}) & \text{if } \mathcal{A}_{i\epsilon} \geq 0; \\ 1 & \text{if } \mathcal{A}_{i\epsilon} < 0. \end{cases}$
Gaussian kernel	$K_{jp} = \exp\left(-\frac{d_{jp}^{(y)}}{2\alpha} \right)$
Non-normalized probability	$\mathcal{P}_j = \sum_{p=1}^{P} w_{jp} K_{jp}$
Distance in eigentissue space	$\Delta_{jkp} = y_{jk} - y_{\mathrm{ref},k}^{(p)}$
Intermediate expression	$\mathcal{D}_{jkp} = w_{jp} \Delta_{jkp} \pi_{jp}^{\frac{1}{2}}$
Forward projection of the search direction	$\varphi_{ik} = \sum_{j=1}^{J} a_{ij} d_{jk}$

$$\mathbf{d}^{\mathrm{T}}\mathbf{H}\mathbf{d} = \mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathbf{L}}\mathbf{d} + \mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathbf{P_Q}}\mathbf{d} + \mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathbf{P_S}}\mathbf{d}. \tag{38}$$

Elements of the Hessian for the likelihood are noted as $H^{(L)}_{jkj'k'}$. Defining the quantities

$$\begin{cases} \mathcal{W}_{i\epsilon} = \sum_{k=1}^{K} \varphi_{ik} g_{k\epsilon}; \\ \mathcal{X}_{i\epsilon} = \sum_{k'=1}^{K} \varphi_{ik'} \left[\left(\sum_{n=1}^{N} \psi_{n\epsilon} \left(1 - \frac{t_{in}}{\bar{t}_{in}} \right) \right) g_{k'\epsilon} \theta(\mathcal{A}_{i\epsilon}) + \sum_{n=1}^{N} \psi_{n\epsilon} t_{in} \frac{\sum_{\epsilon'=1}^{N_e} g_{k'\epsilon} \psi_{n\epsilon'} \exp(-\mathcal{A}_{i\epsilon'})}{\bar{t}_{in}^2} \right], \end{cases} \tag{39}$$

where $\theta(\mathcal{A}_{i\epsilon})$ denotes the Heaviside step function, a compact expression for the likelihood term is

$$\mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathrm{L}}\mathbf{d} = \sum_{i=1}^{I} \sum_{\epsilon=1}^{N_e} \mathcal{W}_{i\epsilon} \mathcal{X}_{i\epsilon} F_{i\epsilon}. \tag{40}$$

For the quantitative log prior, the Hessian matrix is largely sparse as it operates on a voxel-wise basis. The update direction is generally written as

$$\mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathbf{P_Q}}\mathbf{d} = \sum_{j=1}^{J} \sum_{k=1}^{K} \sum_{k'=1}^{K} d_{jk} \left[\frac{1}{\mathcal{P}_j} \left(\frac{1}{\sigma^2_{y_k}} \left(\sum_{p=1}^{P} \left(\frac{1}{\alpha} \frac{\partial \mathcal{H}_1}{\partial y_{j'k'}} + 2\nu \frac{\partial \mathcal{H}_2}{\partial y_{j'k'}} \right) \right) - \frac{\partial \mathcal{P}_j}{\partial y_{j'k'}} \frac{\partial \mathscr{L}_{P,Q_y}(\mathbf{y})}{\partial y_{jk}} \right) \right] d_{jk'}. \tag{41}$$

With the terms $\mathcal{H}_1$ and $\mathcal{H}_2$ defined as

$$\mathcal{H}_1 = w_{jp} K_{jp} \Delta_{jkp}, \tag{42}$$

$$\mathcal{H}_2 = K_{jp} \left(\mathcal{D}_{jkp} - w_{jp} \left(\sum_{p'=1}^{P} \mathcal{D}_{jkp'} \right) \right). \tag{43}$$

Equation (41) is based on the derivatives of $\mathcal{H}_1, \mathcal{H}_2$ and $\mathcal{P}_j$:

$$\begin{cases} \frac{\partial \mathcal{H}_1}{\partial y_{j'k'}} = \frac{\partial w_{jp}}{\partial y_{j'k'}} K_{jp} \Delta_{jkp} + w_{jp} \frac{\partial K_{jp}}{\partial y_{j'k'}} \Delta_{jkp} + w_{jp} K_{jp} \frac{\partial \Delta_{jkp}}{\partial y_{j'k'}}; \\ \frac{\partial \mathcal{H}_2}{\partial y_{j'k'}} = \frac{\partial K_{jp}}{\partial y_{j'k'}} \left(\mathcal{D}_{jkp} - w_{jp} \left(\sum_{p'=1}^{P} \mathcal{D}_{jkp'} \right) \right) + K_{jp} \left(\frac{\partial \mathcal{D}_{jkp}}{\partial y_{j'k'}} - \left(\frac{\partial w_{jp}}{\partial y_{j'k'}} \left(\sum_{p'=1}^{P} \mathcal{D}_{jkp'} \right) \right. \right. \\ \qquad \left. \left. + w_{jp} \left(\sum_{p'=1}^{P} \frac{\partial \mathcal{D}_{jkp'}}{\partial y_{j'k'}} \right) \right) \right); \\ \frac{\partial \mathcal{P}_j}{\partial y_{j'k'}} = \sum_{p=1}^{P} \left(\frac{\partial w_{jp}}{\partial y_{j'k'}} K_{jp} + w_{jp} \frac{\partial K_{jp}}{\partial y_{j'k'}} \right), \end{cases} \tag{44}$$

which are functions of $\frac{\partial w_{jp}}{\partial y_{j'k'}}$, $\frac{\partial \Delta_{jkp}}{\partial y_{j'k'}}$, $\frac{\partial K_{jp}}{\partial y_{j'k'}}$ and $\frac{\partial \mathcal{D}_{jkp}}{\partial y_{j'k'}}$:

$$\begin{cases} \frac{\partial w_{jp}}{\partial y_{j'k'}} = \frac{2\nu}{\sigma^2_{y'_k}} \delta_{jj'} \left(w_{jp} \left(\sum_{p'=1}^{P} \mathcal{D}_{jk'p'} \right) - \mathcal{D}_{jk'p} \right); \\ \frac{\partial K_{jp}}{\partial y_{j'k'}} = -\frac{1}{\alpha \sigma^2_{y_k}} \delta_{jj'} K_{jp} \Delta_{jk'p}; \\ \frac{\partial \Delta_{jkp}}{\partial y_{j'k'}} = \delta_{jj',kk'}; \\ \frac{\partial \mathcal{D}_{jkp}}{\partial y_{j'k'}} = \frac{\partial w_{jp}}{\partial y_{j'k'}} \Delta_{jkp} \pi^{\frac{1}{\nu}}_{jp} + w_{jp} \delta_{jj',kk'} \pi^{\frac{1}{\nu}}_{jp} - \frac{2}{\sigma^2_{y_k}} \delta_{jj'} w_{jp} \Delta_{jkp} \Delta_{jk'p} \pi^{\frac{2}{\nu}}_{jp}. \end{cases} \tag{45}$$

in Eq. (45), $\delta_{jj'}$ and $\delta_{jj',kk'}$ refer to the Kronecker delta. It is left to the reader to combine Eqs. (44) and (45) into Eq. (41), to obtain a final analytical expression. Finally, the contribution of the spatial prior is

$$\mathbf{d}^{\mathrm{T}}\mathbf{H}_{\mathbf{P_S}}\mathbf{d} = \sum_{j=1}^{J}\sum_{k=1}^{K} d_{jk}\beta_k \sum_{r=1}^{J} D_{jr}^{\mathrm{T}}\phi''\left(\sum_{c=1}^{J} D_{rc}y_{ck}, \delta_k\right)\left(\sum_{k'=1}^{K}\left(\sum_{j'=1}^{J} D_{rj'}d_{j'k'}\right)\right), \tag{46}$$

where $\phi''(\cdot)$ is the second derivative of the Huber prior.

Disclosures

No conflicts of interest, financial or otherwise, are declared by the authors.

Acknowledgments

We acknowledge financial support by the Natural Sciences and Engineering Research Council of Canada (NSERC RGPIN/04178-2015 and scholarship PGDS3-490047-2016, as well as a Hydro-Québec fellowship.

References

1. E. Bär et al., "The potential of dual-energy CT to reduce proton beam range uncertainties," *Med. Phys.* **44**(6), 2332–2344 (2017).
2. H. Paganetti, "Range uncertainties in proton therapy and the role of Monte Carlo simulations," *Phys. Med. Biol.* **57**(11), R99 (2012).
3. C. Möhler et al., "Experimental verification of stopping-power prediction from single-and dual-energy computed tomography in biological tissues," *Phys. Med. Biol.* **63**(2), 025001 (2018).
4. V. T. Taasti et al., "Theoretical and experimental analysis of photon counting detector CT for proton stopping power prediction," *Med. Phys.* **45**(11), 5186–5196 (2018).
5. A. Lalonde et al., "The impact of dual-and multi-energy CT on proton pencil beam range uncertainties: a Monte Carlo study," *Phys. Med. Biol.* **63**(19), 195012 (2018).
6. M. Simard et al., "Quantitative imaging performance of mars spectral photon-counting CT for radiotherapy," *Med. Phys.* **47**(8), 3423–3434 (2020).
7. A. Lalonde and H. Bouchard, "A general method to derive tissue parameters for Monte Carlo dose calculation with multi-energy CT," *Phys. Med. Biol.* **61**(22), 8044 (2016).
8. A. Lalonde, E. Bär, and H. Bouchard, "A Bayesian approach to solve proton stopping powers from noisy multi-energy CT data," *Med. Phys.* **44**(10), 5293–5302 (2017).
9. M. Simard et al., "The potential of photon-counting CT for quantitative contrast-enhanced imaging in radiotherapy," *Phys. Med. Biol.* **64**(11), 115020 (2019).
10. S. Faby et al., "Performance of today's dual energy CT and future multi energy CT in virtual non-contrast imaging and in iodine quantification: a simulation study," *Med. Phys.* **42**(7), 4349–4366 (2015).
11. M. Simard, A. Lalonde, and H. Bouchard, "Parametrization of multi-energy Ct projection data with eigentissue decomposition," *Phys. Med. Biol.* **65**(15), 155001 (2020).
12. T. G. Schmidt, R. F. Barber, and E. Y. Sidky, "A spectral CT method to directly estimate basis material maps from experimental photon-counting data," *IEEE Trans. Med. Imaging* **36**(9), 1808–1819 (2017).
13. F. Jolivet, G. Schramm, and J. Nuyts, "One-step inversion algorithms for spectral CT, with application to dynamic cone beam CT," https://arxiv.org/abs/2112.09375 (2021).
14. C. Cai et al., "A full-spectral Bayesian reconstruction approach based on the material decomposition model applied in dual-energy computed tomography," *Med. Phys.* **40**(11), 111916 (2013).
15. Y. Long and J. A. Fessler, "Multi-material decomposition using statistical image reconstruction for spectral CT," *IEEE Trans. Med. Imaging* **33**(8), 1614–1626 (2014).

16. T. Weidinger et al., "Polychromatic iterative statistical material image reconstruction for photon-counting computed tomography," *Int. J. Biomed. Imaging* **2016**, 5871604 (2016).
17. R. F. Barber et al., "An algorithm for constrained one-step inversion of spectral CT data," *Phys. Med. Biol.* **61**(10), 3784 (2016).
18. G. Vilches-Freixas et al., "Comparison of projection-and image-based methods for proton stopping power estimation using dual energy CT," *Phys. Imaging Radiat. Oncol.* **3**, 28–36 (2017).
19. L. A. Lehmann et al., "Generalized image combinations in dual KVP digital radiography," *Med. Phys.* **8**(5), 659–667 (1981).
20. S. Tilley et al., "A general CT reconstruction algorithm for model-based material decomposition," *Proc. SPIE* **10573**, 105731E (2018).
21. A. Malusek et al., "The potential of dual-energy computed tomography for quantitative decomposition of soft tissues to water, protein and lipid in brachytherapy," *Phys. Med. Biol.* **58**(4), 771 (2013).
22. X. Liu et al., "Quantitative imaging of element composition and mass fraction using dual-energy CT: three-material decomposition," *Med. Phys.* **36**(5), 1602–1609 (2009).
23. P. R. S. Mendonça, P. Lamb, and D. V. Sahani, "A flexible method for multi-material decomposition of dual-energy CT images," *IEEE Trans. Med. Imaging* **33**(1), 99–116 (2013).
24. J.-É. Tremblay, S. Bedwani, and H. Bouchard, "A theoretical comparison of tissue parameter extraction methods for dual energy computed tomography," *Med. Phys.* **41**(8), 081905 (2014).
25. H. Q. Woodard and D. R. White, "The composition of body tissues," *Br. J. Radiol.* **59**(708), 1209–1218 (1986).
26. E. Y. Sidky et al., "Three material decomposition for spectral computed tomography enabled by block-diagonal step-preconditioning," https://arxiv.org/abs/1801.06263 (2018).
27. M. Simard et al., "Electron density and effective atomic number estimation in a maximum a posteriori framework for dual-energy computed tomography," *Med. Phys.* **47**(9), 4137–4149 (2020).
28. C. Mory et al., "Comparison of five one-step reconstruction algorithms for spectral CT," *Phys. Med. Biol.* **63**(23), 235001 (2018).
29. S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**(6), 721–741 (1984).
30. S. Geman and D. McClure, "Bayesian image analysis: an application to single photon emission tomography," *Am. Stat. Assoc.* **12**–18 (1985).
31. P. J. Green, "Bayesian reconstructions from emission tomography data using a modified EM algorithm," *IEEE Trans. Med. Imaging* **9**(1), 84–93 (1990).
32. P. J. Huber, "Robust estimation of a location parameter," in *Breakthroughs in Statistics*, S. Kotz and N. L. Johnson, Eds., pp. 492–518, Springer, New York (1992).
33. E. Polak and G. Ribiére, "Note sur la convergence de méthodes de directions conjuguées," *ESAIM: Math. Modell. Numer. Anal.-Modél. Math. Anal. Numér.* **3**(R1), 35–43 (1969).
34. M. Yang et al., "Comprehensive analysis of proton range uncertainties related to patient stopping-power-ratio estimation using the stoichiometric calibration," *Phys. Med. Biol.* **57**(13), 4095 (2012).
35. J. P. Schlomka et al., "Experimental feasibility of multi-energy photon-counting K-edge imaging in pre-clinical computed tomography," *Phys. Med. Biol.* **53**(15), 4031 (2008).
36. J. A. Fessler, "Michigan image reconstruction toolbox," https://web.eecs.umich.edu/fessler/code/index.html (accessed 16/03/18).
37. M. J. Berger and J. H. Hubbell, *XCOM: Photon Cross Sections on a Personal Computer (No. NBSIR-87-3597)*, National Bureau of Standards, Center for Radiation Research, Washington, DC (1987).
38. C. O. Schirra et al., "Statistical reconstruction of material decomposed data in spectral CT," *IEEE Trans. Med. Imaging* **32**(7), 1249–1257 (2013).
39. P. K. Kijewski and B. E. Bjärngard, "Correction for beam hardening in computed tomography," *Med. Phys.* **5**(3), 209–214 (1978).
40. C. Lemmens, D. Faul, and J. Nuyts, "Suppression of metal artifacts in CT using a reconstruction procedure that combines map and projection completion," *IEEE Trans. Med. Imaging* **28**(2), 250–260 (2008).

41. K. Nakada et al., "Joint estimation of tissue types and linear attenuation coefficients for photon counting CT," *Med. Phys.* **42**(9), 5329–5341 (2015).
42. Z. Yu et al., "Spectral prior image constrained compressed sensing (spectral PICCS) for photon-counting computed tomography," *Phys. Med. Biol.* **61**(18), 6707 (2016).

Mikaël Simard received his BEng degree in engineering physics from École Polytechnique de Montréal in 2014, and his MS degree in medical physics from McGill University in 2016. He received his PhD in physics from Université de Montréal in February 2021, and is now a postdoctoral fellow at University College London. His research interests involve the use of quantitative imaging (multi-energy CT and proton radiographs) for radiotherapy applications.

Hugo Bouchard is a professor of physics at Université de Montréal, where he received his PhD in 2010. He currently holds a Canada Research Chair on emerging applications of spectral computed tomography. His research interests include photon-counting CT and various applications of the Monte Carlo method for radiation transport simulation.