Spectral image clustering on dual-energy CT scans using functional regression mixtures

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

28-01-2022 / 04-02-2022

CITATION

Brivet, Segolene; Chamroukhi, Faicel; Coates, Mark; Forghani, Reza; Savadjiev, Peter (2022): Spectral image clustering on dual-energy CT scans using functional regression mixtures. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.19089995.v1>

DOI

[10.36227/techrxiv.19089995.v1](https://doi.org/10.36227/techrxiv.19089995.v1)

Spectral image clustering on dual-energy CT scans using functional regression mixtures

Segolene Brivet, Faicel Chamroukhi, Mark Coates,
Reza Forghani, and Peter Savadjiev

Abstract—Dual-energy computed tomography (DECT) is an advanced CT scanning technique enabling material characterization not possible with conventional CT scans. It allows the reconstruction of energy decay curves at each 3D image voxel, representing varying image attenuation at different effective scanning energy levels. In this paper, we develop novel functional data analysis (FDA) techniques and adapt them to the analysis of DECT decay curves. More specifically, we construct functional mixture models that integrate spatial context in mixture weights, with mixture component densities being constructed upon the energy decay curves as functional observations. We design unsupervised clustering algorithms by developing dedicated expectation maximization (EM) algorithms for the maximum likelihood estimation of the model parameters. To our knowledge, this is the first article to adapt statistical FDA tools and model-based clustering to take advantage of the full spectral information provided by DECT. We evaluate our methods on 91 head and neck cancer DECT scans. We compare our unsupervised clustering results to tumor contours traced manually by radiologists, as well as to several baseline algorithms. Given the inter-rater variability even among experts at delineating head and neck tumors, and given the potential importance of tissue reactions surrounding the tumor itself, our proposed methodology has the potential to add value in downstream machine learning applications for clinical outcome prediction based on DECT data in head and neck cancer.

Index Terms—Spectral image clustering; Dual-Energy CT data; Mixture Models; Mixtures of Regressions; Maximum Likelihood Estimation; Functional Data Analysis

I. INTRODUCTION

Computed Tomography (CT) has been one of the most common and widespread imaging techniques used in the clinic for the last few decades. There is increasing interest in a more advanced CT technique known as dual-energy CT (DECT)

This research is supported by Fonds de recherche du Québec - Santé (FRQS) and Fondation de l'association des radiologistes du Québec (FARQ) (R.F.); by the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference number 260250] (M.C.); and by the French National Research Agency (ANR), [grant SMILES ANR-18-CE40-0014] (F.C.).

S. Brivet and M. Coates are in the Electrical and Computer Engineering Department, McGill University, Montreal, QC H3A 0G4, Canada (e-mail: segolene.brivet@mail.mcgill.ca; mark.coates@mcgill.ca).

F. Chamroukhi is with Normandie Univ, UNICAEN, CNRS, LMNO, 14000 Caen, France (e-mail: faicel.chamroukhi@unicaen.fr).

R. Forghani and P. Savadjiev are in the Augmented Intelligence and Precision Health Laboratory (AIPHL), Department of Radiology, McGill University, Montreal, QC H3G 1A4, Canada (reza.forghani@mcgill.ca; peter.savadjiev@mcgill.ca)

or spectral CT that enables additional material or tissue characterization beyond what is possible with conventional CT. In particular, DECT makes possible the computation of image attenuation levels at multiple effective energy levels. This results in the association of a decay curve with each reconstructed image voxel, representing energy-dependent changes in attenuation at that body location.

In this paper, we evaluate the application of DECT to head and neck squamous cell carcinoma (HNSCC). Current image-based evaluation of HNSCC tumors in clinical practice is largely qualitative, based on a visual assessment of tumor anatomic extent and basic one- or two-dimensional tumor size measurements. However, the frequently complex shape of mucosal head and neck cancers and at times poorly defined boundaries and potential adjacent tissue reactions can result in a high inter-observer variability in defining the extent of the tumor [1]–[4], especially among radiologists without subspecialty expertise in head and neck imaging. Furthermore, there is strong evidence that alterations of gene expression and protein-protein interactions in the peri-tumoral tissue or normal adjacent tissue may play a critical role in the evolution and risk of recurrence of HNSCC tumors (e.g., [5], [6]). Yet these adjacent regions may not be obviously salient upon visual image examination. For all these reasons, the accurate and consistent determination of a predictive region around the tumor is essential, both for conventional staging which determines patient management and for future automated quantitative image-based predictive algorithms based on machine learning. Such a predictive region may include the tumor itself, but also surrounding tissues of biological relevance.

With these considerations in mind, DECT provides new and unexplored opportunities to answer an important question: what can be learned from DECT about tumor heterogeneity and its associations with surrounding tissues? As a first step towards answering this question, in this paper, we adapt functional data analysis techniques to DECT data in order to explore underlying patterns of association in and around the tumor. Functional data analysis (FDA) is a classical branch of statistics dedicated to the analysis of functional data, in situations where each data object is considered a function. This is particularly appropriate for DECT image data, as each 3D image voxel is associated with a curve of image intensity decay over multiple reconstructed energy levels (more details are provided below in Section II-A). Thus, we adapt FDA statistical models for the clustering of 3D image voxels based

on the full functional information provided by the decay curves associated with each voxel. More specifically, the architecture underpinning our proposed method is a functional mixture model, where the mixture component densities are built upon functional approximation of the spectral decay curves at each image voxel, and the mixture weights are constructed to integrate spatial constraints. We then derive an Expectation-Maximization (EM) algorithm to perform the maximum likelihood estimators of the model parameters.

To our knowledge, this is the first article to propose spatial clustering utilizing the full spectral information available in DECT data, based on an appropriate FDA statistical framework. Existing methods for automatic tumor delineation in DECT (reviewed in detail in Section II-B) are mostly based on deep learning techniques and utilize only a small subset of the available information, due to the sheer amount of 4D (spatial + spectral) data available in a single DECT scan.

We apply the proposed methodology on 91 DECT scans of HNSCC tumors and we compare our results to manually traced tumor contours performed by an experienced expert radiologist. We also compare to other baseline clustering methods. However, tumor segmentation on its own is not a clinical outcome. A full demonstration of the clinical utility of our method necessitates an analysis of its ability to predict actual clinical outcomes, and how this prediction performance compares to the performance in the case of manually drawn contours, or contours drawn using alternative automatic methods. We leave this prediction analysis for a subsequent paper. As the first article to adapt FDA statistical tools to DECT data, the main focus of the present paper is on the statistical methodology and on algorithm development. As such, we can summarize our contributions as follows:

- 1) we extend the statistical framework of functional mixture models to the spatio-spectral heterogeneous DECT data. In particular, DECT energy decay curves observed at each image voxel are modeled as spatially-distributed functional observations.
- 2) the proposed model is a novel generative mixture density model, with an efficient EM learning algorithm.
- 3) we develop an unsupervised clustering of voxels incorporating full spectral information from DECT data.
- 4) to our knowledge, this is the first time that FDA statistical tools are applied to the analysis of DECT data.

The source codes with summarising pseudo-codes for the proposed methods are publicly available¹, free of charge.

II. BACKGROUND AND RELATED WORK

A. Dual-energy CT

DECT data may be viewed as a 4D image of a patient: a 3D body volume over a range of energy levels. The dual-energy image acquisition using two x-ray energy peaks at the source provides enough attenuation information to be combined and to be able to reconstruct a curve at multiple “virtual monochromatic” energy levels. These simulate what the attenuation (in Hounsfield units; HU) would be if the

(a) 2D slices of VMIs at 40,65,140keV with tumor contour in red. At lower energy levels, VMIs are more contrasted; at higher levels, VMIs are less noisy. A VMI at 65keV is similar to a standard CT scan.

(b) Examples of decay curves for different body locations. A blue (resp. red, green) curve represents attenuation information stored at one voxel within bone (resp. tumor; tissue).

study was acquired with a monochromatic x-ray beam at that energy value (in kilo-electron-Volt; keV). The reconstructed curve of attenuation numbers over each energy level translates the energy-dependent changes and is commonly called the spectral Hounsfield unit attenuation curve, or an *energy decay curve* [7]. In our method, we will make use of this spectral information through functional approximations, and thus consider the curves as functional observations. An energy decay curve has been calculated for each image voxel, thus a DECT scan is represented as a 4D image with 3 dimensions for X, Y and Z spatial coordinates and 1 dimension for energy level coordinates. The virtual monochromatic image (VMI) is the 3D image representation at a given energy level. See Fig. 1a for examples of a 2D slice from different VMIs and Fig. 1b for examples of decay curves for different tissue characteristics.

B. Segmentation of dual-energy CT data

Segmentation is a process of delineating an image region of interest. For example, radiation oncologists usually

¹https://github.com/segobrivet/DECT_clustering

manually segment tumors for radiation planning. Automatic tumor segmentation has a long history of developments: from knowledge-driven early techniques to data-driven newer techniques, algorithms aim to extract image features to make a decision on region boundaries [8]. However this process remains challenging in medical imaging due to the heterogeneity over the image or the acquisition process, most of current algorithms need manual adjustments on the result [9].

In head and neck CT imaging, the difficulty to contour precisely a tumor region results in large inter-observer variability in segmentation results, even among trained radiation oncologists. A study among radiologists from 14 different institutions obtained a median Dice Similarity Score (DSC) ranging from 0.51 to 0.82 [3], depending on the delineation criteria used. Another study assessing the same variability among 3 experienced radiologists over 10 tumors obtained a mean DSC of 0.57 [4].

To the best of our knowledge, only a few studies have focused on DECT segmentation. These employ deep learning approaches [10]–[12]. The four dimensions of the data required workarounds in order to apply neural networks. For example, using two VMIs sampled from the energy level spectrum, one at a low- and one at a high-energy level, Chen et al. in [10] merge the two VMIs in a layer connected to a U-Net architecture [13]. Wang et al. in [12] learn features from two pyramid networks on the two VMIs independently and combine them through deep attention into a mask scoring regional convolutional neural network (R-CNN). They achieved good performance in segmenting large-sized organs (DSC larger than 0.8); and performance is less impressive for small-sized organs (DSC between 0.5 and 0.8).

C. Decay curve clustering via functional data analysis

FDA aims to represent infinite-dimensional functional data into a finite-dimensional vector of coefficients [14]. To achieve this, FDA consists in expanding functional data into function bases. One approach relies on projection on bases which consists in projecting functional data onto finite dimensional function bases (e.g. splines, B-splines, polynomials, Fourier, wavelet). It associates a finite vector of projection coefficients. This is what we use in this paper. Analogously another common approach would be to run a functional principal component analysis (fPCA) to get a basis of eigenfunctions of the covariance of the process describing our functional data. It associates a (truncated) projection vector of PCA coefficients.

Our objective is to partition our data, modeled with FDA, in different groups of voxels having similar decay curve characteristics. Among the available clustering approaches (e.g. centroid-based clustering such as k-means, connectivity-based clustering such as hierarchical clustering, density-based clustering such as DBSCAN [15], distribution-based clustering with model-based methods), since we have a model for each decay curve, a model-based approach is preferred.

Model-based clustering is a thoroughly developed field [16], [17], in particular for multivariate analysis. Model-based clustering approaches rely on the finite mixture modeling framework [18] to represent the density of a set of independent

multivariate observations and on an optimization algorithm to automatically find a partition into groups of such observations.

To represent different groups of data, mixture models assume each datum to follow a known distribution (e.g. Gaussian), and build a mean representation (i.e. model) for each group of data. The mixture model calculates, for each data point, a value defined by the sum over $k = \{1 \dots \# \text{groups}\}$, of the probability distribution function (pdf) that this point belongs to group k model, emphasized by a weight giving a higher or lower chance of belonging to this group (derived in Section III-A). Mixture models have the advantage of being interpretable, parametric thus well-understood, and flexible as the pdf modeling the data in each cluster is chosen explicitly.

The Expectation-Maximization (EM) algorithm [19] is a popular and adapted tool with desirable properties that can be used to conduct an iterative estimation of the model parameters and thus the cluster membership probabilities.

Mixture models for clustering have been applied and adapted to different kind of data, including time-series data [20], gene expression data [21], or 3D noisy medical images [22], spatio-temporal data (non image data) [23]. They also have been recently investigated for functional data [24], and thus provide an avenue to model the spectral decay curves, but in this context of spectral images, we also need to incorporate the spatial information into the clustering.

A related idea was proposed in [25] to develop a spatio-temporal mixture of Hidden Process Models for fMRI analysis. The authors built a temporal probabilistic model, and reshaped the prior probability with spatial constraints to determine a “region of influence” for the temporal model. A specification of our model covers this approach, and our model goes further in generalizing it via the construction of more flexible Gaussian-mixture weights around the spatial coordinates. The resulting model enjoys better numerical learning properties with faster convergence due to closed-form updating rules, for the spatial weights parameters. Alternative constructions of the proposed mixture model are also presented in order to validate the method and to accommodate potential user specifications.

III. METHODOLOGY

A. Generative modeling framework

We adopt the framework of generative modeling for image clustering using different families of extended mixture distributions. The general form of the generative model for the image assumes that the i th datum (i.e., pixel, voxel) in the image has the general semi-parametric mixture density:

$$f_i(\boldsymbol{\theta}) = \sum_{k=1}^K \pi_{ik} f_i(\boldsymbol{\theta}_k), \quad (1)$$

which is a convex combination of K component densities, $f_i(\boldsymbol{\theta}_k)$, $k \in [K] = \{1, \dots, K\}$, weighted by non-negative mixture weights π_{ik} that sum to one, that is $\sum_{k=1}^K \pi_{ik} = 1$ for all i , $i \in [n]$. The unknown parameter vector $\boldsymbol{\theta}$ of density (1) is composed of the set of component density parameters $\{\boldsymbol{\theta}_k\}$ and their associated weights $\{\pi_{ik}\}$, i.e., $\boldsymbol{\theta} = \{\pi_{ik}, \boldsymbol{\theta}_k\}_{k=1}^K$.

From the perspective of model-based clustering of the image, each component density can be associated with a cluster,

and hence the clustering problem becomes one of parametric density estimation. Suppose that the image has K segments and let $Z_i \in [K]$ be the random variable representing the unknown segment label of the i th observation in the image. Suppose that the distribution of the data within each segment $k \in [K]$ is $f(\theta_k)$, i.e., $f_{i|Z_i=k}(\theta) = f_i(\theta_k)$. Then, from a generative point of view, model (1) is equivalent to *i*) sampling a segment label according to the discrete distribution with parameters being the mixture weights $\pi = \{\pi_1, \dots, \pi_K\}$, then *ii*) sampling an observation Im_i from the conditional distribution $f(\theta_k)$. Given a model of the form (1) represented by $\hat{\theta}$, typically fitted by maximum likelihood estimation (MLE) from the n observations composing the image Im_n , as

$$\hat{\theta} \in \arg \max_{\theta \in \Theta} \log L(\theta | \text{Im}_n) \quad (2)$$

where $L(\theta | \text{Im}_n)$ is the likelihood of θ given the image data Im_n , then, the segment labels can be determined via the Bayes' allocation rule,

$$\hat{Z}_i = \arg \max_{k \in [K]} \mathbb{P}(Z_i = k | \text{Im}(i); \hat{\theta}), \quad (3)$$

which consists of maximizing the conditional probabilities

$$\mathbb{P}(Z_i = k | \text{Im}(i); \hat{\theta}) = \frac{\hat{\pi}_{ik} f_i(\hat{\theta}_k)}{f_i(\hat{\theta})}. \quad (4)$$

that the i th observation originates from segment k , $k \in [K]$, given the image data and the fitted model.

Model (1) has many different specifications in the literature depending on the nature of the data generative process, resulting in a multitude of choices for the mixture weights and for the component densities. Mixtures of multivariate distributions [18] are in particular more popular in model-based clustering of vectorial data using multivariate mixtures. These include multivariate Gaussian mixtures [16], [17], where $\pi_{ik} = \pi_k, \forall i$, are constant mixture weights, and the component densities $f_i(\theta_k) = \phi_i(\mu_k, \Sigma_k)$ are multivariate Gaussians with means μ_k and covariance matrices Σ_k . Mixtures of regressions models, introduced in [26], are common in modeling and clustering of regression-type data. For example, in the widely used Gaussian regression mixture model [27], we have constant mixing proportions, i.e. $\pi_{ik} = \pi_k$, and the mixture components $f_i(\theta_k)$'s are Gaussian regressors $\phi(\cdot, \beta_k^\top x_i, \sigma_k^2)$ with typically linear means $\beta_k^\top x_i$ and variances σ_k^2 , in the case of univariate response.

In this paper, we consider a more flexible mixture of regressions model in which both the mixture weights and the mixture components are covariate-dependent, and are constructed upon flexible semi-parametric functions. More specifically, in this full conditional mixture model, the mixture weights π_{ik} are constructed upon parametric functions $\pi_{ik} = \pi_k(\cdot, x_i; \alpha)$ of some covariates x_i 's represented by a parameter vector α , and the regression functions $f_i(\theta_k)$ are Gaussian regressors $\phi(\cdot, \mu(x_i; \beta_k), \sigma_k^2)$ with semi-parametric (non-)linear mean functions $\mu(x_i; \beta_k)$. This flexible modeling allows us to better capture more non-linear relationships in the functional data via the semi-parametric mean functions. Heterogeneity is accommodated via the mixture distribution, and spatial organization

can be captured via spatial-dependent mixture weights.

B. Spatial mixture of functional regressions for dual-energy CT images

We propose a spatialized mixture of functional regressions model, adapted to the given type of image data, for the model-based clustering of dual-energy CT scans. The images we analyse include spectral curves for each 3D-voxel. Each image, denoted Im , is represented as a sample of n observations, $\text{Im} = \{v_i, x_i, y_i\}_{i=1}^n$ where $v_i = (v_{i1}, v_{i2}, v_{i3})$ is the i th voxel 3D-spatial coordinates. The i th voxel is represented by the curve (x_i, y_i) composed of HU attenuation values $y_i = (y_{i1}, \dots, y_{im})$ measured at energy levels (covariates) $x_i = (x_{i1}, \dots, x_{im})$, with m being the number of energy levels.

To accommodate the spatial organization of the image together with the functional nature of each of its voxels, we propose spatialized conditional extensions of the general family of model (1), in which we model the i th voxel observation of the image using the conditional density $f(y_i | x_i, v_i; \theta)$ that relates the attenuation curve levels y_i given the associated energy levels x_i , and spatial location v_i , via a convex combination of (non-)linear (semi-)parametric functional regressors $f(y_i | x_i; \theta_k)$ with spatial weights $\pi_k(v_i; \alpha)$, that is:

$$f(y_i | x_i, v_i; \theta) = \sum_{k=1}^K \pi_k(v_i; \alpha) f(y_i | x_i; \theta_k). \quad (5)$$

To this purpose, we consider two different spatial constructions of the mixing weights (gating functions) $\pi_k(v_i; \alpha)$: (i) softmax gates; and (ii) normalized Gaussian gates. The latter is an appropriate choice if more approximation quality is needed, and facilitates the computations in the learning process. We also consider different families to model the functional regressors, including spline and B-spline regression functions that enjoy better curve approximation capabilities, compared to linear or polynomial regression functions.

1) Functional regression components: We have a 3D image volume over a range of energy levels that provide, for each voxel i , an attenuation curve (x_i, y_i) which represents energy dependent changes in attenuation, which enables a better characterization of the tissue at voxel i . We therefore model the component densities $f(y_i | x_i; \theta_k)$ as functional regression models constructed upon the attenuation curves as functional observations. This allows us to accommodate the spectral curve nature of the data. More specifically, in the case of univariate energy levels, we use smooth functional approximations to model, for the i th voxel, the mean spectral curve of the k th component $\mu(x_i; \beta_k) = \mathbb{E}_\theta[Y_i | Z_i = k, x_i]$, that is $\mu(x_i; \beta_k) = (\mu(x_{i1}; \beta_k), \dots, \mu(x_{im}; \beta_k))$ using polynomial or (B)-spline functions, whose coefficients are β_k 's.

The conditional density model for each regression is then modeled as a functional Gaussian regressor defined by $f(y_i | x_i; \theta_k) = \phi_m(y_i; \mu(x_i; \beta_k), \sigma_k^2 \mathbf{I})$, with $\mu(x_i; \beta_k) = \mathbf{B}(x_i) \beta_k$ being the function approximation onto polynomial or (B)-spline bases $\mathbf{B}(x_i)$, and the matrix form of the functional regression model is then given by

$$f(y_i | x_i; \theta_k) = \phi_m(y_i; \mathbf{B}(x_i) \beta_k, \sigma_k^2 \mathbf{I}), \quad (6)$$

where $\theta_k = (\beta_k^\top, \sigma_k^2)^\top \in \mathbb{R}^{p+q+2}$ is the unknown parameter vector of regression k .

2) *Spatial gating functions*: The constructed functional mixture of regressions model (5) specifically integrates the spatial constraints in the mixture weights $\pi_k(\mathbf{v}_i; \alpha)$ via functions of the spatial locations $\mathbf{v}_i$ parametrized by vectors of coefficients α . We investigate two choices to this end. The first proposed model is a spatial softmax-gated functional mixture of regression and is defined by (5) with a softmax gating function:

$$\pi_k(\mathbf{v}_i; \alpha) = \frac{\exp(\alpha_k^\top \mathbf{v}_i)}{1 + \sum_{k'=1}^{K-1} \exp(\alpha_{k'}^\top \mathbf{v}_i)}, \quad (7)$$

where $\alpha = (\alpha_1^\top, \dots, \alpha_K^\top)^\top$ is the unknown parameter vector of the gating functions. We will refer to this model, defined by (5), (6) and (7), as the Spatial Softmax-gated Mixture of Functional Regressions, abbreviated as **SsMFR**. The softmax modeling of the mixture weights is a standard choice known in the mixtures-of-experts community. However, its optimization performed at the M-Step of the EM algorithm, is not analytic and requires numerical Newton-Raphson optimization. This can become costly, especially in larger image applications, such as the one we address.

In the second proposed model, we use a spatial Gaussian-gated functional mixture of regressions, defined by (5) with a Gaussian-gating function:

$$\pi_k(\mathbf{v}_i; \alpha) = \frac{w_k \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k)}{\sum_{\ell=1}^K w_\ell \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_\ell, \mathbf{R}_\ell)}, \quad (8)$$

in which w_k are non-negative weights that sum to one, $\phi_d(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k)$ is the density function of a multivariate Gaussian vector of dimension d with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$, and $\alpha = (\alpha_1^\top, \dots, \alpha_K^\top)^\top$ is the parameter vector of the gating functions with $\alpha_k = (w_k, \boldsymbol{\mu}_k^\top, \text{vech}(\mathbf{R}_k)^\top)^\top$.

We will refer to this model, defined by (5), (6) and (8), as the Spatial Gaussian-gated Mixture of Functional Regressions, abbreviated as **SgMFR**. This Gaussian gating function was introduced in [28] to bypass the need for an additional numerical optimization in the inner loop of the EM algorithm. We obtain a closed form updating formula at the M-Step, that is detailed in the next section presenting the derived EM algorithm.

3) *MLE of the SgMFR model via the EM algorithm*: Based on equations (5), (6) and (8), the SgMFR joint density $f(\mathbf{y}_i, \mathbf{x}_i, \mathbf{v}_i; \theta)$ is then derived and the joint log-likelihood we maximize via EM is

$$\begin{aligned} \log L(\theta) &= \sum_{i=1}^n \log f(\mathbf{y}_i, \mathbf{x}_i, \mathbf{v}_i; \theta) \\ &= \sum_{i=1}^n \log \sum_{k=1}^K w_k \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k) \phi_m(\mathbf{y}_i; \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I}). \end{aligned} \quad (9)$$

The complete-data log-likelihood, upon which the EM algorithm is constructed, is

$$\begin{aligned} \log L_c(\theta) &= \\ &= \sum_{i=1}^n \sum_{k=1}^K Z_{ik} \log [w_k \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k) \phi_m(\mathbf{y}_i; \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I})], \end{aligned} \quad (10)$$

where Z_{ik} is an indicator variable such that $Z_{ik} = 1$ if $Z_i = k$ (i.e., if the i th pair $(\mathbf{x}_i, \mathbf{y}_i)$ is generated from the k th regression

component) and $Z_{ik} = 0$, otherwise. The EM algorithm, after starting with an initial solution $\theta^{(0)}$, alternates between the E- and the M- steps until convergence (when there is no longer a significant change in the log-likelihood (9)).

a) *The E-step*: Compute the expected complete-data log-likelihood (10), given image Im_n and the current estimate $\theta^{(t)}$:

$$\begin{aligned} Q(\theta; \theta^{(t)}) &= \mathbb{E} [L_c(\theta) | \text{Im}_n; \theta^{(t)}], \\ &= \sum_{i=1}^n \sum_{k=1}^K \tau_{ik}^{(t)} \log [\alpha_k \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k) \phi_m(\mathbf{y}_i; \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I})], \end{aligned} \quad (11)$$

where $\tau_{ik}^{(t)} = \mathbb{P}(Z_i = k | \mathbf{y}_i, \mathbf{x}_i, \mathbf{v}_i; \theta^{(t)})$ given by

$$\tau_{ik}^{(t)} = \frac{w_k^{(t)} \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k^{(t)}, \mathbf{R}_k^{(t)}) \phi_m(\mathbf{y}_i; \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k^{(t)}, \sigma_k^{2(t)} \mathbf{I})}{f(\mathbf{v}_i, \mathbf{x}_i, \mathbf{y}_i; \theta^{(t)})} \quad (12)$$

is the posterior probability that the observed pair $(\mathbf{x}_i, \mathbf{y}_i)$ is generated by the k th regressor. This step therefore only requires the computation of the posterior component membership probabilities $\tau_{ik}^{(t)}$ ($i = 1, \dots, n$), for $k = 1, \dots, K$.

b) *The M-step*: Calculate the parameter vector update $\theta^{(t+1)}$ by maximizing the Q -function (11), i.e. $\theta^{(t+1)} = \arg \max_{\theta} Q(\theta; \theta^{(t)})$. By decomposing the Q -function as

$$Q(\theta; \theta^{(t)}) = \sum_{k=1}^K Q(\alpha_k; \theta^{(t)}) + Q(\theta_k; \theta^{(t)}), \quad (13)$$

where $Q(\alpha_k; \theta^{(t)}) = \sum_{i=1}^n \tau_{ik}^{(t)} \log [w_k \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k, \mathbf{R}_k)]$ and $Q(\theta_k; \theta^{(t)}) = \sum_{i=1}^n \tau_{ik}^{(t)} \log [\phi_m(\mathbf{y}_i; \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k, \sigma_k^2 \mathbf{I})]$, the maximization can then be done by K separate maximizations w.r.t. the parameters of the gating and the regression functions. *Updating the gating functions parameters*: Maximizing (13) w.r.t. α_k 's corresponds to the M-Step of a Gaussian Mixture Model [18]. The closed-form expressions for updating the parameters are given by:

$$w_k^{(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} / n, \quad (14)$$

$$\boldsymbol{\mu}_k^{(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} \mathbf{v}_i / \sum_{i=1}^n \tau_{ik}^{(t)}, \quad (15)$$

$$\mathbf{R}_k^{(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} (\mathbf{v}_i - \boldsymbol{\mu}_k^{(t+1)}) (\mathbf{v}_i - \boldsymbol{\mu}_k^{(t+1)})^\top / \sum_{i=1}^n \tau_{ik}^{(t)}. \quad (16)$$

Updating the regression functions parameters: Maximizing (13) w.r.t. θ_k 's corresponds to the M-Step of standard mixtures of experts with univariate Gaussian regressions. The closed-form updating formulas are given by:

$$\boldsymbol{\beta}_k^{(t+1)} = \left[\sum_{i=1}^n \tau_{ik}^{(t)} \mathbf{B}^\top(\mathbf{x}_i) \mathbf{B}(\mathbf{x}_i) \right]^{-1} \sum_{i=1}^n \tau_{ik}^{(t)} \mathbf{B}(\mathbf{x}_i)^\top \mathbf{y}_i, \quad (17)$$

$$\sigma_k^{2(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} (\mathbf{y}_i - \mathbf{B}(\mathbf{x}_i) \boldsymbol{\beta}_k^{(t+1)})^2 / \sum_{i=1}^n \tau_{ik}^{(t)} m_i. \quad (18)$$

C. Alternative two-fold approaches

We also investigate an alternative approach to the one derived before, which consists of a two-fold approach, rather than simultaneous functional approximation and model estimation for segmentation. We first construct approximations of the functional data onto polynomial or (B-)splines bases $\mathbf{B}(\mathbf{x}_i)$

via MLE (ordinary least squares in this case) to obtain:

$$\hat{\beta}_i = [\mathbf{B}(\mathbf{x}_i)^\top \mathbf{B}(\mathbf{x}_i)]^{-1} \mathbf{B}(\mathbf{x}_i)^\top \mathbf{y}_i. \quad (19)$$

Then, we model the density of the resulting coefficient vectors $\hat{\beta}_i$, which is regarded as the i th curve representative, by a mixture density with spatial weights of the form

$$f(\hat{\beta}_i, \mathbf{v}_i, \boldsymbol{\theta}) = \sum_{k=1}^K \pi_k(\mathbf{v}_i; \boldsymbol{\alpha}) \phi_d(\hat{\beta}_i; \mathbf{m}_k, \mathbf{C}_k), \quad (20)$$

where $\mathbf{m}_k$ and $\mathbf{C}_k$ are the mean and the covariance matrix of each component. The spatial weights $\pi_k(\mathbf{v}_i; \boldsymbol{\alpha})$ are normalized Gaussians as in (8) or softmax as in (7). We will refer to these methods as Spatial Gaussian-gated (resp. softmax-gated) Mixtures of Vectorized Functional Regressions, **SgMVFR** (resp. **SsMVFR**).

The EM algorithm for fitting this mixture of spatial mixtures, constructed upon pre-computed polynomial or (B-)spline coefficients with its two variants for modeling the spatial weights, takes a similar form to the previously presented algorithm, and is summarized as follows. The conditional memberships of the E-Step are given for the softmax-gated model by

$$\tau_{ik}^{(t)} = \frac{\pi_k(\mathbf{v}_i; \boldsymbol{\alpha}^{(t)}) \phi_d(\hat{\beta}_i; \mathbf{m}_k^{(t)}, \mathbf{C}_k^{(t)})}{f(\hat{\beta}_i; \mathbf{v}_i; \boldsymbol{\theta}^{(t)})}, \quad (21)$$

and for the Gaussian-gated model by

$$\tau_{ik}^{(t)} = \frac{w_k^{(t)} \phi_3(\mathbf{v}_i; \boldsymbol{\mu}_k^{(t)}, \mathbf{R}_k^{(t)}) \phi_d(\hat{\beta}_i; \mathbf{m}_k^{(t)}, \mathbf{C}_k^{(t)})}{f(\mathbf{v}_i; \hat{\beta}_i; \boldsymbol{\theta}^{(t)})}. \quad (22)$$

The latter has the same advantage as explained above. In the M-Step, the gating functions parameter updates are given by (14), (15), and (16) for the Gaussian-gated model, or through a Newton-Raphson optimization algorithm for the softmax-gated model. The component parameter updates are those of classical multivariate Gaussian mixtures

$$\mathbf{m}_k^{(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} \hat{\beta}_i / \sum_{i=1}^n \tau_{ik}^{(t)}, \quad (23)$$

$$\mathbf{C}_k^{(t+1)} = \sum_{i=1}^n \tau_{ik}^{(t)} (\hat{\beta}_i - \mathbf{m}_k^{(t+1)}) (\hat{\beta}_i - \mathbf{m}_k^{(t+1)})^\top / \sum_{i=1}^n \tau_{ik}^{(t)}. \quad (24)$$

In a nutshell, to compute a clustering of the image, the label of voxel i given the fitted parameters $\boldsymbol{\theta}$ is calculated by the Bayes' allocation rule (3), in which $\text{Im}(i)$ is the spatial coordinates of voxel i with, either its direct spectral curve representative $(\mathbf{x}_i, \mathbf{y}_i)$, or its pre-calculated functional approximation coefficients $\hat{\beta}_i$ given by (19).

IV. EXPERIMENTAL METHODS

In this section we describe the evaluation of different versions of our method: mixtures of B-spline and polynomial functional regressions with spatial Gaussian gates (resp. SgMFR-Bspl, SgMFR-poly), mixture of B-spline regressions with softmax gates (SsMFR-Bspl), mixture of vectorized B-spline regressions with Gaussian gates (SgMVFR-Bspl).

A. Data

91 head and neck DECT scans were evaluated, consisting of head and neck squamous cell carcinoma (HNSCC) tumors

of different sizes and stages from different primary sites. 34% of tumors in our dataset are located in oral cavity, 26% in oropharynx, 21% in larynx, 8% in hypopharynx and 11% in other locations. The tumors T-stage [29] ranges from T1 to T4. 75% of patients were coming for a first diagnostic while 25% were recurrent patients. Institutional review board approval was obtained for this study with waiver of informed consent. Tumors were contoured by an expert head and neck radiologist. All scans were acquired using a fast kVp switching DECT scanner (GE Healthcare) after administration of IV contrast and reconstructed into 1.25mm sections of axial slices with a resolution of 0.61mm, as previously described [30]. Multienergy VMIs were reconstructed at energy levels from 40 to 140 keV in 5 keV increments at the GE Advantage workstation (4.6; GE Healthcare).

In each DECT scan, we crop Volumes of Interest (VOIs) of size 150*150*6 containing a tumor, along with the 21-point-spectral curve associated to each selected voxel, in order to reduce the computational demands for an exploratory study, and to exclude regions containing a majority of air voxels around the body. A pre-processing step is also applied to mask any remaining air voxels in the VOI to focus the clustering on tissues.

B. Regularization parameter

In our study, we augmented the statistical estimator in equation (16) of the covariance matrix of the spatial-coordinates within cluster k , with a regularization parameter $\lambda \in (0, 1]$, which controls the amount of spatial dispersion (neighborhood) taken into account in the spatial mixture weights, by

$$\tilde{\mathbf{R}}_k^{(t+1)} = \lambda \mathbf{R}_k^{(t+1)}. \quad (25)$$

By doing so, we can numerically control the amount of data within cluster k (i.e. its volume). Indeed, if we decompose $\mathbf{C}_k = \lambda \mathbf{D} \mathbf{A} \mathbf{D}^T$ where $\mathbf{A}$ is the $\mathbf{C}_k$'s eigenvectors matrix, and $\mathbf{D}$ a diagonal matrix whose diagonal elements are the eigenvalues in decreasing order, then λ is the volume of cluster k . Since the tumor cluster has in general no strong spatial dispersion, then in practice, we take small values of order 0.1 (see examples in Fig. 4).

C. Parameter initialization

For the sake of reproducibility, we start by initializing the regression mixture and weight parameters of the EM algorithm with a coarse clustering solution given by a Voronoi diagram. We build up Voronoi tiles over the selected voxels in the VOI (a square region where air voxels are deleted) with k-means algorithm applied only on spatial coordinates. To avoid any randomness in the workflow, we start k-means with fixed, regularly spaced, cluster centers, and in the rare cases when EM does not converge following this initialization, we can re-start k-means with random cluster centers.

Then to fix K , the parameter defining the number of clusters, and to fix λ , the spatio-spectral hyper-parameter, we assess on a small training set (10 patients) the three metrics described in Section IV-E. Because computational time can be long, a full grid search would be too expensive, so an iterative

process of fixing K and varying λ , then fixing λ and varying K , then fixing again K and varying λ is applied. The process is run similarly for both methods of mixtures of functional regressions and mixtures of vectorized functional regressions. The range of search values is adapted for each method, and 5 search values are taken in each range. In the end, we use the mean of optimal values over the patients in the train set.

D. Baselines

We compare the quantitative and qualitative performance of our methodology with three baseline algorithms. First, we implement a Gaussian Mixture Model (GMM) with the iterative EM algorithm, using the standard non-reshaped algorithm [18] to cluster the spectral curves, thus leading to not include spatial coordinates. Because several clusters can become empty through the optimization, we fix an initial number of clusters ($K = 150$) that ends up in providing, in average, the same resulting number of clusters than in our method (i.e. $K = 40$). Second, we implement k-means clustering, using all vector information available, that is, the input vector is build with spectral information (i.e. the energy decay curve points) concatenated to a vector of spatial information (i.e. the 3D coordinates). The number of clusters is picked to be also $K = 40$, the number of clusters being stable throughout the optimization. We use the Matlab k-means implementation for images with a reproducible initialization through the built-in ‘imsegkmeans’ function. Third and last, we implement Selective Search, a machine learning graph-based segmentation method for object recognition. Using a region merging hierarchical approach with an SVM classifier to select the hierarchical rank of the resulting regions, the authors published an open-source code [31]. We apply Selective Search on a low, an intermediate and a high energy level (resp. 40, 65 and 140keV). These energy levels are used instead of the three RGB image channels. We note that Selective Search does not predetermine the number of clusters but specifies a cluster minimum size or favors smaller or larger cluster sizes.

E. Metrics

Our clustering methods, as well as the baseline clustering algorithms are all evaluated using the following three metrics:

- 1) a cluster separation index: *Davies-Bouldin index (DB)*, computed on spatial content and on spectral content,
- 2) a clustering separation index focused only on the relationship between tumor clusters and other clusters: *Davies-Bouldin index on tumor (DB_t)*, an adaptation of DB, computed on spatial and on spectral content,
- 3) a segmentation score computed on tumor clusters versus ground truth region: *the Dice similarity score*, that can be computed only on spatial content.

To define tumor clusters, we select the cluster(s) which best cover the tumor, i.e. the ones that give the best Dice score when being merged together. Several clusters segmenting the tumor area can indeed represent different tumor subparts, but we only know the tumor primary site contour as ground truth.

When C is the ensemble of clusters, $d(\cdot, \cdot)$ is the Euclidean distance operator, $\bar{c}_k$ is the centroid of cluster c_k , the Davies-Bouldin index is defined as:

$$DB(C) = 1/|C| \sum_{c_k \in C} \max_{c_l \in C \setminus c_k} (S(c_k) + S(c_l)) / d(\bar{c}_k, \bar{c}_l) \quad (26)$$

where $S(c_k) = 1/|c_k| \sum_{x_i \in c_k} d(x_i, \bar{c}_k)$.

The ‘tumor’ Davies-Bouldin index is adapted as:

$$DB_t(C) = \max_{c_l \in C \setminus c_{tumor}} (S(c_{tumor}) + S(c_l)) / d(\bar{c}_{tumor}, \bar{c}_l), \quad (27)$$

where c_{tumor} is the region of the merged tumor clusters.

The Dice score is defined as:

$$Dice(c_{tumor}, c_{truth}) = 2 * |c_{tumor} \cap c_{truth}| / (|c_{tumor}| + |c_{truth}|), \quad (28)$$

where c_{truth} is the ground truth region. We summarize the distribution of these index values across the population by computing the mean, median and interquartile range.

V. RESULTS

Fig. 2 shows a qualitative overview of our results for one tumor example. We visualize the results on a 2D slice when the model has been run on the 3D VOI containing this slice. The top row shows the performance of the baseline algorithms, whereas the bottom row shows our proposed methods. While baseline approaches attribute a high number of clusters to bone regions containing big spectral variations and miss smaller variations in tissue regions, our methods with Gaussian gates in Fig. 2 (e, f, g) are able to adapt to relative variations and split the image with more spatial coherence. The results also demonstrate that our method is able to capture tissue characteristics invisible in Fig. 2 (b, c, d, h). Note that DB and DB_t scores depend on the number of clusters and SgMFR in Fig. 2 (h) has a very low number of clusters (softmax having vanishing clusters in the optimization). GMM and Selective Search in Fig. 2 (b) and (d) have around 40 clusters (varying number as explained in Section IV-D). The results in Fig. 2 (c, e, f, g) were obtained with 40 clusters.

Table I and Fig. 3 present the quantitative results obtained with the three clustering metrics defined in Section IV-E. Among the three proposed methods that outperform the baseline methods in terms of Dice score (SgMFR-Bspl, SgMFR-poly, and SgMVFR-Bspl), we compared the Dice score distribution obtained with SgMFR-poly (which has the lowest median Dice score among the three) to that obtained with the k-means-like baseline (which has the highest median Dice score) with a two-sample t-test, obtaining $p=0.0014$.

Fig. 4 showcases the clustering results obtained when varying the tuning of λ . Here, we understand that a smaller λ gives a higher preference to the spatial information: clusters are compact and are defining well-separated areas. On the other hand, a larger λ gives a higher priority to the spectral information: clusters more closely match tissue characteristics, but one cluster can be split into tiny voxel groups spread all over the image. The ideal λ choice would be a λ that prioritizes spectral information, but still achieves some cluster spatial compactness. We assess this through metrics calculated on spectral and spatial content as explained in Section IV-E. The general tuning of $\lambda = 0.075$ has been determined to be, on average, the optimal hyper-parameter. However, we can see

Fig. 2: Qualitative clustering results for each approach in one tumor (DB_t = spatial/spectral index). Our proposed approaches are on the bottom row. One random color is assigned per cluster, ground truth tumor contour is in blue (a) or white (b,c,d,e,f,g,h).

Fig. 3: Boxplots for each metric per method. Note that SsMFR-Bspl method gives few outliers out of reach (order of 10^{13}) on the DB spectral index, and one outlier of 135 for tumor DB spatial index. These values are shifted in the displayed range and exhibit a top arrow.

TABLE I: Average, median (interquartile range) of metrics and runtime over 81 patient scans. DB cluster separation index is computed for spatial content (spat-DB) and spectral content (spec-DB); idem for DB_t indices focused on tumor separability; runtime is given in seconds.

	GMM			k-means			S.Search			SgMFR-Bspl			SgMFR-poly			SgMVFR-Bspl			SsMFR-Bspl		
	avg	med	itq	avg	med	itq	avg	med	itq	avg	med	itq	avg	med	itq	avg	med	itq	avg	med	itq
Dice	0.17	0.15	(0.15)	0.45	0.44	(0.28)	0.35	0.28	(0.44)	0.56	0.58	(0.32)	0.57	0.56	(0.31)	0.61	0.66	(0.28)	0.20	0.15	(0.20)
spat-DB	3.32	2.69	(2.38)	6.00	5.57	(3.08)	1.30	1.28	(0.24)	1.75	1.75	(0.45)	1.74	1.72	(0.45)	1.40	1.32	(0.35)	1.84	1.22	(1.31)
spec-DB	3.17	2.80	(2.21)	8.41	7.75	(3.51)	14.36	14.36	(5.66)	7.52	6.95	(1.83)	7.66	7.13	(2.27)	8.69	8.58	(2.81)	3.78e12	2.96	(3.48)
spat- DB_t	14.97	10.95	(13.94)	5.07	3.68	(4.30)	6.57	4.07	(6.00)	2.49	2.05	(1.27)	2.59	2.00	(1.33)	2.05	1.70	(0.95)	11.07	7.12	(7.89)
spec- DB_t	1.97	1.11	(1.71)	16.52	11.42	(12.17)	28.47	19.86	(20.11)	9.64	6.78	(6.90)	9.46	7.52	(5.46)	11.57	9.53	(8.64)	2.52	1.36	(2.20)
runtime	342	165	(210)	3.64	3.78	(0.78)	0.062	0.047	(0.014)	1423	1321	(600)	902	901	(345)	437	361	(277)	2425	2121	(2049)

strong improvements in tumor separation, on a case by case basis, with small variation of λ . As shown in Fig. 4 (e) and (f), the Dice score increases from 0.36 to 0.64. This shows an example out of several results belonging to the lower quartile

in the boxplot of Dice scores in Fig. 3 that could be highly improved simply with a specific tuning. Some other examples of results in the lower quartile could be due to small tumors (size inferior to 1cm), although half of these small tumors are

Fig. 4: Clustering results for our SgMFR method with different λ tuning. One random color is assigned per cluster, ground truth tumor contour is in blue (a,d) or white (b,c,e,f).

Fig. 5: Clustering results with our SgMFR for a small tumor. Note the robustness of the result in the presence of a metallic artifact in the right hand side of the anatomical image.

actually well separated with our method, reaching a Dice score as high as 0.84 in the best case (see in Fig. 5).

VI. DISCUSSION AND CONCLUSION

In this paper, we have developed a statistical methodology to cluster intensity attenuation curves in DECT scans. We applied our proposed methods, together with other alternative clustering algorithms used as baselines, to a set of 91 DECT scans of HNSCC tumors. The classical manner of evaluating algorithms for clustering/segmentation is via measures of overlap (such as the Dice score) with a ground truth segmentation. However, as mentioned in the Introduction above, the manual segmentations of HNSCC tumors that are used as “ground truth” can suffer from large inter-rater variability, and do not incorporate in any systematic manner regions immediately adjacent to the tumor that may be biologically important for

determining the course of evolution of the tumor. Because of this inherent uncertainty in the appropriate contours of an HNSCC tumor, the main objective in our paper was to compare our clustering results to the manual contouring, but also to explore associations between voxels within the ground truth tumor contour and voxels in the surrounding tissue areas.

Compared to the baseline algorithms, it is clear both visually and quantitatively that our methods using Gaussian gates (SgM(V)FR) produce results that match better the manual segmentation contours. Our method using softmax gates (SsMFR) is less flexible compared to the one with Gaussian gating functions, and thus sometimes leads to non-satisfactory results. Although in terms of qualitative assessment clusters of SsMFR-Bspl are indeed more spatially compact, quantitative performance in some situations stays similar to GMM baseline. Thus this variant of the algorithm does not appear to perform well in practice. Using Gaussian gates, however, Dice score distributions are significantly better than with the k-means-like algorithm, the best of our baseline methods.

That being said, it is also clear that with Dice scores ranging from nearly 0 to nearly 1, our proposed methods do not recover the “ground truth” segmentations in a reliable and consistent manner. Several reasons may explain this finding. First, our clinical dataset of DECT scans is not uniform, i.e. it includes tumors of highly variable characteristics, in highly variable sizes, locations and environments, which makes it particularly challenging. Moreover, as seen in Fig. 4, changes in parameter tuning can lead to substantial improvement in Dice scores for some tumors. Finally, because of their intricate morphology and often small sizes, HNSCC tumors are inherently difficult to segment. In a recent international challenge, Dice scores of head and neck tumor segmentation ranged across different competition entries between 0.56 and 0.76 [32].

Most importantly, as argued throughout this paper, the clinical value of recovering the manual segmentations of HNSCC tumors as an objective criterion for evaluating the algorithm is also not clear. In fact, it has recently been argued in the clinical literature that AI methods in medical imaging would be more meaningful if evaluated against clinical outcomes, as opposed to an evaluation against radiologists’ performance, due to inherent subjectivity and variability of the latter [33]. For all the reasons, the objective of this paper was moved away from reproducing the manual contours produced by the radiologist, and was focused instead on developing tools that discover patterns of association in the DECT data.

Our study has several limitations. As discussed above, in our view, the appropriate way of evaluating the methodology’s clinical utility is not by computing Dice scores relative to manually drawn contours. Rather, a more clinically informative evaluation would determine the performance of the recovered clusters in predicting clinical outcome in a machine learning setting, compared to the same predictive algorithm applied with the manual tumor segmentations. Such an evaluation is missing from the present paper; it will be part of a subsequent paper in future work. Another limitation stems from the lack of an automated identification of those clusters that are associated with the tumor region. Right now, we choose those clusters that maximize overlap with the manually segmented tumor

region. Ideally, however, the abnormal tumor clusters should be identified automatically, by selecting those clusters that have the highest association with clinical outcomes. In this manner, the automated cluster identification can be naturally made part of a single machine learning pipeline for predicting clinical outcomes. Yet another limitation comes from the very small size of the subset of tumors ($N=10$) over which we estimated the algorithm parameters (λ and number of clusters), before applying the algorithm to the remaining 81 tumors in our dataset. A larger dataset, together with additional patient-specific tuning will help tune the algorithm's performance.

The need for the improvements described above is clear, and they will be made part of a subsequent publication. In the present article, we chose to focus on the theoretical and algorithmic developments. As mentioned in the Introduction, this is the first time to our knowledge that statistical tools from the Functional Data Analysis field are put into practice with DECT data. As such, the present paper remains an inherently exploratory one in its experimental framework.

Nevertheless, we believe we provide several important technical and methodological contributions. We constructed a functional regression mixture model that integrates spatial content into the mixture weights, and we developed a dedicated EM algorithm to estimate the optimal model parameters. Our mixture-based model is a highly flexible statistical approach allowing for many choices of the parametric form of the component densities. We proposed two candidate designs for the mixture weights, normalized Gaussian gates and softmax gates. The Gaussian-gate closed-form solution for spatial mixture weight updates considerably reduces the computation time while also providing solutions with better clustering index values, compared to the Newton-Raphson optimization algorithm needed at each update of the softmax-gating parameters.

REFERENCES

- [1] S. K. Vinod, M. G. Jameson, M. Min, and L. C. Holloway, "Uncertainties in volume delineation in radiation oncology: A systematic review and recommendations for future studies," *Radiother Oncol*, vol. 121, no. 2, pp. 169–179, 2016.
- [2] T. Hong, W. Tome, R. Chappell, and P. Harari, "Variations in target delineation for head and neck IMRT: An international multi-institutional study," *Int. J. Radiat. Oncol. Biol. Phys.*, vol. 60, no. 1, Supplement, pp. S157–S158, 2004.
- [3] J. van der Veen, A. Gulyban, and S. Nuyts, "Interobserver variability in delineation of target volumes in head and neck cancer," *Radiother Oncol*, vol. 137, pp. 9–15, 2019.
- [4] S. Gudi *et al.*, "Interobserver variability in the delineation of gross tumour volume and specified organs-at-risk during IMRT for head and neck cancers and the impact of FDG-PET/CT on such variability at the primary site," *J Med Imaging Radiat Sci*, vol. 48, pp. 184–192, 2017.
- [5] F. Ganci *et al.*, "Altered peritumoral microRNA expression predicts head and neck cancer patients with a high risk of recurrence," *Mod. Pathol.*, vol. 30, pp. 1387–1401, Oct 2017.
- [6] D. Aran *et al.*, "Comprehensive analysis of normal adjacent to tumor transcriptomes," *Nat. Commun.*, vol. 8, p. 1077, Oct 2017.
- [7] R. Forghani, B. D. Man, and R. Gupta, "Dual-energy computed tomography: Physical principles, approaches to scanning, usage, and implementation: Part 2," *Neuroimaging Clin N Am*, vol. 27, no. 3, pp. 385–400, Aug 2017.
- [8] P. Savadjiev, C. Reinhold, D. Martin, and R. Forghani, "Knowledge based versus data based: A historical perspective on a continuum of methodologies for medical image analysis," *Neuroimaging Clin. N. Am.*, vol. 30, no. 4, pp. 401–415, 2020.
- [9] P. Savadjiev *et al.*, "Image-based biomarkers for solid tumor quantification," *Eur Radiol*, no. 29, pp. 5431–5440, 2019.
- [10] S. Chen *et al.*, "Automatic multi-organ segmentation in dual-energy CT (DECT) with dedicated 3D fully convolutional DECT networks," *Med Phys*, vol. 47, no. 2, pp. 552–562, Feb 2020.
- [11] —, "Improving generalization capability of multi-organ segmentation models using dual-energy CT," *IEEE trans radiat plasma med sci*, 2021.
- [12] T. Wang *et al.*, "Head and neck multi-organ segmentation on dual-energy CT using dual pyramid convolutional neural networks," *Phys. Med. Biol.*, vol. 66, no. 11, 2021.
- [13] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Med Image Comput Comput Assist Interv, MICCAI 2015*, N. Navab *et al.*, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.
- [14] J.-L. Wang, J.-M. Chiou, and H.-G. Müller, "Functional data analysis," *Annu. Rev. Stat. Appl.*, vol. 3, no. 1, pp. 257–295, 2016.
- [15] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise." AAAI Press, 1996.
- [16] G. McLachlan and K. Basford, *Mixture Models: Inference and Applications to Clustering*. New York: Marcel Dekker, 1988.
- [17] C. Fraley and A. E. Raftery, "Model-based clustering, discriminant analysis, and density estimation," *JASA*, vol. 97, pp. 611–631, 2002.
- [18] G. McLachlan and D. Peel, *Finite mixture models*. N.Y.: Wiley, 2000.
- [19] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J R Stat Soc Series B Stat Methodol*, vol. 39, no. 1, pp. 1–38, 1977.
- [20] A. Samé, F. Chamroukhi, G. Govaert, and P. Aknin, "Model-based clustering and segmentation of time series with changes in regime," *Adv Data Anal Classif*, vol. 5, no. 4, pp. 301–321, Dec 2011.
- [21] K. Y. Yeung, C. Fraley, A. Murua, A. Raftery, and W. Ruzzo, "Model-based clustering and data transformations for gene expression data," *Bioinformatics*, vol. 17, no. 10, pp. 977–987, Oct 2001.
- [22] M. A. Balafar, "Spatial based expectation maximizing (EM)," *Diagnostic Pathol.*, vol. 6, no. 103, Oct 2011.
- [23] J. Vanhatalo, S. D. Foster, and G. R. Hosack, "Spatiotemporal clustering using gaussian processes embedded in a mixture model," *Environmetrics*, vol. 32, no. 7, 2021.
- [24] F. Chamroukhi and H. D. Nguyen, "Model-based clustering and classification of functional data," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, Jan. 2019.
- [25] Y. Shen, S. D. Mayhew, Z. Kourtzi, and P. Tiño, "Spatial-temporal modelling of fMRI data through spatially regularized mixture of hidden process models," *NeuroImage*, vol. 84, pp. 657–671, 2014.
- [26] R. E. Quandt and J. B. Ramsey, "Estimating mixtures of normal distributions and switching regressions," *JASA*, vol. 73, no. 364, pp. 730–738, 1978.
- [27] L. Montuelle and E. L. Pennec, "Mixture of Gaussian regressions model with logistic weights, a penalized maximum likelihood approach," *Electron. J. Stat.*, vol. 8, no. 1, pp. 1661 – 1695, 2014.
- [28] M. I. Jordan and L. Xu, "Convergence results for the EM approach to mixtures of experts architectures," *Neural Netw*, vol. 8, no. 9, pp. 1409–1431, 1995.
- [29] *AJCC Cancer Staging Manual*, 8th ed. New York: Springer International Publishing, 2017.
- [30] R. Forghani *et al.*, "Head and neck squamous cell carcinoma: Prediction of cervical lymph node metastasis by dual-energy CT texture analysis with machine learning," *Eur Radiol*, vol. 29, no. 11, pp. 6172–6181, 2019.
- [31] J. Uijlings, K. Sande, T. Gevers, and A. Smeulders, "Selective search for object recognition," *Int. J. Comput. Vis.*, vol. 104, pp. 154–171, 2013.
- [32] V. Andrearczyk *et al.*, "Overview of the HECKTOR challenge at MICCAI 2020: Automatic head and neck tumor segmentation in PET/CT," in *Head and Neck Tumor Segmentation*, 2021.
- [33] P. Savadjiev *et al.*, "Demystification of ai-driven medical image interpretation: past, present and future," *Eur Radiol*, no. 29, pp. 1616–1624, 2019.