Visualization and Exploration of Time-Varying Medical Image Data Sets

Zhe Fang
GrUVi Lab
Simon Fraser University
zfang@cs.sfu.ca

Torsten Möller
GrUVi Lab
Simon Fraser University
torsten@cs.sfu.ca

Ghassan Hamarneh
Medical Image Analysis Lab
Simon Fraser University
hamarneh@cs.sfu.ca

Anna Celler
MIRG Lab
Vancouver General Hospital
aceller@phas.ubc.ca

ABSTRACT

In this work, we propose and compare several methods for the visualization and exploration of time-varying volumetric medical images based on the temporal characteristics of the data. The principle idea is to consider a time-varying data set as a 3D array where each voxel contains a time-activity curve (TAC). We define and appraise three different TAC similarity measures. Based on these measures we introduce three methods to analyze and visualize time-varying data. The first method relates the whole data set to one template TAC and creates a 1D histogram. The second method extends the 1D histogram into a 2D histogram by taking the Euclidean distance between voxels into account. The third method does not rely on a template TAC but rather creates a 2D scatter-plot of all TAC data points via multi-dimensional scaling. These methods allow the user to specify transfer functions on the 1D and 2D histograms and on the scatter plot, respectively. We validate these methods on synthetic dynamic SPECT and PET data sets and a dynamic planar Gamma camera image of a patient. These techniques are designed to offer researchers and health care professionals a new tool to study the time-varying medical imaging data sets.

CR Categories: J.3 [Computer Application]: Life and Medical Sciences—Medical Information Systems;

Keywords: medical imaging, time-varying data, multi-dimensional scaling, volume rendering, transfer function

1 INTRODUCTION

While 3D medical imaging has become the standard in recent years for capturing structural information about the anatomy, functional imaging studies, which additionally include temporal dynamics, are still on the forefront of medical imaging research. For example, Single Photon Emission Computed Tomography (SPECT), Positron Emission Tomography (PET) and functional Magnetic Resonance Imaging (fMRI) make it possible to trace changing distributions of radiolabeled substances or examine magnetic phenomena associated with spatio-temporal brain activation. Physicians may use temporal in addition to 3-dimensional spatial information to diagnose pathological tissue based on their altered perfusion dynamics. They may track disease progress or its response to therapy over a short time or over several months or years. Further, they may examine correlations between diseases and drugs, and organ function. However, exploring and visualizing this multi-dimensional data remain challenging tasks.

For the purposes of our paper, we consider time-varying 3D medical image data (TVMID). Using data from dynamic imaging alone, there is an inherit ambiguity in deciding whether the change in spatial location of activation is due to patient motion or due to change in function of a stationary subject. Therefore, during image acquisition, patients are typically instructed to remain still or even hold their breath, whenever possible. Alternatively data from other modalities (e.g. tracking devices or structural imaging) can be acquired and used to compensate for subject motion. There are alternative approaches and a huge body of literature dealing with motion correction and image registration. In this work, we make the assumption that volumes at different times are already spatially registered to each other and do not attend to this issue any further.

One way of understanding our time-varying image data is that each voxel, with coordinates (x,y,z), does not contain a scalar value, as is common for 3D image data, but rather a temporal function, which is typically referred to as the Time-Activity Curve (TAC). The source of this temporal behaviour varies with the particular modality. For example, in blood oxygen level-dependent fMRI, the captured data reflect the metabolic processes in the brain. Specifically, in active brain regions, blood oxygenation is higher than in resting regions. Since deoxygenated haemoglobin has a stronger reaction to the magnetic field than oxygenated haemoglobin, the temporal signals in fMRI at each voxel are indicative of brain activity. In dynamic SPECT, on the other hand, the recorded TACs are based on the detection of electromagnetic radiation, whose source is an injected radioactive pharmaceutical. Using a dedicated Gamma camera, photons are detected from different angles around the patient. The 2D information about their spatial distribution and the times they were detected is recorded. In a subsequent step, a series of 3D images corresponding to different times are reconstructed yielding a time-varying 3D image. Planar Gamma images from a single angle also provide useful functional information, even without subsequent 3D reconstruction. PET is somewhat similar to SPECT. However, instead of injecting the subject with Gamma-emitting drugs, positron emitting pharmaceuticals are used. Upon annihilation with electrons, these positrons emit two 511keV photons simultaneously in two opposite directions and several rings of detectors are used to record these signals arriving in coincidence from different directions.

Different regions of the body may exhibit different radiotracer uptake and washout, resulting in different temporal behaviour and related TACs. It is usually expected that all voxels corresponding to tissues with the same physiology, performing similar function, and/or belonging to the same healthy organ will display similar TACs. Nevertheless, there can exist voxels that are spatially close to each other but with very different TACs (e.g. at boundaries of different organs). Vice versa, distant voxels can have very similar TACs (e.g. in the left and right lung). As we will explain later, our proposed techniques allow the investigation and visual exploration of these scenario.

More importantly, the altered physiology within an organ caused by disease may result in temporal behavior different from that of a healthy tissue. Recognizing abnormal behavior is important for the purpose of patient diagnosis [28]. However, it is not easy to discover such temporal changes and abnormalities in a TVMID by using conventional segmentation or visualization tools designed for scalar fields. The basic motivation of this paper is that detecting and quantifying the difference in TACs allows the segmentation, quantification, and visualization of different organs or regions of interest.

The objective of this work is to investigate new methods to visualize TVMID and build an application which allows the user to easily explore these TAC-based data. One primary purpose of functional imaging is to distinguish diseased tissue from healthy tissue. Hence, in contrast to conventional time-varying visualization, where each time step is visualized separately, we consider all time steps at once and offer a comprehensive (non-animated) visualization of the time-behaviour. Visualizing each time-step separately does not provide comprehensive information about the dynamic profiles of different tissue and thus impairs the ability of researchers and health care professionals to perform an accurate and clear diagnosis.

The general approach of our work is to define a distance metric for the space of TACs (Section 3.1). Given a particular distance metric, we propose two principle ways to explore the TVMID. The user can either select a voxel (and hence its TAC) interactively or can input an arbitrary TAC into the system. Our system will then render all TACs based on their similarity to the desired TAC (Section 3.2 and Section 3.3). As an alternative, the distances from each TAC to all other TACs in the volume can be computed. We then apply multi-dimensional scaling (MDS) to layout the high-dimensional graph of all TACs on a 2D palette which forms the basis of interactive exploration of the underlying TVMID (Section 3.4).

Our main contributions in this paper are: (1) The interactive query of a particular dynamic tissue behaviour (tracer concentration or tissue magnetization). (2) The assessment of different distance metrics for the discrimination of temporal tissue behavior. (3) The assessment of the influence of a spatial relationship between TACs for the determination of different tissue behavior. (4) Development of comprehensive methods for the visualization of the temporal behaviour of medical data.

These visualization methods can be applied to most time-varying medical image data sets: such as dynamic PET, dynamic SPECT, planar dynamic single photon studies and fMRI. We show the effectiveness of our visualization methods on dynamic SPECT and PET data sets created by using physical or computational phantoms as well as on a dynamic planar Gamma camera image of a patient. We believe that visualizing the TVMID by our new methods is intuitive, provides new insight into these types of data sets and, consequently, helps improve diagnosis and treatment evaluation.

2 Related Work

The rendering community thus far has mostly focused on the technical aspects (accuracy and speed) of rendering very large data sets. This problem has been elevated by time-varying data, since many data structures cannot easily be updated from one time frame to the next in real-time. Several data structures have been proposed to alleviate this problem like the T-BON [27], time-space partition trees [25], differential time-histogram tables [34] and others. Ma et al. [16] give a survey of different time-varying data visualization techniques. The motivation of most of that research lies behind dealing with massive data volumes, common in engineering problems (see e.g. the Richtmeyer-Meshkov data [19]).

However, the goal of medical diagnosis is not to observe a dynamic behaviour in real-time as much as to differentiate healthy from diseased tissue. Hence, it is not so important to provide interactive selection of arbitrary time-slices, but rather the comprehensive evaluation of the time-behaviour of each tissue sample and its classification into different tissue types and tissue behaviours. Therefore, novel transfer function designs are needed, that take the whole data set into account.

2.1 Time-Varying Volume Data Visualization

Recently, several approaches have been proposed to extend the aspect of transfer function design to time-varying data in a comprehensive way [3, 8]. A time histogram is given by concatenating a series of conventional 1D histograms together for each step. Hiroshi et al. [3] present a visualization method based on the time histogram, which allows simultaneous classification of the entire time series. In their system, the user can specify transfer functions either directly or indirectly using the time histogram and achieve interactive feedback.

The approach taken in this paper is to treat each voxel as a TAC, as opposed to reducing it to a scalar. We then cluster these functions in groups with similar behaviour, creating a 1D or 2D transfer function. The user then specifies the regions of interest using an interface similar to that described by Doleish et al. [9]. However, in our case the 2D space in which we specify the regions of interest is obtained by using a clustered representation provided via MDS.

MDS is suited for this task and is well studied in database applications and visualization [4, 21, 22]. However, the size of our data is unprecedented for such algorithms (our data may have more than one million data points). A fast and interactive MDS system would facilitate interactive data exploration by the user. Hence, we employ the MDSteer system [32], a steerable mass-spring model MDS system.

Tory et al. discuss several methods for visualizing time-varying medical data [29]. They explore visualization methods, which include iso-surfaces, direct volume rendering and vector visualization using glyphs, to produce an animation of consecutive time steps that illustrate how the intensity and gradient change with time. They argue that these visualization techniques could provide new insight into the information contained in the TVMID and therefore help improve diagnosis or treatment. However, their system is not capable of classifying data into different tissue types.

2.2 Segmentation of time-varying medical data

Our proposed framework allows for identification and visualization of regions with particular dynamic behaviours. Several techniques have been proposed in the past for the segmentation of time-varying medical image data. What is characteristic about these techniques is that once the time-varying image data, along with some initial parameters, are set, the user has no influence over the segmentation outcome. Our framework, on the other hand, provides the user with the means to intervene, by leveraging their domain specific expertise. This is accomplished via an intuitive graphical interface, while providing immediate visual feedback of alternative segmentation results. This keeps the clinical expert in-the-loop, which allows for producing validated and useful results.

A majority of previous segmentation techniques are based on clustering [1]. In [33], a k-means like algorithm is used to cluster TACs in PET images, whereas in [2] fuzzy c-means (FCM) clustering is used to segment SPECT images. Velamuru et al. [31] used hierarchical clustering as well as inter- and intra-cluster measures to assess tissue activity. In [11], pre-clustering and hierarchical clustering are used with extracted kinetic parameters to label PET data.

Our approach is based on reducing the dimensionality of the data by embedding the TACs of all the voxels in a two dimensional space. The separation between the TAC observations in 2D is proportional to TAC dissimilarity or spatial distance between TAC sources. Effectively, this iteratively groups the TAC observations while continuously producing an update of the 2D view. Toennies et al. [28] discuss a dynamic SPECT segmentation scheme for aggregating voxels with similar TACs by using principal component analysis (PCA). The similarity of voxels is measured on the most significant components of the transformation of TACs in eigenspace, in order to bring the signal-to-noise ratio to a level which would

improve the diagnostic potential. In [13], FCM is compared to other techniques based on expectation-maximization (EM) and independent component analysis (ICA). The use of ICA is somewhat related to our approach as both perform dimensionality reduction. Other previous work incorporated spatial information in clustering PET data through Markov Random Fields (MRF) [7, 15]. In our framework, the geometric distance between TACs can also be taken into account. During the dimensionality reduction and clustering, the user can visualize, on-the-fly, the results of interactively changing the degree of influence of spatial TAC proximity on the final segmentation.

Other approaches for segmenting time-varying medical images are based on energy minimization of spatio-temporal explicit, parametric deformable models and are typically applied to dynamic cardiac data [17, 20] and tagged cardiac MRI data [18]. Implicit, level-set based deformable models have also been proposed more recently to segment time-varying medical image data [12]. Deformable registration-based segmentation has also been proposed to address the segmentation of dynamic image data. Here, a labelled initial 3D volume is provided and then non-rigidly deformed to subsequent time frames in order to transfer the labels throughout the sequence of volumes [6]. In a more recent work, Song et al. [26], combine segmentation and registration for dynamic renal MR images. These deformable shape model based techniques have been generally designed for modalities that exhibit a non-rigidly varying shape with time, rather than stationary anatomy that exhibits dynamic activity, which is the main focus of our work. Analyzing functional MRI data, provides a somewhat unique perspective. The data is typically reduced to scalar statistical parametric maps based on confidence values of activation and then simple thresholding is applied to highlight regions of interest [10].

3 METHODS

In order to understand, analyze, and render the TAC volume, we need to define operations on these TACs that will enable us to use standard volume rendering tools common for 3D scalar data. Additionally, in order to be able to assign colour and opacity to a particular TAC, we need to adapt classical scalar field based transfer functions to time-varying data. One of the key ideas of transfer function design is to select a particular feature (i.e. iso-value for scalar volumes) and assign it a particular colour and opacity. Then, typically, voxels with similar feature values are assigned a slightly lower opacity (see Levoy [14]). The same principle can be applied to non-scalar data, including TAC-based data, as long as the notion of *similarity* is well defined. Hence, defining a distance metric among TACs is key.

3.1 Similarity measures for TACs

For traditional (static) 3D data, $V_s : \mathbb{R}^3 \rightarrow \mathbb{R}$ is a scalar volume. However, in this paper we deal with a dynamic volume $V_d : \mathbb{R}^3 \rightarrow C^0(\mathbb{R})$. Here $C^0(\mathbb{R})$ is the space of all continuous functions. This differs from the traditional approach $V_t : \mathbb{R}^4 \rightarrow \mathbb{R}$. The representation of V_d is not unreasonable, since the change of activity concentration observed in the tissue will be gradual. The possibility for distance measures on the spaces $\mathbb{R}$ and $C^0(\mathbb{R})$ include the L_p metric, of which we will use the L_1 and L_2 metrics:

$$d_1(v_i, v_j) = \int_{-\infty}^{\infty} |v_i(t) - v_j(t)| dt$$

$$d_2(v_i, v_j) = \sqrt{\int_{-\infty}^{\infty} (v_i(t) - v_j(t))^2 dt}$$

where $v_i, v_j \in V_d$, i.e. they are image TACs.

Assuming we have N time samples, the discretized version of L_1 and L_2 (denoted by l_1 and l_2) are given by

$$d_1(v_i, v_j) = \frac{1}{N} \sum_{n=1}^{N} |v_i[t_n] - v_j[t_n]| \quad (1)$$

$$d_2(v_i, v_j) = \sqrt{\frac{1}{N} \sum_{n=1}^{N} (v_i[t_n] - v_j[t_n])^2} \quad (2)$$

Further, we consider the cross-correlation between two signals, which is given by:

$$v_i \star v_j(t) = \int_{-\infty}^{\infty} v_i(\tau) v_j(\tau + t) d\tau$$

For the purposes of defining a *similarity* measure, we consider the maximum cross-correlation, i.e.

$$s_C(v_i, v_j) = \max_t \int_{-\infty}^{\infty} v_i(\tau) v_j(\tau + t) d\tau$$

Again, in its discretized version we have:

$$s_C(v_i, v_j) = \max_k (\frac{1}{N} \sum_{n=1}^{N} v_i[t_n] v_j[t_n + k]) \quad (3)$$

where $k \in \mathscr{Z}$ and $v[t_k] = 0$ for $k \notin [0, N]$.

Last, but not least, the Euclidean distance between voxels will also play an important role in the classification of the tissue. Hence, in a slight abuse of our notation, we will also define the Euclidean metric as:

$$d_E(v_i, v_j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2} \quad (4)$$

where (x_i, y_i, z_i) and (x_j, y_j, z_j) are the spatial locations of v_i and v_j respectively.

3.2 Template TAC

As in scalar volume visualization, where features are typically specified as a particular iso-surface value, we can similarly determine a particular template TAC. Such a template TAC can be determined by interactively probing the underlying data by specifying an arbitrary voxel via its spatial location (x_k, y_k, z_k) or by simply sketching a template TAC. Often the clinician knows or expects a particular temporal behaviour from a particular tissue. Hence the a priori specification of a template TAC is a valuable feature of our system. The template TAC can also be generated from the result of some previous analysis, such as finding an average of a region of interest, principal component analysis (PCA) or nonlinear dimensionality reduction.

Figure 1 shows the widget which we created for the inspection and modification of TACs. Figure 1a shows the volume-rendered image of some classification of a phantom data set. Figure 1b displays three different TACs associated with different tissue types: lungs (blue), heart (green) and liver (red). Results in Figure 1 were obtained by using the d_1 distance.

After a particular template TAC is specified, we can compute the distance of the TACs at all other voxels to this template TAC, which results in a scalar volume (i.e. each voxel is associated with one scalar distance). This volume can now be rendered using classical (scalar) volume rendering algorithms. For example, Figure 2 shows the histogram and template TAC determined for a synthetic data set. Using standard transfer functions, we can now render the underlying data set.

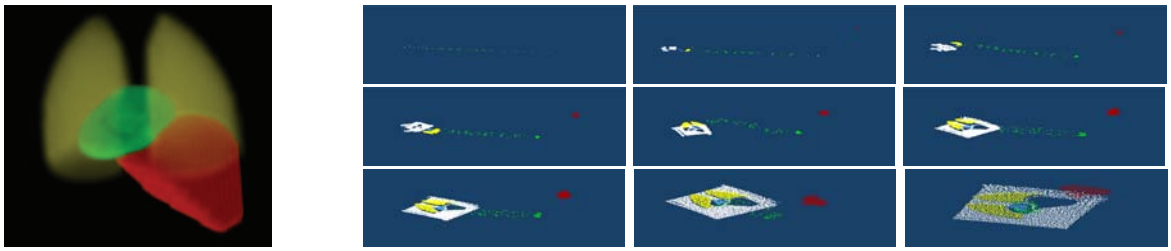

(a) (b)

Figure 6: The MDS of TACs obtained by assigning different weight. (a) The synthetic data set used for this experiment. (b) The result of the MDS layout using different weights (from left to right): (top row) $1d_1+0d_E, 0.9d_1+0.1d_E, 0.8d_1+0.2d_E$; (middle row) $0.7d_1+0.3d_E, 0.6d_1+0.4d_E, 0.5d_1+0.5d_E$; (bottom row) $0.4d_1+0.6d_E, 0.3d_1+0.7d_E, 0.2d_1+0.8d_E$;

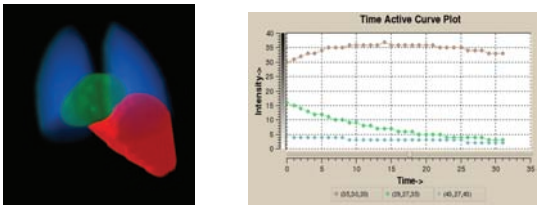

(a) (b)

Figure 1: The time activity curve visualization and design interface. (a) Volume rendering of a distance volume obtained by using the d_1 distance between each TAC and the template TAC. The template TAC used was a TAC from the heart (green) (b) Representative TACs corresponding to the lungs (blue), heart (green), and liver (red).

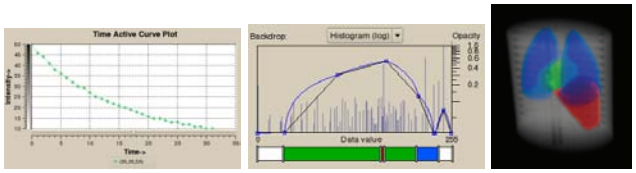

(a) (b) (c)

Figure 2: Transfer function assignment and volume rendering of distance volume to a template TAC. (a) The template TAC. (b) The histogram of the distance volume (using d_1) and an associated transfer function. (c) Volume rendering of the distance volume.

3.3 Dissimilarity and Euclidean Distance to the Template TAC

While discriminating simply on the basis of the TAC distance can be quite effective, we may also wish to distinguish between body parts of functional regions, that have similar TACs, but are separated from each other (e.g. the right and the left lung). In order to do so, we also include the Euclidean distance (d_E) and create a 2D histogram. Here, d_E is depicted along the x-axis and the TAC distance (d_1, d_2 or s_C) along the y-axis. Hence, we can specify colour and opacity using a 2D histogram (Figure 3). The user simply outlines a region over the data of interest. A colour and opacity value can then be assigned to this region, which will be inherited by all data points included in this region.

As one can see in Figure 4, we are not able to separate the right from left lung. This is due to the fact that the orthogonal separation of the Euclidean distance and TAC distance is not sufficient. Hence, we need an alternative way to visualize the data.

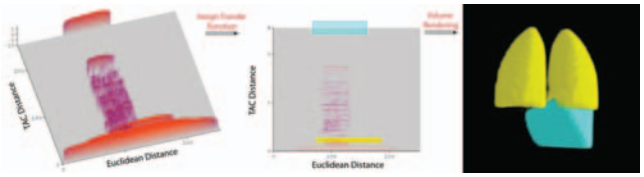

Figure 3: A 2D histogram formed by d_E and TAC distance d_1. A transfer function can be assigned directly to the 2D histogram in order to highlight the region of interest.

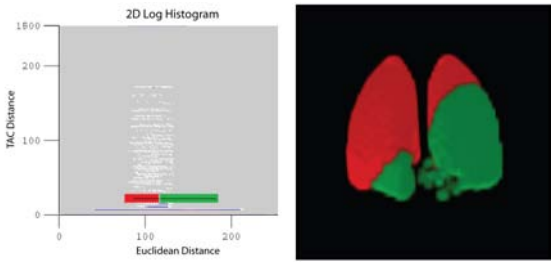

Figure 4: Transfer function assignment (left) on TAC and Euclidean distance 2D histogram fails to segment the left and right lungs (right).

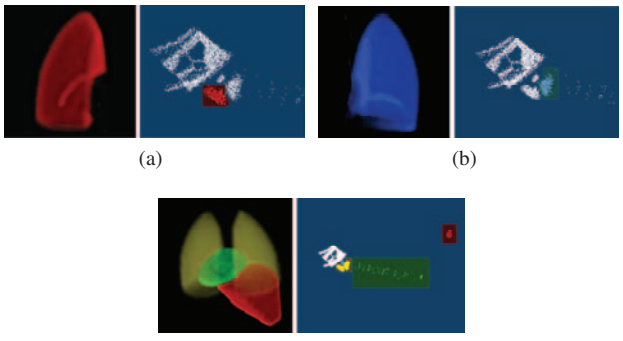

(a) (b)

(c)

Figure 5: The MDS layout using a similarity distance of $0.67d_1+0.33d_E$ and a transfer function assignment on the right with the corresponding volume rendering on the left. (a) The left lung segmentation (b) The right lung segmentation (c) a complete segmentation of the dynamic SPECT data set.

3.4 Layout by Multi-dimensional Scaling

In scalar volume rendering one often uses the histogram (or modified histograms [5, 23]) as an aid to determine features in the dataset. This approach is not straightforward for TAC-based volumes. Assuming no template TAC is available, we simply compute a distance from each TAC in the volume to all of the other TACs. This gives an $n \times n$ distance matrix, where n is the size of the data set (i.e number of spatial voxels). This distance matrix is the input to a standard MDS algorithm. MDS embeds the space of all TACs into lower dimensions, 2D in our case, such that the distance distortions are minimized. Given this 2D embedding we can now use it in order to specify a 2D colour and opacity transfer function. Here we use the same transfer function specification method as for the 2D histograms in Section 3.3. Figure 5 illustrates the transfer function assignment and corresponding volume rendering.

While we also could have chosen an embedding into 1D or 3D space, we concluded that a 1D embedding would sacrifice the separability of the clusters too much and that the specification of a transfer function in 3D would be too difficult. Specifying the transfer function in 2D is still manageable and can be provided through an intuitive interface.

Revisiting the problem of differentiating body parts that are separated in space but have similar tracer behaviour, we assign a weighted TAC distance, that balances Euclidean and TAC distances:

$$d(v_i, v_j) = (1-\alpha)d_E(v_i, v_j) + \alpha d_p(v_i, v_j) \quad (5)$$

where $p \in \{1,2\}$ and $\alpha \in [0,1]$. One drawback of the similarity measure s_C is that it cannot be applied in such a setting directly, since there is no obvious calibration with the Euclidean distance. Whereas d_1 and d_2 will always be zero for two equal TACs, s_C can assume any arbitrary value.

Here we use α in order to weight between Euclidean distance and TAC distance. Both these distances are normalized between zero and one. Figure 6 shows the influence of the Euclidean distance on the layout by assigning different weights. In order to assess the impact of this weighting, we colour coded the layout according to a pre-classified synthetic data set, as shown in Figure 6a. However, normally the assignment of colour and opacity happens after the initial 2D embedding using MDS, which is illustrated by Figure 5.

4 Results

4.1 System environment

In order to implement our ideas we use the vuVolume framework. This framework has been developed in our lab and is available on sourceforge for download[1]. Since vuVolume focuses on volume rendering of scalar data, we extended its capabilities to deal with time-varying data.

The basic environment can be seen in Figure 8. First we load the data using the widget (Figure 8(1)). Then the user determines a template TAC by specifying its (x,y,z) location (Figure 8(2)). A white marker appears in the volume in order to guide the user's selection (see the render window in Figure 8(3)). Further, the TAC will be displayed in the TAC widget Figure 8(4). Next, the user can pick one of the three distance metrics (d_1, d_2 or s_C) (Figure 8(5)). After the distance volume is computed, the histogram is being displayed in the histogram widget (Figure 8(6)). Now the user can interactively explore the scalar volume in real time, similar to several other volume rendering environments. Instead of using the traditional 1D histogram-based transfer functions, the user can switch to a 2D histogram, spanned by the Euclidean distance and the chosen TAC distances by picking the proper tab (Figure 8(7)). Other GUI elements have been explained in Figure 3.

Alternatively, when no template TAC is chosen, the user still has to pick a similarity measure before switching to the MDS view using the widget Figure 8(7).

[1] http://sourceforge.net/projects/vuvolume/

4.2 Assessment of the different TAC metrics

In the following experiments we are assessing the strengths and weaknesses of different TAC similarity measures with noise interference. In order to do so, we generate a template TAC $v(t)$ with log-normal distribution behaviour $v(t) = 4e^{-8(\ln t)^2}/t\sqrt{2\pi}$. We add uniform noise $n(t)$ to this synthetic template TAC $v(t)$ in order to judge the robustness of the similarity measure. Hence, we are producing a signal $v_n(t) = v(t) + \beta/100 * n(t)$, where we are assuming, that $n(t) \in [-1,1]$. In Figure 7(a) we plot the similarity measure $d(v(t), v_n(t))$ as β is changing between 0 and 100 (where d stands for d_1, d_2 or s_C). While both, d_1 and d_2 will start at point $(0,0)$, the maximum correlation s_C naturally is not zero. To remove the effect of the value of the maximum correlation of $v(t)$, we normalize this curve by $s_C(v(0), v(0))$. Further, to make the trend clearer, we do a least-squares fit of all three curves, and compute their slopes. The result of this test is displayed in Figure 7(b).

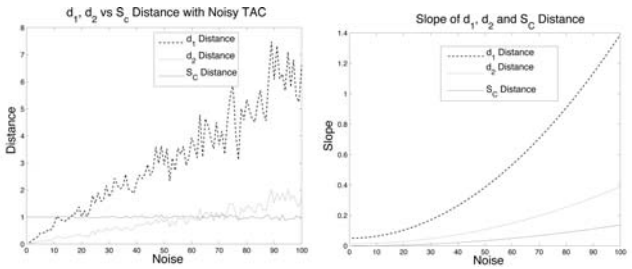

(a) (b)

Figure 7: Given template TAC $v(t) = \frac{4e^{-8(\ln t)^2}}{t\sqrt{2\pi}}$ (a) d_1, d_2 and s_C distance between noise TACs and template TAC (b) The slope of d_1, d_2 and s_C distance between noise TACs and template TAC

The measure s_C is clearly preferable over d_1 and d_2 as it is more resilient to increasing levels of noise (lower slope values in Figure 7(b)). Further, d_1 is more sensitive to uniform noise compared to d_2.

4.3 Phantom and patient data

Our visualizations and studies presented earlier in this paper were all based on a dynamic version of the MCAT phantom [30]. The adaptation of the static MCAT phantom to a dynamic setting was done by the Medical Imaging Research Group (MIRG) at the Vancouver General Hospital. The dynamic data set, which we will denote dMCAT, corresponds to a series of $64 \times 64 \times 64$ matrices spanning over 32 time frames. This phantom models a myocardial perfusion SPECT study. It reproduces a distribution of the dynamic radiotracer teboroxime in the heart and other organs of the thorax.

The challenge in SPECT heart studies is to differentiate photons that originate from the liver from those originating from the heart. Scattered liver photons are often mis-identified as coming from the heart, which may conceal myocardial perfusion defects. Since the dynamic behaviour of the tracer in the liver and the heart are different, their respective TACs can be used to differentiate between the two groups of photons. Figure 9 shows the best separation of liver and heart for the three similarity measures d_1, d_2, and s_C. The

initial template TAC was chosen to be a simple flat zero TAC created from a voxel outside any region of interest. This resulted in an easy distinction between heart and liver TACs for the d_1 and d_2 distances. The constant zero TAC is not usable for the s_C distance since the correlation of any TAC to a zero TAC will result in a zero value. Therefore, an alternative template TAC must be selected. We used a TAC extracted from the heart region.

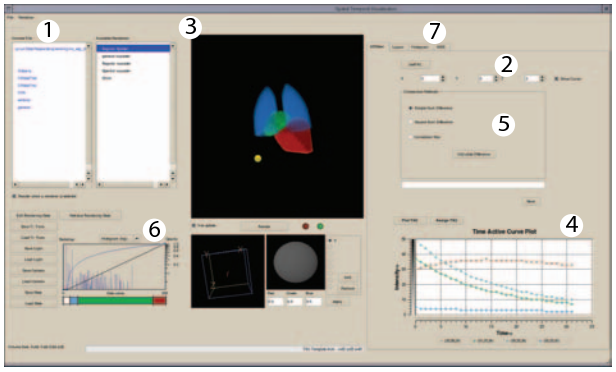

Figure 8: Screen-capture of the TVMID visualization and exploration application. The GUI includes panes for volume rendering (3), viewing TACs (4), 1D histograms and specifying a 1D transfer function (6), 2D histograms and a 2D transfer function (under a different tab (7)), MDS layout (7), choice of TAC dissimilarity metric (5) and more. The left half part is for loading a TVMID, specifying a transfer function and performing volume rendering and the right half is where the user can choose one of the different methods which can be used to analyze data sets.

Methods	Heart	Liver	Lung
original	1947 (voxels)	7148 (voxels)	17544 (voxels)
d_1 distance	100%	100%	100%
d_2 distance	100%	100%	100%
s_C distance	96.8%	100%	98.8%

Table 1: Comparing the performance of segmentation of the dynamic MCAT data with different similarity measures. The table depicts the percentage of correctly labeled voxels.

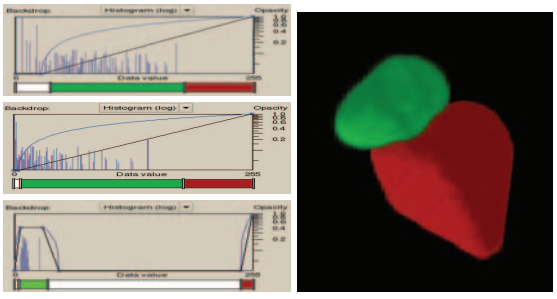

(a) (b)

Figure 9: Using different similarity measures for the separation of heart, liver, and lung tissue (Table 1). (a) the distance histogram and the assigned transfer function. (b) volume rendered image of the heart (green) and liver (red). Since all three rendered images look very similar, we only show one representative. (a-top) d_1 distance and a zero TAC as template. (a-middle) d_2 distance and a zero TAC as template. (a-bottom) s_C distance and a heart TAC as template.

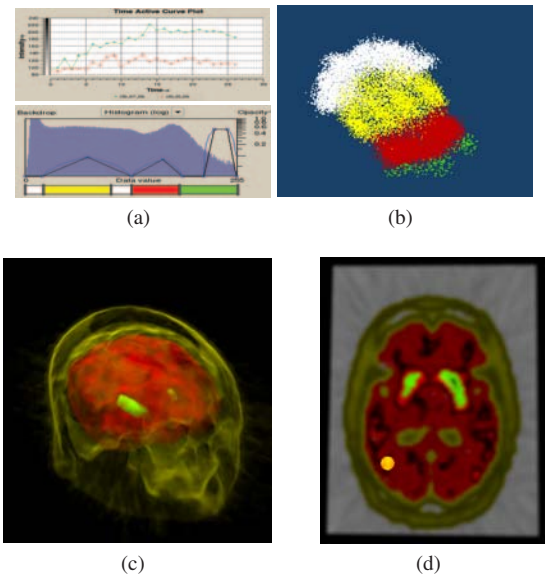

(a) (b) (c) (d)

Figure 10: Simulated dynamic PET data set (based on PET-SORTEO): (a-top) The TAC of a voxel in the striatum (green) and the template TAC in the brain area (red). (a-bottom) The histogram (based on the red template TAC) and the transfer function used to render the image in (c). (b) The MDS layout using a similarity distance proportional to $0.5d_1 + 0.5d_E$. (c) Volume rendering based on the transfer function in (a-bottom) using a template TAC in the brain and the d_1 distance. (d) The position (yellow dot) of the template TAC.

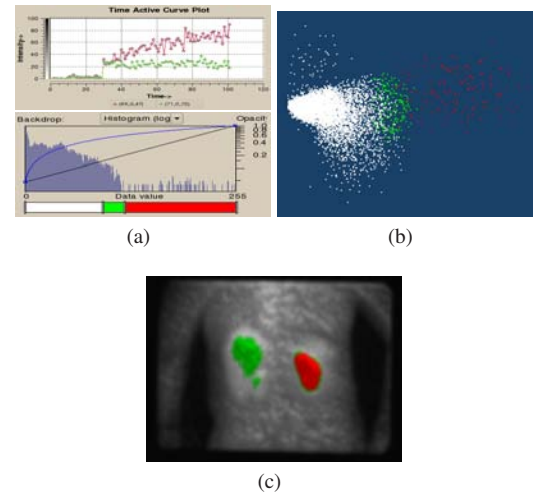

(a) (b) (c)

Figure 11: Data set of a dynamic nuclear medicine (single photon) study: (a-top) The TAC of a voxel in the healthy kidney area (red) and abnormal area (green). (a-bottom) The transfer function used to render the image in (c). (b) The MDS layout using a similarity distance proportional to $1d_1 + 0d_E$ (c) display of the slice (not rendered) with the colour assignment set in part (a-bottom).

Since dMCAT represents a synthetic data set, we know exactly which voxels correspond to the liver tissue and which to the heart tissue. Table 1 presents the accuracy of our segmentation using our three different similarity measures. Since this data set does not contain any noise, we can see that the d_1 and d_2 similarity measures perform best and the noise robustness of s_C is not being exploited.

The second type of data we have worked with was a simulation of a dynamic PET brain study from PET-SORTEO [24]. The data set that we used has $128 \times 128 \times 63$ spatial voxels and 27 time frames. In clinical practice the striatum and cerebellum are typically regions of interest which one would want to analyze and visualize [13]. Hence, we select a template TAC, corresponding to a voxel which belongs to the white matter of the brain (Figure 10d). We calculate the d_1 distance between the template TAC and all other TACs; the resulting distance volume is visualized by volume rendering (Figure 10c). The striatum shows a different d_1 distance and is highlighted (green) inside the brain (red) in Figure 10c. A TAC (green) from the striatum is extracted and displayed in Figure 10a in order to compare it with a TAC (red) from the brain. Figure 10b shows the MDS layout of TACs with a combined similarity measure of $0.5d_1 + 0.5d_E$ respectively. The points are colour coded to match the colour of the volume rendering. The points in green are associated with TACs in the striatum area; points in red are associated with TACs in the brain area; points in yellow are associated with TACs in the the skin, while background TACs are white.

The third type of data we have worked with was a dynamic series of planar images corresponding to a renal patient study. The planar image has $128 \times 128 \times 1$ spatial voxels and 102 time frames. Physicians use these dynamic planar images to evaluate the kidney function. In this case a zero template TAC was used. We first calculate the d_1 distance between the template TAC and all other TACs, which is then displayed in the histogram at the bottom of Figure 11a. We display the resulting 2D image in Figure 11c. This image was created with a transfer function presented in the bottom of Figure 11a. Figure 11b shows the MDS layout of TACs by using a similarity distance proportional to $1d_1 + 0d_E$. Figure 11c clearly shows that one of this patient's kidneys (red) is healthy as it has normal TAC. The other kidney (green) displays abnormal temporal behaviour (as can be seen in Figure 11a and b) which is related to the fact that its function is altered by some disease. This example shows how our method can potentially help in clinical diagnosis by identifying organs or areas which are characterized by TAC's different from that of a normal tissue.

5 Conclusions

We proposed several new analysis and visualization methods that are designed to help the user effectively explore time-varying medical image data sets. All our methods are based on the analysis of temporal information (TAC) associated with the dynamic data in contrast to previous visualization methods [29]. We define three similarity measures between different TACs and base our visualizations on these. We find that the maximum correlation is robust under noise, which is important for nuclear medical imaging data which often have low statistics thus high noise. However, this similarity metric is not well suited for a combination with the Euclidean distance. We find the Euclidean distance to be an important instrument to differentiate between tissues with the same temporal behaviour but belonging to different, spatially separated organs.

We proposed three principle approaches for the exploration of time-varying medical image data: using 1D TAC distance histogram; 2D geometric vs TAC distance histogram; and the 2D layout of spring mass based nonlinear dimensionality reduction. The first two methods are fast and intuitive. However, they require the user to have some previous knowledge about the underlining data in order to choose the proper template TAC. The third method does not require such a template TAC and provides an overview of the complete data set. It takes, however, longer to create such a layout. We use multi-dimensional scaling to layout the whole data set as a backdrop for specifying a 2D transfer function. Our interface provides flexibility to the user and can be a powerful tool for gaining an understanding of time-varying data.

6 Future Work

The purpose of the tool we have developed is not clinical use because in a typical clinical department there is little time to specify transfer functions etc. Therefore, our tool is rather intended for use by a researcher, a medical physicist for example, who is involved in development of new dynamic reconstruction and analysis algorithms or investigation of new tracers. While the ideas presented in this paper have been developed and tested in collaboration with such a physicist, at this point we have only applied them to a limited set of simulated data created for research purposes and only a few patient studies. At the next stage, however, we plan to extend these tests on several different types of clinical data and to involve our medical collaborators in the assessment and optimization of the method.

Further, it takes presently too long to compute the MDS layout (several hours for a 64^3 data set). We are actively collaborating with the creators of MDSteer [32] to incorporate into our program the next version of their system that is faster by several orders of magnitudes from the current version.

References

[1] P Acton, L Pilowsky, D Costa, and P Ell. Multivariate cluster analysis of dynamic iodine-123 iodobenzamide SPET dopamine D2 receptor images in schizophrenia. In *European journal of nuclear medicine*, pages 111–118, 1997.

[2] P Acton, L Pilowsky, H Kung, and P Ell. Automatic segmentation of dynamic neuroreceptor single-photon emission tomography images using fuzzy clustering. In *European journal of nuclear medicine*, pages 581–90, 1999.

[3] Hiroshi Akiba, Nathaniel Fout, and Kwan-Liu Ma. Simultaneous classification of time-varying volume data based on the time histogram. In *Proc. of EuroVis 2006*, 2006.

[4] Nina Amenta and Jeff Klingner. Case study: Visualizing sets of evolutionary trees. In *IEEE Symposium on Information Visualization*, pages 71–74, 2002.

[5] Chandrajit L. Bajaj, Valerio Pascucci, and Daniel R. Schikore. The contour spectrum. In *Proceedings of the 8th IEEE Conference on Visualization '97*, pages 167–175. IEEE Computer Society Press, 1997.

[6] Raghavendra Chandrashekara, Daniel Rueckert, and Raad Mohiaddin. Cardiac motion tracking in tagged MR images using a 4D B-spline motion model and nonrigid image registration. In *ISBI*, pages 468–471, 2004.

[7] Jun Chen, Steve Gunn, Mark Nixon, Roger Gunn, Michael Insana, and Richard Leahy. Markov random field models for segmentation of PET images. In *information processing in medical imaging*, pages 81–92, 2001.

[8] H. Doleisch, M. Mayer, M. Gasser, R. Wanker, and H. Hauser. Case study: Visual analysis of complex, time-dependent simulation results of a diesel exhaust system. In *In Proc. of the 6th Joint IEEE TCVG - EUROGRAPHICS Symposium on Visualization*, pages 91–96, 2004.

[9] Helmut Doleisch, Martin Gasser, and Helwig Hauser. Interactive feature specification for focus+context visualization of complex simulation data. In *VISSYM '03: Proceedings of the symposium on Data visualisation 2003*, pages 239–248, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[10] K Friston, A Holmes, K Worsley, J Poline, C Frith, and R Frackowiak. Statistical parametric maps in functional imaging: A general linear approach. In *Human Brain Mapping*, pages 189–210, 1995.

[11] Hongbin Guo, Rosemary Renaut, Kewei Chen, and Eric Reiman. Clustering huge data sets for parametric PET imaging. In *Biosystems*, pages 81–92, 2003.

[12] Timo Kohlberger, Daniel Cremers, Mikaël Rousson, Ramamani Ramaraj, and Gareth Funka-Lea. 4D shape priors for a level set segmentation of the left myocardium in SPECT sequences. In *MICCAI (1)*, pages 92–100, 2006.

[13] Heidi Koivistoinen, Jussi Tohka, and Ulla Ruotsalainen. Comparison of pattern classification methods in segmentation of dynamic PET brain images. In *Signal Processing Symposium*, pages 73–76, 2004.

[14] Marc Levoy. Display of surfaces from volume data. In *IEEE Computer Graphics and Applications*, 1988.

[15] Kang-Ping Lin, Shyhliang A. Lou, Chin-Lung Yu, Being-Tau Chung, Liang-Chi Wu, and Ren-Shyan Liu. Markov random field method for dynamic PET image segmentation. *Medical Imaging 1998: Image Processing*, 3338(1):1198–1204, 1998.

[16] Kwan Liu Ma and Eric B. Lum. Techniques for visualizing time-varying volume data. In *The Visualization Handbook*, pages 511–531, 2005.

[17] T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis. *IEEE Transactions on Computerized Medical Imaging and Graphics*, pages 69–83, 1994.

[18] D. Metaxas, T. Chen, X. Huang, and L. Axel. Cardiac segmentation from MRI-tagged and CT images. In *8th WSEAS International Conf. on Computers, special session on Imaging and Image Processing of Dynamic Processes in biology and medicine*, 2004.

[19] A. A. Mirin, R. H. Cohen, B. C. Curtis, W. P. Dannevik, A. M. Dimits, M. A. Duchauneau, D. E. Eliason, D. R. Schikore, S. E. Anderson, D. H. Porter, P. R. Woodward, L. J. Shieh, and S. W. White. Very High Resolution Simulation of Compressible Turbulence on the IBM-SP System. In *ACM/IEEE Supercomputing 1999 Conference (SC'99)*, page 70, Los Alamitos, CA, USA, 1999. IEEE Computer Society.

[20] Johan Montagnat and Herve Delingette. Space and time shape constrained deformable surfaces for 4D medical image segmentation. In *MICCAI '00: Proceedings of the Third International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 196–205, London, UK, 2000. Springer-Verlag.

[21] A. Morrison, G. Ross, and M. Chalmers. A hybrid layout algorithm for sub-quadratic multidimensional scaling. In *Proc. IEEE Information Visualization*, pages 152–158, 2002.

[22] Alistair Morrison, Greg Ross, and Matthew Chalmers. Fast multidimensional scaling through sampling. In *Proc. IEEE Information Visualization*, pages 68–77, 2003.

[23] Vladimir Pekar, Rafael Wiemker, and Daniel Hempel. Fast detection of meaningful isosurfaces for volume data visualization. In *Proceedings of the IEEE Conference on Visualization '01*, pages 223–230. IEEE Computer Society, 2001.

[24] A. Reilhac, G. Batan, C. Michel, C. Grova, J. Tohka, N. Costes, and A. C. Evans. PET-SORTEO: validation and development of database of simulated PET volumes. In *IEEE Trans. Nucl. Sci., 2004*, pages 1321– 1328, 2004.

[25] Han-Wei Shen, Ling-Jan Chiang, and Kwan-Liu Ma. A fast volume rendering algorithm for time-varying fields using a time-space partitioning (TSP) tree. In *IEEE Visualization '99*, pages 371–378, 1999.

[26] Ting Song, Vivian S. Lee, Henry Rusinek, Samson Wong, and Andrew F. Laine. Integrated four dimensional registration and segmentation of dynamic renal MR images. In *MICCAI (2)*, pages 758–765, 2006.

[27] Philip Sutton and Charles D. Hansen. Isosurface extraction in time-varying fields using a temporal branch-on-need tree (T-BON). In *Proceedings of IEEE Visualization '99*, pages 147–153, 1999.

[28] Klaus D. Toennies, Anna Celler, Stephan Blinder, and Ronald Harrop Torsten Möller. Scatter segmentation in dynamic SPECT images using principal component analysis. In *Proc. of SPIE International Symposium on Medical Imaging*, pages 507–516, 2003.

[29] Melanie Tory, Niklas Röber, Torsten Möller, Anna Celler, and M. Stella Atkins. 4D space-time techniques: A medical imaging case study. In *Proc. IEEE Visualization 2001*, pages 473–476, 2001.

[30] B.M. Tsui, J.A. Terry, and G.T. Gullberg. Evaluation of cardiac cone-beam SPECT using observer performance experiments and ROC analysis. *Investigative Radiology*, 28(12):1101–1112, July 1993.

[31] P. K. Velamuru, R. A. Renaut, H. Guo, and K. Chen. Robust clustering of positron emission tomography data. In *Joint Interfce CSNA*, 2005.

[32] Matt Williams and Tamara Munzner. Steerable, progressive multidimensional scaling. In *Proc. IEEE Information Visualization 2004*, pages 57–64, 2004.

[33] K Wong, D Feng, S Meikle, and M Fulham. Segmentation of dynamoic PET images using cluster analysis. In *IEEE Transactions on Nuclear Science*, pages 200–207, 2002.

[34] Hamid Younesy, Torsten Möller, and Hamish Carr. Visualization of time-varying volumetric data using differential time-histogram table. In *Workshop on Volume Graphics 2005*, pages 21–29, 2005.