

Scatter segmentation in dynamic SPECT images using principal component analysis

Klaus D. Toennies^{1*}, Anna Celler^{**}, Stephan Blinder^{**}, Torsten Möller^{***}, Ronald Harrop^{***}

^{*}Computer Vision Group, Otto-von-Guericke-Universität Magdeburg, Germany,

^{**}Medical Imaging Research Group, Vancouver General Hospital, Vancouver, Canada

^{***}School of Computing Science, Simon Fraser University, Burnaby, Canada

ABSTRACT

Dynamic single photon emission computed tomography (dSPECT) provides time-varying spatial information about changes of tracer distribution in the body from data acquired using a standard (single slow rotation) protocol. Variations of tracer distribution observed in the images might be due to physiological processes in the body, but may also stem from reconstruction artefacts. These two possibilities are not easily separated because of the highly underdetermined nature of the dynamic reconstruction problem. Since it is expected that temporal changes in tracer distribution may carry important diagnostic information, the analysis of dynamic SPECT images should consider and use this additional information.

In this paper we present a segmentation scheme for aggregating voxels with similar time activity curves (TACs). Voxel aggregates are created through region merging based on a similarity criterion on a reduced set of features, which is derived after transformation into eigenspace. Region merging was carried out on dSPECT images from simulated and patient myocardial perfusion studies using various stopping criteria and ranges of accumulated variances in eigenspace. Results indicate that segmentation clearly separates heart and liver tissues from the background. The segmentation quality did not change significantly if more than 99% of the variance was incorporated into the feature vector. The heart behaviour followed an expected exponential decay curve while some variation of time behaviour in liver was observed. Scatter artefacts from photons originating from liver could be identified in long as well as in short studies.

Keywords: segmentation, dynamic imaging, principal component analysis, scatter detection, medical imaging

1. INTRODUCTION

Single photon emission tomography (SPECT) is a medical diagnostic imaging technique which visualises organ function using the three-dimensional activity distribution of a radioactive tracer which was injected prior to the measurement. The spatial resolution of this technique is of the order of 1cm^3 . The signal-to-noise ratio is fairly low because of dosimetry and instrumentation considerations. Dynamic SPECT extends the 3D-




Figure 1: Two different time steps of a dSPECT data set. Different behaviour of left ventricle of the heart and the liver is clearly visible.

¹ Computer Vision Group, Dept. Computer Science, PO Box 4120, D-39016 Magdeburg, Germany, e-mail klaus@isg.cs.uni-magdeburg.de

spatial imaging capabilities of standard SPECT into the fourth, temporal dimension by creating 3-dimensional “movies” representing changes of the tracer distribution in the patient body. In particular, the dSPECT method, introduced by Celler et al.^{1,2} provides such time-varying spatial information from data acquired using a standard (single slow rotation) SPECT protocol (see Figure 1).

Temporal variations of tracer distribution observed in the images might be due to physiological processes in the body and in this case may carry important diagnostic information about kinetic processes but might alternatively be due to reconstruction artefacts. These two possibilities are not easily identified and separated. The fact that the dynamic reconstruction seeks a solution to a highly underdetermined optimisation problem using only the low number of counts available in the acquired data causes the reconstruction to be susceptible to noise and other artefacts. The purpose of our work is to create a framework, which supports the separation of legitimate image information from reconstruction artefacts and provides tools for extracting this diagnostic information from the data.

The basic assumption of our approach is that, for dynamic tracers, activity changes in all voxels belonging to a region with the same physiology will display similar temporal behaviour. Different organs will generally correspond to different time activity curves (TACs). However, for purposes of patient diagnosis, another aspect of this assumption and of our image analysis procedure may be more important: the changes in physiology within an organ caused by disease could lead to different temporal behaviour thus correspond to different TACs and be identifiable by these TACs. Another important application of temporal analysis of the dynamic images may prove to be very valuable in scatter correction. Since photons originating from different organs will display temporal behaviour and have TACs characteristic to those organs, it may prove possible to use this fact to separate scattered photons from primary photons. This would be particularly important in cardiac studies, where scattered liver photons (one TAC) contribute to the image of the heart (different TAC) and might obscure myocardial perfusion defects.

Ideally, we endeavour to segment our image into a number of regions with “different” time activity behaviour. This is based on the assumption that the 3-d image can be described by an unknown number of compartments, the TACs within each compartment being “the same” for all voxels of the compartment. A compartment would be a functional unit, which is assumed to be enclosed within some organ’s boundary. An organ may contain more than one compartment in order to allow for different time behaviours within the organ, for example when blood and tissue are analysed separately or healthy and diseased tissues are being considered.

Segmentation would ideally find the compartments based on some homogeneity criterion applied to the data. In real life situations, the criteria, which define the meanings of the concepts “different”, “the same” and “region” get determined as regions are generated. In the absence of information about organ boundaries it cannot be expected that such criteria will enable exact separation of compartments. In order to reduce the likelihood of regions crossing organ boundaries, the criteria are set low which usually leads to the regions being smaller than normally hypothesised functional compartments.

The remainder of the paper is organised as follows. First, we describe the computation of those features of the time activity curves, which are used to separate the temporal effects related to physiology from artefacts due to noise or reconstruction errors. These features will be subjected to a segmentation procedure, which is presented in the following section. Features for scatter separation are discussed in section four. Section five describes results of our methods and is followed by concluding remarks in section six.

2. FEATURE COMPUTATION FROM DYNAMIC SPECT DATA

Dynamic SPECT (dSPECT) data used in this work corresponded to those from myocardial perfusion studies and were generated using dynamic tracer Tc99m-Tebroxime³. Each of five data sets was reconstructed as a time-series of voxel cubes consisting of 22 to 28 slices with 96x96 voxels in each slice. Each voxel in the series represents a single time activity curve (TAC). All studies were initiated at 2 minutes post injection and, depending on the study in question, their

TACs consisted of either 34 time steps spanning 12 minutes or of 16 time steps spanning 3 minutes. Artefacts distorting the data consist of noise due to a low photon count per voxel, possible reconstruction artefacts and artefacts due to deviations of the true situation from the model assumed in the reconstruction. An example of the latter would be scatter artefacts. The goal of segmentation is to suppress noise and reconstruction artefacts while creating segments of functional compartments for recognition of reliable diagnostic information as well as of scatter artefacts.

The dSPECT data may be considered to be a three-dimensional image with each voxel having its TAC as a feature vector. The dimension of the feature vector is the number of time steps for which 3D activities were reconstructed. We make the following assumptions about the data:

- The image consists of an unknown but finite number m of compartments c_i .
- Voxels of each compartment c_i have a common time activity curve TAC_i (feature) associated with them.
- Ideally, the TAC of a compartment is the result of a single cause. However, scatter and partial volume effects may lead to the time activity curve of a compartment being a linear combination of p_i partial time activity curves $pTAC_{ik}$, which are weighted by weights w_{ik} : $TAC_i = \sum_{k=1}^{p_i} w_{ik} pTAC_{ik}$.
- Noise N_i in the i -th compartment is gaussian distributed with zero mean.
- Artefacts from reconstruction can be expressed as an additive component $A(x,y,z)$. Its amplitude is significantly lower than that of the signal. For a compartment c_i , the artefact is a sum of all artefacts A_{ij} of the n_i voxels in the compartment.

The total observed signal of all TACs under these assumptions is a weighted sum of all time activity curves plus reconstruction artefacts and noise:

$$TAC_{total} = \sum_{i=1}^m \left(\sum_{k=1}^{p_i} w_{ik} pTAC_{ik} + \sum_{j=0}^{n_i} A_{ij} + N_i \right)$$

Segmentation should ideally be carried out using only the features of the time activity curves. However, as *a priori* we do not know the compartments (nor the features associated with them), the relevant features can only be identified based on the summed signal TAC_{total} . As the sum is a linear combination of signals, decomposition into linearly independent functions is appropriate. Such decomposition can be derived using principal component analysis (PCA)⁴. The result of PCA is a set of orthogonal eigenvectors and corresponding eigenvalues of the covariance matrix of the feature vectors. Each eigenvector is an enumeration of values of a function (called a base function), which is orthogonal to all other eigenvectors. The set of functions generated from a set of n -dimensional feature vectors span an n -dimensional eigenspace. Each vector can be expressed as a linear combination of base functions. Eigenvalues indicate the variance of all feature vectors from the corresponding eigenvector. Base functions are ordered by eigenvalues. The eigenvalue gives the variance that the corresponding base function contributes to the original signal. Base functions, with high corresponding eigenvalue will be called high significance functions. Functions with lower corresponding eigenvalue are called medium or low significance functions.

The base functions from PCA do not directly correspond to partial TACs, artefact functions or noise assumed in our model because these are not necessarily independent of each other. Each function $pTAC_{ik}$ in a compartment and each artefact function A_{ij} is an unknown linear combination of base functions. However, we may use our knowledge about the reconstruction technique to estimate which base functions mainly model the signal, which model reconstruction artefacts, and which model noise.

The reconstruction of the time signal from projection data was carried out by an iterative optimisation technique¹. The solution for each time step is computed from projections taken over a limited range of angles. Regularisation is usually needed because the system of equations is seriously underdetermined. Regularisation may use smoothness in time, smoothness in space and the requirement that the TAC has at most one peak. Consequently, TACs should consist of linear combinations of low frequency functions. The main difference between different areas in the image is the intensity of signal, i.e. the number of counts in the corresponding voxels. This number is related to the total number of

photons emitted from the corresponding area in the body over the time of the experiment, i.e. the average activity. This effect is reflected in the amplitude of TAC's. Thus, much of the variance is expected to be represented by a base function, which is almost constant. Noise, on the other hand, as it is assumed to be approximately gaussian and with zero mean, should have low amplitude in the average signal over all TACs. A base function from PCA, which mainly represent noise, should have a low variance. Biased reconstruction artefacts, which ideally should not exist at all, should contribute only little to the reconstructed signal. If such artefact remains constant over time, it is represented by a base function, which is almost constant. If it depends on the rotation angle, then its position should change between subsequent time frames. It should then be characterised by high frequency components.

In summary, the following conclusions can be drawn:

- The true TACs should be representable by low-frequency base functions. The corresponding eigenvalues for these base functions should be high.
- Noise is evenly distributed over the frequency range.
- Reconstruction artefacts are either constant over time or have high frequency.

This gives rise to the following predictions for the base functions of PCA:

- The base function with highest variance should be almost constant. It represents mainly the true activity averaged over time plus some noise and possibly reconstruction artefacts.
- The main variance should be represented by base functions with low frequency. These base functions mainly represent time-dependent activity plus some noise and reconstruction artefacts.
- The base functions with small corresponding eigenvalues should represent the high frequency components of the signal. They mainly represent reconstruction artefacts and noise.

If these hypotheses are reflected in the result of the PCA analysis applied to the data sets then the relevant features, which are to be input to the segmentation step, should correspond to the first few vectors in eigenspace.

Principal component analysis was applied separately to the TACs of each of five data sets. PCA was carried out on all data points that did not belong to the background. Background suppression was done so that its constant activity does not dominate the result. A voxel is considered to be background, if its average activity is below 1% of the maximum average activity of all voxels of the image. An evaluation of the eigenvalues from the PCAs of the five different data

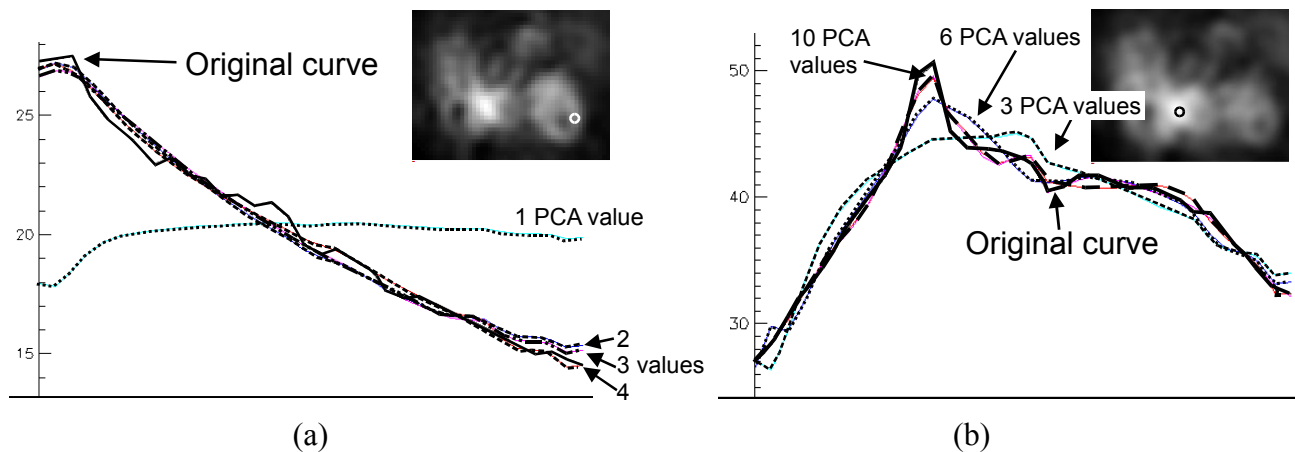


Figure 2: The TAC signal can be represented by few components of in eigenspace: (a) Heart TACs are mostly represented by the first three to four components. (b) A liver TAC requires about ten components. The original curve is drawn with a full line and the location of the samples are marked by the circles on the images.

sets showed that, as expected, most of the signal variance could be expressed by a linear combination of the first three to ten base functions (see Figure 2). The first base function closely resembled an averaging function over time. The next functions were mainly representing low frequency components of the signal. Signal decay, for instance, was represented mostly by the second function (see Figure 2a). Single peaks in the activity curve, being a clinically relevant signal in the liver, were represented by a linear combination of the first ten functions in the 34-dimensional eigenspace (see Figure 2b). The remaining functions described high frequency components at a low variance.

3. SEGMENTATION

The purpose of segmentation is to create groups of voxels with similar time activity behaviour. Voxel aggregates shall be smaller than the (unknown) compartments but big enough to sufficiently improve the signal-to-noise ratio of the average TAC in the segment. Segmentation is carried out by voxel aggregation through region merging^{5,6}. Input to the merging process was a 3D region adjacency graph (RAG), which was created from the voxel scene. Each node of the RAG represents a region (initially the voxels themselves). Two nodes are set to be connected by an edge if the corresponding regions share at least one common edge. Region information of a region R consists of

- a list of all voxels belonging to the region.
- values $PCA_{p,R}^{avg}$ for the $l=\{1,\dots,p\}$ principal components of the average time curve
- minimum and maximum of the PCA values $PCA_{p,R}^{min}$ and $PCA_{p,R}^{max}$ of all voxels in the region.

Each edge carries as attribute a similarity value for the two incident nodes and an “active” flag. The flag is set, if the two regions may be merged. Initially, the flag is set “true” for all edges in the RAG. Similarity is computed from a user-specified number of PCA values of the voxels in the scene. The first principal component is weighted with 20% of the weight of all other components. This is done in order to reduce the contribution of the average signal in favour of the activity decay.

Two different criteria were tested for computing the similarity between adjacent regions. The first was the inverse of the norm of the difference vector between minimum and maximum value of each feature of the feature vector:

$$S_1(R1, R2) = 1 / \left(1 + \sum_{p=1}^l \left\| \max(PCA_{p,R1}^{max}, PCA_{p,R2}^{max}) - \min(PCA_{p,R1}^{min}, PCA_{p,R2}^{min}) \right\| \right).$$

The second measure computed the likelihood that the average feature of one region could be found in the distribution of the features of the second region. For this measure, we estimate the means $\mu_{p,R1}$ and variances $\sigma_{p,R1}^2$ of a multivariate gaussian distribution of PCA values from the current members of a region $R1$. Then the probability of the average PCA value $PCA_{p,R2}^{avg}$ of a region $R2$ to be member of this group is computed based on this distribution:

$$S_2(R1, R2) = \prod_{l=1}^n \exp\left(-\left(PCA_{p,R2}^{avg} - \mu_{p,R1}\right)^2 / \sigma_{p,R1}^2\right).$$

Although the second measure more closely models the desired attributes of a region of similar TACs, we favour the first measure. This is for two reasons: The reliability of estimates of mean and variance for a region with few members is very low. Tests (using the total variance of all TACs as initial estimate for the variance for all regions consisting of a single voxel) revealed some instability in the results using S_2 . The unstable behaviour is caused by the grossly over-estimated variance for single-voxel regions. Results improve if the voxel variance is estimated from the variance in a local neighbourhood. However, variance is over-estimated at boundaries, which may still cause undesirable results. Using similarity measure S_1 leads to good and stable results. They are comparable to a segmentation using S_2 in those regions which are far from boundaries.




Figure 3: Regions of similar TACs were segmented with sufficient reliability even though only the first four components in eigenspace were used for computing similarity between TACs.

It is difficult to define a stopping criterion for the merging process as there is no absolute threshold known which would prevent merging of different compartments. However, our goal is to create regions that are large enough to increase the reliability of the segmentation based on the analysis of the TAC signal to an acceptable level. Using the static SPECT image as a

guideline, a region is considered to be large enough, when it contains about as many counts per time step as there are counts in a single voxel of the region of a corresponding static SPECT image. We estimated this number for one region of interest, e.g. the heart, and stopped the merging, if either of the following conditions were true:

- The number of counts in the region exceeds this threshold.
- The size of the region exceeds a certain maximum.

Region merging is carried out by repeatedly selecting the edge in the RAG, which is set “active” and has the highest similarity value. The region is created and, if either of the two conditions listed above is true, the activity flag for all edges of the new node are set to “false”.

Alternatively, selection of the next edge can be based on the similarity criterion alone. In this case, the stopping criterion is the number of regions. It is guaranteed that the similarity between voxels of all regions is better than the similarity of the last region that was merged.




Figure 4: Segmentation of reconstructed simulated data (with noise added to the measurement signal) showed varying agreement between true and reconstructed TACs. TACs of the original data are shown as dashed curves, average TACs of the segments are shown as solid curves. The liver was reconstructed as several regions with slightly different TACs. TACs of the large regions were similar to TACs in the original data. Original and reconstructed TACs in the heart were very similar to each other.

4. SCATTER DETECTION

The simplest model used for SPECT image reconstruction assumes a parallel, unattenuated propagation of gamma rays of the radioactive tracer. In reality, these rays are attenuated and scattered causing severe artefacts in the images. Attenuation effects can be relatively easily corrected using an iterative algorithm¹ but inclusion of scatter correction is much more difficult because it greatly decreases the sparsity of the projection matrix to be inverted.

Scattered photons are attributed to a line whose origin and direction is given by the collimator at the location where the photon is detected. Scattered photons lose some of their energy, and this fact can be used for partial removal of scatter by setting a window on the energy of the accepted photons. However, this removal is not complete and artefacts are still present in the data. In our case, the liver is expected to cause such artefacts because it has a particularly high

uptake of the Teboroxime and, as was mentioned in the Introduction, scattered photons originating from the liver may be attributed to other organs, such as the heart causing diagnostic problems. Fortunately, in the case of Teboroxime study, the typical TAC of the heart is very different from that of the liver. The maximum activity in the heart occurs usually shortly after the injection and in the liver at 6-8 minutes. In our particular case, acquisitions were started at 2 minutes post injection so the heart activity decreases monotonically during the time of the scan and the peak of the liver happens much later and is visible in TACs at the locations in the liver.

Following the argumentation above, normal behaviour of the Teboroxime in the heart in long and short studies is expected to show decrease from the first to the last time frame. The main components in eigenspace should be as following: a medium high value for the first component modelling the average activity and a high second component representing most of the decay curve. Most of the signal should be representable by the first few components. Liver with its late peak should exhibit a different behaviour. The first component should be higher because of the higher uptake of Teboroxime in the liver, the second component modelling the decrease of activity should be lower and the peak should be represented by higher values in components of medium significance.

5. RESULTS

As mentioned, the method was applied to dSPECT images reconstructed from simulated and patient studies of myocardial perfusion performed with the dynamic tracer Teboroxime-Tc99m (see Celler et al.³).

Segmentation of images reconstructed from simulated phantom data was carried out for validating our segmentation




Figure 5: Segmentations from a varying number of base functions of a 12-minute patient study (data set DYNR24).

procedure. Average TACs of segments in the original phantom data were compared with average TACs of segments containing the corresponding location in the reconstructed phantom. Simulated data had a Poisson noise at a level corresponding to the real-life clinical studies added to the projections prior to reconstruction. Segmentation was carried out on the first four base functions in eigenspace after applying principal component analysis. Segmentation was stopped using the total number of segments as a simplified stopping criterion. Results from a segmentation into 200 regions show that curves corresponding to single voxels were very similar to the average curve although only the four most significant base functions in eigenspace were used for similarity computation (see Figure 3). It can also be seen that the reconstruction process is not perfect. Reconstructed liver TACs deviate from those in the original phantom data (see Figure 4) although in many locations the average reconstructed signal indeed closely resembles the original TACs. The reconstructed signal in each segment of the heart was very similar to that in the original phantom data. This is depicted in Figure 4. It shows a slice of the reconstructed phantom data (top left: reconstruction from noisy simulated data, middle: segmentation from




Figure 6: Heart and liver tissues were clearly separable in the data. The left image shows the average activity of a slice of data set DYNR24. On the right grey-value-coded segments are shown. Coding was according to differences in the first four base functions of the PCA on the TACs.

60 counts, respectively. These numbers are clearly too low to provide reliable diagnostic information on a voxel per voxel basis. The region merging process was carried out using various stopping criteria (minimum counts varying between 200 and 2000 counts, size varying between 20-200 voxels). For the results shown in figures 6 to 9, large segments were generated with a count value of 2000 counts or 200 voxels per segment. This was done because larger regions are easier recognised in the figures of this paper. For generating reliable TACs, smaller regions are segmented in order to avoid segments exceeding compartment boundaries. However, results obtained were similar to those presented in the figures.

For segmentation, we compared various numbers of base functions representing between 98% and 99.9% of the total variance of the data. Although there is a noticeable difference in segmentations using two, three or four base functions this difference diminishes when more than 99.9% of the signal variance is represented (see Figure 5). This reduces the dimensionality of the feature vector to values between three and eight. Segmentation was carried out on five different data sets, of which one data set was a three-minute-study with 16 time steps while the others were 12-minute-studies consisting of 34 time steps.




Figure 7: The common attribute of all liver activity was the single peak in the TACs. However, position of the peak as well as steepness of increase and of decrease varied for different locations in the liver.

reconstructed data, right: simulated data without noise).

In segmentation of patient data, the goal was to separate heart and liver in the images and also to attempt to identify the contribution of scattered photons, which originate in liver and are recorded in the heart image, something that cannot be reliably estimated by studying single voxels. We also investigated the potential for anatomical classification of normal and diseased tissues based on their time behaviour. The average reconstructed activity per voxel per time frame in the heart and liver tissues were 15-20 and 40-

Our preliminary results indicate that segmentation clearly separates heart and liver tissues from the background even if the regions were chosen to be as large as mentioned above (see Figure 6). The average TAC of normal heart tissue (left and right ventricle) of the longer studies was easily differentiable from liver and spleen. In shorter studies the component corresponding to the average uptake over time dominated the display. The heart behaviour followed an expected exponential decay curve while some variation of time behaviour in liver was observed (see Figure 7). This agreed with our findings from the simulation study. However, in the patient studies the number of differently behaving compartments of the liver was higher and the differences were bigger than in simulated data. It may be possible that part of the variation is truly due to a different tracer uptake and blood flow in different parts of the liver.

Scatter artefacts of photons originating from liver were expected to be found in the heart and, indeed, were identified in long as well as in short studies (see Figure 8 and 9). This effect was visible in the first four principal components of the PCA images which showed typical liver behaviour in heart wall regions. Furthermore, it appears that liver compartments possessing a similar TAC are spatially close to regions of the heart wall that are mostly affected by scatter, therefore liver contribution is more easily identified there.

Classification results based on different time activity curve behaviours are promising but we find that presently the variations of the time behaviour in different areas of the liver are not understood sufficiently to provide information reliable enough for a purely data-driven classification.

6. DISCUSSION AND FUTURE WORK

We have presented a segmentation scheme for generating voxel aggregates in dynamic SPECT data in order to bring the signal-to-noise ratio to a level which would improve diagnostic potential. Segments are generated through three-dimensional region merging on the most significant components of the transformation of time activity curves in eigenspace. Analysis of segmented dSPECT data showed scatter artefacts in the left ventricle of the heart, that is in the region anatomically lying close to the liver, which were detectable without a specific scatter model. Separation of liver tissue from heart tissue was possible although purely data-driven classification of anatomic information in the functional data set is hampered by unexplained TAC variation in the liver.

Future work will concentrate on removal of scatter contribution from the heart images by using segmented TACs of the liver as a model. We will investigate the use of liver TACs in order to describe reconstructed heart segments as a linear combination of true heart behaviour and scattered photons. We will further investigate our finding that scatter segments in the heart are close to regions of similar behaviour in the liver. We will attempt to exploit this for an automatic generation of model curves for liver TACs and removal of liver scatter from the heart image. A second goal will be to develop a diagnostically meaningful method to display dSPECT information after segmentation and scatter removal. Segmentation results may be used to provide reliable gradient information for anisotropic filtering of scatter-corrected data. Scatter-corrected and noise-reduced heart TACs will be presented on a surface rendered representation, which is generated, from applying a deformable geometric heart model to the data.




Figure 8: Scatter is clearly visible in this 12-minute study (data set DYNR13). The graph below shows average TACs in segmented regions. It can be seen that the characteristics of compartments with scatter artefacts in the heart are similar to those in nearby locations in the liver.

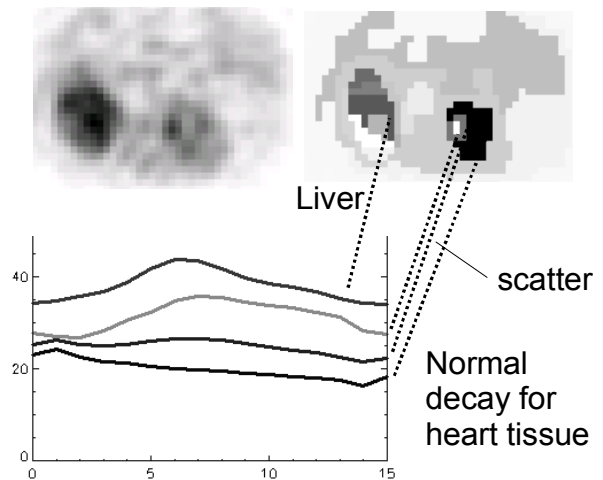


Figure 9: Scatter artefacts in a 3-minute study (data set DYN207). A small region in the heart shows a peak in its TAC which usually is attributed to the activity behaviour of the liver.

REFERENCES

1. T. Farncombe, A. Celler, D. Noll, R. Harrop, J. Maeght. Dynamic SPECT imaging using single camera rotations (dSPECT), *IEEE Transactions on Nuclear Sciences* 46, 1999, 1055 – 1061.
2. A. Celler, T. Farncombe, C. Bever, D. Noll, J. Maeght, R. Harrop, D. Lyster. Performance of the dynamic single photon emission computed tomography (dSPECT) method for decreasing or increasing activity changes, *Physics in Medicine and Biology* 45, 2000, 3525 – 3544.
3. A. Celler, S. Blinder, D. Noll, T. Tyler, F. Duclercq, R. Harrop. Investigation of Tc99m-Teboroxime myocardial perfusion using dynamic SPECT and a 4D kinetic thorax model dMCAT. Presented at the *2001 International Meeting on Fully 3D Image Reconstruction in Radiology and Nuclear Medicine*, Pacific Grove, CA, October 2001.
4. K.D.Toennies, L.Remonda, and R.Pohle. Shape-based enhancement of vascular structures in digital subtraction angiography images using local covariance information. *Proceedings of the SPIE (Medical Imaging 2000)*, 3979:630-637, 2000.
5. M.Sonka, V.Hlavak, R.Boyle. *Image Processing, Analysis and Machine Vision*, PWS Publishing, 1999.
6. R.Haris, S.N.Estradiadus, N.Maglaveras, and A.K.Katsaggelos. Hybrid image segmentation using watersheds and fast region merging. *IEEE Transactions on Image Processing*, 7(12): 1684-1699, 1999.