

A multiresolution image based approach for correction of partial volume effects in emission tomography

N Boussion¹, M Hatt, F Lamare, Y Bizais, A Turzo, C Cheze-Le Rest and D Visvikis

INSERM U650, Laboratoire du Traitement de l'Information Médicale (LaTIM), CHU Morvan, Brest, France

Received 11 October 2005, in final form 23 January 2006

Published DD MMM 2006

Online at stacks.iop.org/PMB/51/1

Abstract

Partial volume effects (PVE) are consequences of the limited spatial resolution in emission tomography. They lead to a loss of signal in tissues of size similar to the point spread function and induce activity spillover between regions. Although PVE can be corrected for by using algorithms that provide the correct radioactivity concentration in a series of regions of interest (ROIs), so far little attention has been given to the possibility of creating improved images as a result of PVE correction. Potential advantages of PVE-corrected images include the ability to accurately delineate functional volumes as well as improving tumour-to-background ratio, resulting in an associated improvement in the analysis of response to therapy studies and diagnostic examinations, respectively. The objective of our study was therefore to develop a methodology for PVE correction not only to enable the accurate recuperation of activity concentrations, but also to generate PVE-corrected images. In the multiresolution analysis that we define here, details of a high-resolution image H (MRI or CT) are extracted, transformed and integrated in a low-resolution image L (PET or SPECT). A discrete wavelet transform of both H and L images is performed by using the 'à trous' algorithm, which allows the spatial frequencies (details, edges, textures) to be obtained easily at a level of resolution common to H and L. A model is then inferred to build the lacking details of L from the high-frequency details in H. The process was successfully tested on synthetic and simulated data, proving the ability to obtain accurately corrected images. Quantitative PVE correction was found to be comparable with a method considered as a reference but limited to ROI analyses. Visual improvement and quantitative correction were also obtained in two examples of clinical images, the first using a combined PET/CT scanner with a lymphoma patient and the second using a FDG brain PET and corresponding T1-weighted MRI in an epileptic patient.

¹ Author to whom any correspondence should be addressed.

1. Introduction

Partial volume effects (PVE) are well-known consequences of the limited spatial resolution in emission tomography. PVE is characterized by the loss of signal in tissues of size similar to the point spread function (PSF). In addition, PVE induces a signal cross-contamination in adjacent structures with different amounts of radioactivity (Aston *et al* 2002, Du *et al* 2005). In this latter phenomenon, sometimes referred to as spillover, the high activity in a given region can spread out and contaminate a bordering area of lower activity, leading to either underestimated or overestimated activity concentration measurements.

These effects can be corrected for by using a number of different algorithms that often rely on the use of the PSF of the imaging device and *a priori* anatomical knowledge provided by computed tomography (CT) or magnetic resonance imaging (MRI) (Meltzer *et al* 1990, Muller-Gartner *et al* 1992, Rousset *et al* 2000, Aston *et al* 2002, Matsuda *et al* 2003, Baete *et al* 2004, Bencherif *et al* 2004, Quarantelli *et al* 2004, Kusano *et al* 2005, Rota Kops and Krause 2005). The large majority of these algorithms, which have been evaluated mostly in the context of cerebral imaging, require a segmentation step to delineate the different parts from anatomical images. This step renders their accuracy dependent on the segmentation algorithm used as well as making their application on other clinical investigations outside the brain challenging (Feuardenet *et al* 2003). For example, the pixel-based approach of Meltzer *et al* (1990) is restricted indeed to brain metabolism or neuroreceptor binding, and requires compartmental analysis (Meltzer *et al* 1999). As a rare example, Pretorius and King (2004) proposed an application of PVE correction for cardiac SPECT. Furthermore, and similar to the great majority of PVE correction methods (except in the interesting approach of Kennedy *et al* using Taylor expansion (Kennedy *et al* 2005)), these algorithms offer quantitative correction of ROI (region of interest) intensities without considering the construction of enhanced images. On the other hand, resolution compensation or resolution recovery algorithms can also be used to reduce PVE in emission tomography. However, the majority of these algorithms suffer from being reconstruction algorithm specific (Ardekani *et al* 1996, Som *et al* 1998, Somayajula *et al* 2005), as well as being only tested in limited clinical context such as cardiac SPECT (Hutton and Lau 1998) or FDG PET in the human brain (Baete *et al* 2004).

One of the reference methods (referred to from here onwards as RSF for regional spread function) described by Rousset *et al* (1998, 2000) and recently improved (Frouin *et al* 2002, Du *et al* 2005) was also developed in the brain context and allows estimating the true mean signal in any user-defined series of n homogeneous regions of interest (ROIs), but the images themselves are not enhanced. This approach relies on the inversion of an $n \times n$ matrix called geometric transfer matrix (GTM). The elements w_{ij} of the GTM are the coefficients of activity spillage from ROI i to ROI j , and the true activity T_i in ROI i can be deduced from the measured activity t_i by inverting the equation $[t] = [\text{GTM}] \times [T]$, where $[t]$ and $[T]$ are the vectors containing the t_i and T_i values, respectively. The use of this approach is theoretically possible in various clinical applications even if it was originally designed for cerebral studies where generally only three ROIs are required (white matter, grey matter, cerebrospinal fluid). Actually, the method works satisfactorily when the image is segmented into a series of ROIs that constitute a partition. In other words, ROIs must not overlap and at the same time considering all ROIs together must cover the entire image. As a consequence, when studying tumours in whole-body images, the number of ROIs can dramatically increase (Feuardenet *et al* 2003), thus hampering the clinical use of this methodology.

In general, the aim of all these methods is to provide the user with correct radioactivity concentration estimates in a given ROI. To date however, except in very specific applications (Baete *et al* 2004), little attention has been given to the challenging possibility of creating

improved images through a generic approach. In this paper, a new PVE correction methodology is proposed, based on the multiresolution analysis of images of different spatial resolutions. The main advantage of the proposed methodology is that it not only enables the accurate recuperation of activity concentrations, but is also capable of simultaneously generating PVE-corrected images. This improvement allows (a) performing a visual control of the correction, (b) improving clinical diagnostic studies through a better visual assessment of the images, and most important (c) allowing further image processing (such as, for example, functional volume estimate of tumours, location of epileptogenic foci in cerebral imaging (Boussion *et al* 2003), or wall motion and ejection fraction in cardiac imaging (Hickey *et al* 2004)). Furthermore, the method is not restricted to a particular organ and does not require tedious and time-consuming ROI delineation.

In the following section, a concise presentation of the wavelet transform and the multiresolution analysis serves as an introduction to section 2. The developed PVE algorithm is described in detail including a description of the test images and overall methodology used to validate the developed algorithm.

2. Materials and methods

2.1. Multiresolution image analysis and wavelet transform

Although the theoretical foundations of multiresolution analysis do not constitute the main topic of this study, it is constructive to introduce the basic concepts of the wavelet transform which is an important part of the proposed methodology. Actually, the wavelet transform can be introduced by comparison with the more common Fourier transform with which it has a number of similarities. While the Fourier transform provides global information about the spatial frequencies in an image, the wavelet transform leads to a local representation of these spectral properties. From an image processing point of view, the Fourier transform permits one to switch between the spatial and the frequency domains while the wavelet transform allows one to bring them together in one single image. In practice, the wavelet transform of a given image is another image presenting the areas where one may find either more or less important contrast. In addition, one of the interests of the wavelet transform in image processing is that it enables work at different levels of spatial resolution, operating as a tool of multiresolution analysis. Multiresolution analysis allows retrieving the layers of details that have different sizes by separating the spatial frequencies that the image contains. Basically, a medical image at a given spatial resolution R contains information at different scales, from large structures to small details. For instance, in a cerebral MRI the sharp edges between white and grey matters will be lost when a low-pass filter is applied, but at the same time the skull will stay clearly separated from the brain. Accessing and separating these structures of different sizes is the scope of multiresolution analysis.

If we now consider the mathematic point of view, the wavelet transform allows expressing a signal according to a basis of elementary functions called wavelets. This basis is built from a ‘mother’ wavelet ψ (also referred to as analysing wavelet) on which are applied dilation and translation computations. This process is obtained in one dimension as a result of the following formula:

$$\psi_{a,b}(x) = \frac{1}{\sqrt{a}} \psi \left(\frac{x-b}{a} \right) \quad (a > 0). \quad (1)$$

a is called the scale parameter and is linked to the frequency domain, while b is the position parameter linked to time or space.

The wavelet transform $W(a, b)$ of the function $f(x)$ is defined as

$$W(a, b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} f(x) \psi^* \left(\frac{x-b}{a} \right) dx, \quad (2)$$

where ψ^* stands for the complex conjugate of the analysing wavelet ψ . W is linear, shift invariant and also invariant by dilation. These latter two properties are of interest in image processing involving combination of different images. Actually, on the one hand, shift invariance limits the unavoidable consequences of inaccurate superimposition of images. On the other hand, dilation invariance is valuable for observing ‘objects’ of different sizes in a given signal without changing the analysing wavelet.

There are many algorithms available to perform the discrete wavelet transform of an image. All have particular interests and drawbacks but they must be chosen carefully because the passage to the discrete approach can lead to the loss of interesting properties such as invariance mentioned above. A widely used approach is the pyramidal methodology which consists of reducing the size of the image iteratively to get smoother and smoother versions of the initial image. This is the widespread multiresolution approach that Mallat (1989) developed through his algorithm that permits compression of data by decimating the image. This method is anisotropic in the sense that horizontal, diagonal and vertical details are separated during the process. Another common approach is the algorithm ‘à trous’ (French term that means ‘with holes’). This is an undecimated method inducing shift invariance which is of particular interest when investigating image comparison. The transformation is not pyramidal since the initial image and the images of coarser spatial resolution have identical sizes. For this reason, this particular algorithm is redundant and is of reduced interest in image compression. This algorithm forms however the basis of our PVE correction methodology as it presents several practical advantages, namely (a) the implementation is straightforward and the initial image can be perfectly reconstructed without loss of any kind, (b) there is no selection of specific directions during the analysis since the process is isotropic, (c) the transform is known for each pixel improving accuracy of further processing, and (d) navigation is easy between the different levels of resolution.

This discrete wavelet transform algorithm called ‘à trous’ was introduced by Dutilleul (1987), developed by Holdschneider *et al* (1989) and detailed by Starck *et al* (1998). The process gives an image sequence of coarser and coarser spatial resolution by performing successive convolutions with a low-pass filter h obtained from a scaling function ϕ . At each iteration j , the spatial resolution of the image I_j is degraded to give the approximation image I_{j+1} according to

$$I_{j+1}(k, l) = \sum_{m,n} h(m, n) I_j(k + m2^j, l + n2^j). \quad (3)$$

As already pointed out, there is no decimation involved in the process, which means that all I_j approximations have the size of the initial image I_0 . However, only one pixel out of 2^j is considered during the filtering process, leading to inclusion of zeros in the rows and the columns of the mask. This feature gives its name to the algorithm, i.e. ‘with holes’, and it also explains why the process is dyadic, where the successive approximations I_j have resolutions decreasing by powers of 2.

The difference $I_j - I_{j+1}$ is the wavelet coefficients w_{j+1} containing the details (edges, texture) at a resolution level between I_j and I_{j+1} . Note that the undecimation permits one to follow the local information at a pixel level for any I_j , that is, navigation through all I_j images

is possible at any pixel location. The synthesis procedure that reconstructs the original image from its layers of details w_k is given by

$$I_0 = I_N + \sum_{k=1}^{k=N} w_k, \quad (4)$$

with N the number of iterations from the initial image I_0 to the final approximation I_N of spatial resolution decreased by 2^N . A pixel at location (x, y) can be expressed as the sum of the wavelet coefficients at this position plus the smoothed array at the same (x, y) coordinates:

$$I_0(x, y) = I_N(x, y) + \sum_{k=1}^{k=N} w_k(x, y). \quad (5)$$

The ‘à trous’ algorithm can easily be implemented by performing the following steps (Starck *et al* 1998):

- (1) Initialize j to 0: start with the original image I_0 .
- (2) Increment j and carry out a convolution of I_{j-1} with the low-pass filter h . The distance between the central pixel and the adjacent ones is 2^{j-1} .
- (3) The wavelet coefficients w_j at this level of resolution are given by $I_{j-1} - I_j$.
- (4) If j is less than the required number N of resolutions, go to step 2.
- (5) The set $W = \{w_1, w_2, \dots, w_N, I_N\}$ is the wavelet transform of I_0 .

Provided they satisfy a limited number of properties (compacity, regularity, symmetry) and according to suitable prerequisites, different scaling functions can be constructed. However, several already exist possessing interesting characteristics. The most widely used filters in the ‘à trous’ algorithm are based on linear interpolation and B-splines interpolation. For instance, the bicubic spline is a very smooth function, well suited for isolation of large image structures. On the other hand, linear interpolation is a good compromise, enabling work with both small and large scale characteristics. Another filter, sometimes called low-scale filter, is a sharply peaked function that performs well in isolating very small structures. The normalized coefficients of these different filters are presented in the appendix. Each of these filters were tested under different image characteristics in order to evaluate their behaviour and choose the most appropriate one in the framework of the developed PVE correction algorithm.

2.2. Description of the algorithm implementation

The process employed in the developed algorithm comes from the field of data fusion (Luo *et al* 2002) and as stated above it relies on a wavelet-based image merging. Actually, new approaches to image merging that uses multiresolution analysis procedures based upon the discrete wavelet transform have been proposed recently in as different domains as texture classification (Li and Shawe-Taylor 2005), forensic science (Wen and Chen 2004) or aerial images (Ranchin and Wald 2000). The multiscale fusion that we define here is the process whereby details of a high-resolution image H (MRI or CT typically) are extracted, transformed according to a given model and integrated in a low-resolution image L like PET or SPECT for instance. The challenge is to preserve the global functional characteristic of L while incorporating additional data in it and the mandatory hypothesis is that the tissues examined by L are also present in the high-resolution image H . Contrary to all other applications that have been studied till now, such as aerial imaging or forensic sciences, the visual enhancement is not here the unique goal of the process. Our primary objective is the quantitative improvement of the recovered activity concentrations. This latter is possible by adding detail layers of different resolutions that all have a zero-mean signal.




Figure 1. (a) The high-resolution H image with discs of various sizes and intensities. The five horizontal discs have identical intensities and decreasing sizes in order to evaluate the correction of tissue-fraction effects. The six vertical discs have decreasing intensities to allow studying spillover effects. (b) Low-resolution image L corresponding to the degradation of high-resolution image H after 10-standard-deviation Gaussian noise addition and low-pass filtering. (c) L image after PVE correction using the developed wavelet-based algorithm.

Wavelet analysis allows the spatial frequencies to be obtained easily at any level of resolution, in particular at a level of resolution common to H and L. A model is then inferred to compute the lacking details of L from the high-frequency details' layers of H. If the level of resolution of H is q , referred to as H_q , and that of L is $r = q + p$, referred to as L_r , we can write

$$L_r(x, y) = L_{q+p}(x, y) = L_{q+p+1}(x, y) + w_{q+p+1}^L(x, y) \quad (6)$$

and

$$H_q(x, y) = H_{q+p+1}(x, y) + \sum_{k=1}^{k=p+1} w_{q+k}^H(x, y). \quad (7)$$

The lacking details of L are the wavelet coefficients w_i^L with $q \leq i \leq q + p$. However, we do possess w_{q+p+1}^L and w_{q+p+1}^H and we assume that there exists a more or less simple link between them like $w_{q+p+1}^L = \alpha \times w_{q+p+1}^H$, $\alpha \in \mathbb{R}^*$ for instance. Although, different models can be envisaged, in this study a simple linear model is used where the parameter α is considered equal to the mean pixel-by-pixel division of w_{q+p+1}^L by w_{q+p+1}^H .

L_q can now be reconstructed from L_r by taking w_i^L ($q \leq i \leq q + p$) into account. They are calculated as $w_i^L = \alpha \times w_i^H$ ($q \leq i \leq q + p$) leading to

$$L_q(x, y) = L_{q+p+1}(x, y) + \alpha \sum_{k=1}^{k=p+1} w_{q+k}^H. \quad (8)$$

2.3. Validation studies

2.3.1. Synthetic and simulated images. The developed algorithm was firstly validated using different synthetic and simulated datasets. Synthetic images were composed of a circular container (intensity 50) including discs of different sizes and contrast ratios (figure 1(a)). A first series of five horizontal discs of decreasing diameter (30 mm, 20 mm, 15 mm, 10 mm, 5 mm) and constant intensity (70) was built up to specifically examine the tissue-fraction effect and the recovering of small areas. A second series of six vertical discs of constant diameter (10 mm) but decreasing intensity (70, 60, 40, 30, 20 and 10) was designed specifically to consider

spillover effects. This image of high spatial resolution (1 mm) corresponded to H as stated in the previous section. A n -standard-deviation (SD) Gaussian noise was added (n ranging from 1 to 10) and a 6 mm FWHM Gaussian blur was convolved in order to simulate the L images (figure 1(b)) with a uniform 6 mm spatial resolution (the letter n in n -standard deviation corresponds to the amount of noise added on a given pixel; the Gaussian law built with a zero mean and a standard deviation equal to n times the standard deviation calculated in a 5×5 pixel ROI around the pixel to treat). The convolution induced a contamination of signal between homogeneous areas of the synthetic images very similar to partial volume effects, and the levels of resolution for L and H were chosen according to those of typical PET and MRI studies, respectively. The different amounts of noise in L, introduced by the variable SD Gaussian noise, aimed at investigating the performance of the correction algorithm in PET images of variable statistical quality.

The mean intensity in the discs inside the container was calculated before and after PVE correction in each one of the ten L images and then compared with actual values in H. In practice, ROIs delineating discs in L were obtained automatically by copying the masks of exact discs in H. The mean intensity in the discs of L images was then calculated inside these exact ROIs. These results were also compared with those obtained by applying the RSF method (Rousset *et al* 1998) often considered as the reference numerical approach (Frouin *et al* 2002, Quarantelli *et al* 2004) for PVE correction in emission tomography.

The robustness of the developed algorithm to spatial misalignment between the H and L images was also studied by introducing artificial displacements of L with regards to H. A set of 20 configurations was created, namely 1, 2 and 3 pixels (each pixel 1 mm of size) translation errors in the four directions (up, down, left, right), 1, 2 and 3 degrees of rotations clockwise and anticlockwise, and finally, two scaling errors of 99% and 101%. The L image with 6 SD Gaussian noise was considered for this specific investigation. The error in intensity recovery was calculated in the 11 discs for each of the 20 configurations of misalignment produced, after applying either the RSF PVE correction or our wavelet-based method.

The synthetic images were also used to assess the proposed methodology in cases where image contents no longer correlate. As already mentioned, one of the mandatory prerequisites to the application of the correction method proposed in this paper is the similarity of tissues in the images we are dealing with. For this purpose, and without altering the low-resolution image, we modified the synthetic 'CT' image to create two grossly unfavourable configurations. In the first one, the horizontal series of five discs was completely removed from the synthetic 'CT' image (figure 2(a)), and in the second one (figure 2(c)) the intensity of the first disc in this same series was set to 20 (cold intensity) instead of 70 (hot intensity, as in the 'PET' image).

Finally, simulated images were also included in our study. They consisted of a simplified numerical version (figure 3(a)) of the physical IEC phantom (IEC Publication 61675-1 1998). This phantom consists of a 20 cm diameter by 20 cm long cylinder, containing six spheres of 37 mm, 28 mm, 22 mm, 17 mm, 13 mm and 10 mm in diameter. The numerical version of this phantom was produced as a set of 64 contiguous planes of 64×64 square pixels of $4 \text{ mm} \times 4 \text{ mm}$ in size. This phantom was subsequently combined with a Monte Carlo based simulation of the Philips Allegro PET scanner using GATE (Lamare *et al* 2006). A total of 60 million coincidences were simulated considering a sphere/cylinder activity concentration ratio of 5/1. Images were subsequently reconstructed using the OPLEM algorithm (11 iterations) (Reader *et al* 2002). The high-resolution image serving for PVE correction was the numerical phantom in which values were arbitrarily set to 1000 for the background, 2000 for the cylindrical container and 3000 for the spheres, leading to a 1.5 sphere/cylinder intensity ratio. These realistic ratios in the numerical phantom and in the simulated PET image were chosen in order




Figure 2. Synthetic images with no tissue correlation between high- and low-resolution images; only the high-resolution image is altered. (a) First configuration, without horizontal discs in the high-resolution image, and (b) corresponding PVE-corrected low-resolution image. (c) Second configuration, with contrast modified in the first horizontal disc only, and (d) corresponding PVE-corrected low-resolution image. A local application of the proposed algorithm is shown in (e), with the region of interest surrounding the disc to be corrected, and (f) the whole image after PVE correction demonstrating that only the part inside the specified region of interest is PVE corrected.

to investigate the behaviour of the developed PVE correction methodology in more realistic conditions than the synthetic images.

2.3.2. Quantitative and qualitative assessment of the developed methodology on clinical images. In order to demonstrate the use of the developed algorithm in the clinical context, the technique was applied on two different sets of patients' images. The first one consisted of a whole-body FDG PET and corresponding CT images figure 4 acquired on a lymphoma patient using a dedicated combined PET/CT scanner (GE Discovery LS), while the second dataset consisted of a FDG brain PET and corresponding T1-weighted MRI images acquired during pre-surgical evaluation of refractory epilepsy figure 5. The MRI and PET cerebral




Figure 3. Simulated numerical image of the IEC image quality phantom: (a) uncorrected and (b) PVE corrected.




Figure 4. Clinical PET/CT patient study of the thorax acquired on a combined PET/CT scanner. (a) Original emission PET FDG image. (b) Corresponding CT image at the same anatomical location (identical slice level). (c) FDG PET after PVE correction.

images were acquired separately and spatially co-registered by using mutual information maximization (Wells *et al* 1996) and affine transformation (rotation, translation, scaling).




Figure 5. (a) Transaxial brain PET image obtained with FDG in an epilepsy follow-up study. (b) Corresponding MRI slice (the skull has been removed). (c) Same transaxial plane as in (a), here shown after PVE correction by using the developed wavelet-based multiresolution analysis algorithm.

Apart from the qualitative assessment of the corrected images, a ROI analysis was also performed to quantify the impact of the PVE correction. Different regions were drawn in both the original and corrected whole-body PET images, namely four circular ROIs of 30 mm in diameter placed at the middle of the right lung corresponding to a normal area with homogeneous intensity (ROI_{lung}), two circular ROIs of 20 mm in diameter inside the heart in a normal but visually inhomogeneous area (ROI_{heart}) and a ROI surrounding the lesion in left lung (ROI_{lesion} , size 20 mm \times 8 mm). A similar quantitative investigation was performed in the FDG brain PET, in which grey matter and white matter intensities were calculated before and after PVE correction. As a first step, grey and white matters were delineated automatically in the T1-weighted MRI using the SPM software (Ashburner and Friston 2000), and in a second step the two obtained segmented areas were superimposed on the PET image. They both served as ROIs in which mean intensities were calculated.

3. Results

Figure 1(b) shows an example of the synthetic image L used to assess the developed methodology. The corresponding PVE-corrected image of L is given in figure 1(c), where it can be noted that, aside from quantitative considerations, edges are visually enhanced. Figure 6 illustrates the PVE correction in a semi-quantitative fashion. Profiles across the five horizontal discs in uncorrected and corrected L images along with the related wavelet coefficients built from the H image and representing the correction values are presented. In the profile presented in figure 6(a), one can note that the discs are represented by five coarse Gaussian shapes, the last being so attenuated that it is hardly visible. The corresponding profile in figure 6(b) shows that the correction required for this disc is the most significant among the five discs, leading at the end to equal corrected values as demonstrated by figure 6(c). As a consequence, the process does not only enhance the edges of objects but also increases the global intensity level when needed, especially for smaller objects.

The global recovery of intensity in the low-resolution synthetic images, considering the ten different levels of image noise described in section 2.3.1, is represented in figure 7, where results concerning tissue-fraction and spillover effects are separated in figures 7(a) and 7(b), respectively. As figure 7(a) demonstrates, the intensity level in the five discs with diminishing sizes in the uncorrected L images was found to dramatically decrease clearly demonstrating the tissue-fraction effects. For example, the intensity in disc 1 (30 mm diameter) and disc 5

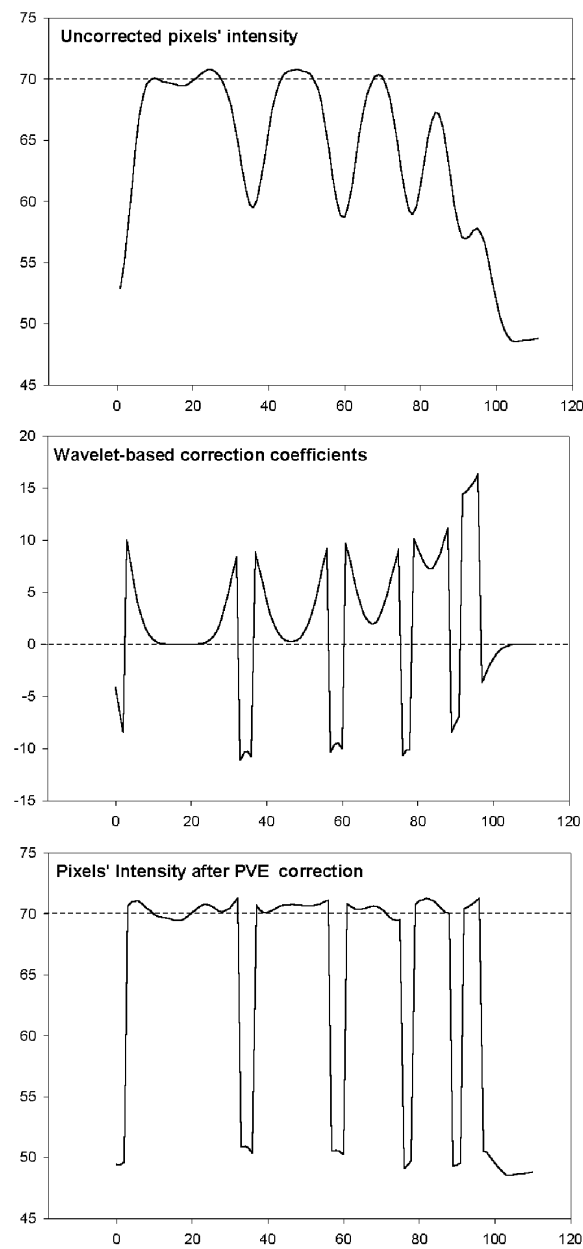


Figure 6. A semi-quantitative assessment of the results obtained from the application of the developed PVE methodology on the synthetic images (considering the L image with 5-standard-deviation Gaussian noise). A 'plot profile' is generated along a line crossing the five horizontal discs, in the uncorrected image (a), in the wavelet-based correction image (b) and in the corrected image which is the pixel-to-pixel addition of the first two (c).

(5 mm diameter) decreased from 70 to 67.2 (−4.0%) and 56.7 (−19.0%), respectively. The spillover effect is illustrated by the six vertical discs, whose intensities were either underestimated or overestimated depending on their initial signal-to-background (S/B) ratios (figure 7(b)). For example, the intensity in disc 6 (10 mm in diameter, initial S/B = 1.4)

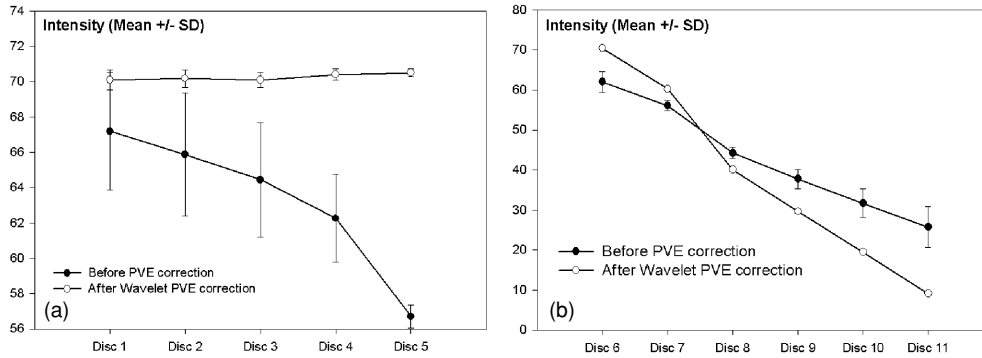


Figure 7. Quantitative performance assessment of the PVE correction. The results from the L synthetic images of the ten different noise levels considered are summarized in this figure, namely (a) mean intensity recovery in the five discs of diminishing sizes, with the errors bars representing the standard deviation in the mean taking into account the ten different noise levels considered (the actual intensity is 70), (b) mean intensity in the six discs of decreasing intensities (the actual values are 70, 60, 40, 30, 20 and 10 for discs 6–11, respectively).

decreased by 11.4% and the intensity in disc 11 (10 mm in diameter, initial S/B = 0.2) increased by 158.0% (from 10 to 25.8). Following the application of the developed algorithm, both phenomena were corrected for as also illustrated in figure 7. It is however of significance to observe the dependence of the correction upon the choice of the filter. In figure 8, for example, the results of the correction are shown for the 11 discs simulating the tissue-fraction effect and spillover, according to the four filters described in section 2.1 and the appendix. For this purpose, we call recovery error the difference between the expected disc intensity (for example, 70 in the discs 1–5) and the intensity in the same disc after PVE correction. A perfect PVE correction would then lead to a recovery error equal to 0. It is clear that the bicubic spline filter and the 5×5 linear filter perform better than the other two. For discs 1–5, the percentage of recovery error was less than 1% for these two filters, while the low-scale filter led to error greater than 4%. For discs 9–11, the percentage of recovery error exceeded 20% with this low-scale filter. According to these results, the bicubic spline filter or the 5×5 linear filter should be preferred. Consequently, all results presented in this paper were obtained with the bicubic spline filter.

The good noise characteristics of the developed algorithm are shown in figure 9 where the percentage of correction error is plotted against the different L images considering variable noise levels. As far as tissue-fraction effects are considered (figure 9(a)), a negligible overestimation error was found, globally increasing with respect to the amount of noise. However, errors never exceeded 1% in the set of images considered. Concerning the spillover, the correction slightly underestimated the true values, with an error up to 4% in high noise conditions figure 9(b). Finally, the comparison with the RSF method is given in figure 10. The graphs show that the two methods have very similar behaviours in correcting for both tissue-fraction effects (discs 1–5) and spillover (discs 6–11) since the two graphic representations almost perfectly overlap. However, a closer comparison reveals a slight difference in favour of the wavelet method. For the discs 4 and 5, which are highly subjected to tissue-fraction effect on account of their small sizes, the errors in intensity recovery are, respectively, 0.07% and 0.33% for the wavelet approach against 0.34% and 0.92% for the RSF method. The difference is more significant when considering the disc 11 which undergoes substantial spill-in from the surrounding area: 0.47% of error for the proposed method against 4.27% for the RSF approach.




Figure 8. Influence of the filter used in the 'à trous' algorithm on the percentage of recovery error. Results are presented for the 11 discs of the synthetic images, dedicated to the study of tissue-fraction effects (discs 1–5) and spillover (discs 6–11). Values are mean percentages obtained from all ten different noise level synthetic images.

Figure 11 summarizes the results on the effects of spatial registration errors, demonstrating that the developed algorithm behaves similarly to the RSF method. Both methods gave satisfactory PVE correction of tissue-fraction effects (discs 1–5). The correction of PVE in discs simulating spillover seemed to be more dependant upon co-registration than the correction of tissue-fraction effects (discs 6–11). However, as figure 11 demonstrates the developed algorithm performed better in correcting spillover effects than the RSF method, since, in some cases, the latter led to errors exceeding 100% for discs 10 and 11.

The effects of grossly unfavourable configurations for the methodology described in this work, where tissues greatly differ between high- and low-resolution images, are demonstrated in figure 2. The image in figure 2(a) corresponds to the case where horizontal discs are erased in the high-resolution image. As a result, in the corresponding corrected image figure 2(b) vertical discs are PVE corrected as is the contour of the container, but the horizontal series of discs stays unmodified since no corresponding information exists in the high-resolution

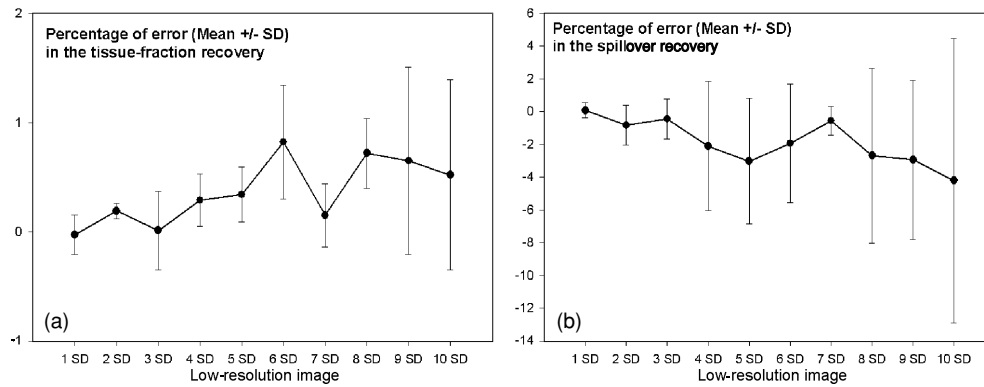


Figure 9. Robustness of the wavelet-based PVE correction according to noise. (a) Error percentage of intensity recovery in the five discs of decreasing sizes, and (b) error percentage of intensity recovery in the six discs of decreasing intensities.




Figure 10. Comparison of the wavelet-based PVE correction with the RSF approach of Rousset *et al*, presented in terms of percentage of recovery error. Results are mean values obtained in the set of ten L images with different levels of noise. For the sake of clarity, the standard deviations (error bars) are presented for the wavelet-based method only.

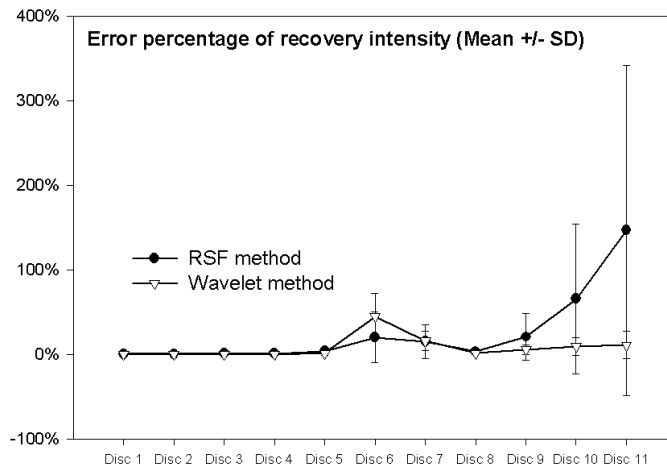


Figure 11. Robustness of the correction according to the alignment accuracy of the images. Twenty different configurations were tested (12 translation movements, 6 rotation movements and 2 inadequate scalings). For each disc, the value presented is the mean error of intensity recovery obtained in the set of 20 configurations considered.

image. In terms of quantitative accuracy, there has been no alteration in the values of the horizontal spheres independently of their sizes. Therefore, the complete lack of a given tissue in the high-resolution image does not modify the corresponding part in the low-resolution image. The image in figure 2(c) relates to the second case where the intensity in an isolated disc is altered compared to the other discs ('cold spot' rather than a 'hot spot'). This leads to a local artefact in the corrected image corresponding to the limits of the 'cold sphere' in the high-resolution image, mainly due to a local inverse contrast in the high-resolution image compared to the same area in the low-resolution image (figure 2(d)). The quantitative errors are also local to the visual artefact with values not altered in the rest of the sphere.

The transaxial slice of the reconstructed PET image of the numerical IEC phantom containing all the lesions is shown in figure 3(a). The corresponding PVE-corrected image is given in figure 3(b), while the quantitative results (an expected ratio between spheres and background of 5/1, irrespective of lesion size) are presented in figure 12. As this latter figure demonstrates, although both correction methods lead to an improvement in the sphere/cylinder ratios, the wavelet-based correction performed better in all spheres, particularly for the smallest ones. For example, the ratios for the 13 mm and 10 mm diameter spheres were improved from 2.19 to 5.93 and from 2.03 to 5.96, respectively, using the wavelet-based algorithm, against 9.08 and 9.80 with the RSF method.

Finally, figures 4(c) and 5(c) clearly demonstrate the visual image quality improvement achieved in both oncology and brain clinical applications, permitted by the generation of PVE-corrected images using the developed algorithm. Aside from these visual improvements, ROI analyses were performed to get a quantitative insight of the correction in clinical images. In the whole-body PET image and before PVE correction, the average intensities in ROI_{lung} , ROI_{heart} and ROI_{lesion} were 35.1, 245.2 and 109.2, respectively. After correction with the proposed multiresolution method, these mean intensities changed to 36.4 (+3.7%), 240.0 (-2.1%) and 129.2 (+18.3%), respectively. This led to an increase of lesion-to-lung ratio of 16.1%. In the brain PET image, the mean intensity in white matter was 113.8 before correction and 90.7 after PVE correction. In the grey matter, the values were 126.6 before correction and 162.3 after correction, representing a 28.2% increase.




Figure 12. Sphere-to-background ratio in each of the different diameter spheres in the transaxial reconstructed slice of the simulated IEC phantom, before PVE correction and after PVE correction by either the RSF method or the developed wavelet-based algorithm. Expected ratios are equal to 5.

4. Discussion

The development of spatial co-registration algorithms in the last decade has allowed the automated and reliable superimposition of multimodality images in a day-to-day practice, in as different domains as neurology or cardiology. Moreover, since the advent of combined PET/CT and more recently SPECT/CT scanners, there has been a widespread acceptance of this new technology. This development has led to easier and direct superimposition of functional and anatomical images (Vogel *et al* 2004) for oncology applications. Consequently, the use of anatomical data has naturally become one of the keys in addressing the problem of partial volume effects in emission tomography. However, the current PVE correction methods focus mainly on cerebral imaging, while PVE remains a major problem in other applications notably in oncology and whole-body studies. Actually, the size of tumours is often close to the PSF of PET scanners and consequently they are significantly exposed to PVE (Soret *et al* 2001). In this case, the activity and the dimension of tumours, which are critical parameters in quantitative accuracy for applications such as response to therapy or radiotherapy treatment planning (Caldwell *et al* 2003), become difficult to assess.

As previously stated, most of the different methods of correction that have been published till now suffer from the need to perform a segmentation of the anatomical information of interest and their subsequent specificity to brain imaging, with very few procedures having been tested on other organs. In addition, the vast majority of developed algorithms focus on the recuperation of accurate activity concentrations and not on yielding PVE-corrected images. The methodology developed by Rousset *et al* (1998) and referred to throughout this paper as the RSF method remains indeed the most widely used approach because it is straightforward to implement and it only requires the knowledge of the PSF. Like many other approaches it however also needs the segmentation of anatomical structures, which may become tedious and time consuming when no automatic algorithm is available. It is important as well to underline that this algorithm needs the ROIs to be accurately delineated (Frouin *et al* 2002, Zaidi *et al*

2005) which may be uncertain in the case of manual segmentation such as may be the case in whole-body PET.

In this paper, a novel approach aiming at overcoming these limitations was proposed. The general concept of the technique is a mutual multiresolution analysis of functional and anatomical images that are supposed to be correctly co-registered. During the process, the details in the high-resolution image are automatically extracted, altered according to a model and introduced in the PET image by using a simple pixel-to-pixel addition. The correction image containing these modified details is built from wavelet coefficients and consequently has a zero mean. As a result, the algorithm adjusts the intensity in boundaries of the organs and greatly enhances the smallest parts like lesions or tumours, but at the same time it hardly modifies the intensity in large and homogeneous structures.

The performance of the developed algorithm was assessed through the use of synthetic and simulated images in direct comparison with the most popular PVE correction methodology developed by Rousset *et al* in the first instance for brain imaging applications. The data obtained from the use of synthetic images demonstrated the correction ability of both aspects of PVE, that is to say tissue-fraction effects and spillover, with as good accuracy as the RSF method. In addition to the quantitatively accurate results, the developed methodology allowed the derivation of enhanced images. One may argue that the synthetic high-resolution image H had ideal properties since the L images were derived from it. Actually, the intensity in the different parts of H was rigorously the same as that we wanted to retrieve in L. However, the simple model that we defined between the wavelet coefficients of L and H allowed us to obtain similar corrected values in L by flipping the contrast in H, illustrating the robustness of the model against the nature of the high-resolution image. Similarly, good results were obtained using simulated PET images of the IEC phantom, considering the presence of lesions from 10 mm to 40 mm in diameter, with the developed methodology leading to a larger improvement in comparison to the RSF method, particularly for smaller lesions. Finally, the improvement that can be derived from the developed technique in the clinical setting was demonstrated by the enhanced PVE-corrected PET images produced as a result of the developed algorithm considering both brain and oncology applications. Such enhanced images may allow a more accurate lesion delineation providing solutions in clinical PET applications of increasing importance such as response to therapy studies and use for radiotherapy treatment planning. In the whole-body PET image, indeed, the improved visual delineation of the lesion came with an increase of intensity leading to a 16% increase of lesion-to-lung ratio. As far as brain PET is considered, the PVE correction induced compensation for both spill-in and spillover effects. The intensity in the white matter decreased in the benefit of grey matter acting as a redistribution of activity from an overestimated area to an underestimated one. The accuracy and the precision of these quantitative results are difficult to evaluate but the global behaviour of the correction is in accordance with an expected 'inverse' cross-contamination.

However, the improvement in grey matter uptake seems to be lower than that obtained in simulated images as described by Quarantelli *et al* (2004). Even if the comparison of results from real and simulated images is questionable, this point justifies some discussion. First, the proposed PVE correction operates in 2D since the wavelet transform that we use (the 'à trous' algorithm) is restricted to 2D images. This is a notable weakness but a potential solution would consist in taking into account the 3D nature of PET images, for example by performing the 'à trous' algorithm in coronal and sagittal slices in addition to the transverse plane. In the same way, the 16% increase of lesion-to-lung ratio in the whole-body PET image may appear moderate. Potential improvement in this domain could be foreseen by performing the multiresolution analysis in a limited ROI surrounding the lesion instead of the whole image containing very different tissues and organs. The simple and global linear

model that was presented in this study (allowing us to build lacking details of the PET image from the details of the CT or MRI images, see section 2.2) may indeed not be adapted to very extended regions presenting with heterogeneous structures, typically like abdominal images where clearly certain structures such as intestines appear differently in CT and PET. Therefore, although a simple and global model may be better adapted to a visual improvement alone, a quantitative study under clinical imaging conditions could potentially benefit from considering a more restricted field of view. In support of this statement, it is worth considering the results that were obtained in the second unfavourable configuration of synthetic images (figures 2(c) and (d)). In this latter case, an intense local artefact was observed after applying our algorithm, thus leading to a corrupted correction. However, if we operate our method in a limited area surrounding the given disc as demonstrated in figure 2(e), the algorithm results in the expected partial volume correction for the specified region of interest (figure 2(f)). Obviously, the remainder of the image is not PVE corrected. It is also important to note that the same kind of artefact as that shown in figure 2(d) will appear in a corrected PET image, considering the use of the global linear model, when a given tissue is present in the CT and not in the PET. A potential solution to this particular case will be the use of a restricted field of view for the application of the global model in combination with the introduction of a measure of similarity between the wavelet components of the high- and low-resolution images within the particular ROI. However, one must keep in mind that the aim is to correct for PVE in a part of the PET image corresponding to an actual tissue of interest. If there is no structure of interest present in PET, there is no associated interest in PVE correction.

Finally, only circular and homogeneous lesions were tested in the synthetic and simulated images that were presented in this paper. In clinical practice concerning whole-body imaging for example, lesions can obviously be of different shape and have non-uniform activity concentration. Here again, a more sophisticated model rather than the linear model used in the implemented algorithm in combination with its application in a more limited region may be more appropriate.

5. Conclusion

In this paper, a novel technique to correct emission tomography for partial volume effects using multiresolution analysis has been presented. The advantages of this approach are three-fold and can be summarized as follows. According to various tests on synthetic images, the efficiency of the correction proved to be as good as a reference method based on regional spread functions. On the other hand, the wavelet-based correction leads to better results in simulated images. In addition, and contrary to the RSF method, the process allows enhancing the images themselves to perform further processing, without any time-consuming step of ROI delineation. Finally, images of any kind of tissues, organs, functions and metabolisms are likely to be corrected, provided an anatomical image of the same object is available and correctly aligned. A simple linear and global link between the wavelet coefficients of the emission image and those of the anatomical one has been defined in this study. Under certain imaging conditions, a local model, defining only a limited area around the tissue of interest, may be more appropriate and will be considered in future developments.

Acknowledgment

This work has been financially supported by the Region of Brittany.

Appendix

$$5 \times 5 \text{ bicubic spline : } \begin{bmatrix} 1/256 & 1/64 & 3/128 & 1/64 & 1/256 \\ 1/64 & 1/16 & 3/32 & 1/16 & 1/64 \\ 3/128 & 3/32 & 9/64 & 3/32 & 3/128 \\ 1/64 & 1/16 & 3/32 & 1/16 & 1/64 \\ 1/256 & 1/64 & 3/128 & 1/64 & 1/256 \end{bmatrix} .$$

$$5 \times 5 \text{ linear interpolation : } \begin{bmatrix} 1/100 & 1/50 & 3/25 & 1/50 & 1/100 \\ 1/50 & 1/25 & 2/25 & 1/25 & 1/50 \\ 1/25 & 2/25 & 4/25 & 2/25 & 1/25 \\ 1/50 & 1/25 & 2/25 & 1/25 & 1/50 \\ 1/100 & 1/50 & 1/25 & 1/25 & 1/100 \end{bmatrix} .$$

$$3 \times 3 \text{ linear interpolation : } \begin{bmatrix} 1/16 & 1/8 & 1/16 \\ 1/8 & 1/4 & 1/8 \\ 1/16 & 1/8 & 1/16 \end{bmatrix} .$$

$$3 \times 3 \text{ low scale : } \begin{bmatrix} 1/172 & 1/86 & 1/172 \\ 1/86 & 160/172 & 1/86 \\ 1/172 & 1/86 & 1/172 \end{bmatrix} .$$

References

- Ardekani B A, Braun M, Hutton B F, Kanno I and Iida H 1996 Minimum cross-entropy reconstruction of PET images using prior anatomical information *Phys. Med. Biol.* **41** 2497–517
- Ashburner J and Friston K J 2000 Voxel-based morphometry—the methods *Neuroimage* **11** 805–21
- Aston J A, Cunningham V J, Asselin M C, Hammers A, Evans A C and Gunn R N 2002 Positron emission tomography partial volume correction: estimation and algorithms *J. Cereb. Blood Flow Metab.* **22** 1019–34
- Baete K *et al* 2004 Evaluation of anatomy based reconstruction for partial volume correction in brain FDG-PET *Neuroimage* **23** 305–17
- Bencherif B, Stumpf M J, Links J M and Frost J J 2004 Application of MRI-based partial-volume correction to the analysis of PET images of mu-opioid receptors using statistical parametric mapping *J. Nucl. Med.* **45** 402–8
- Boussion N, Cinotti L, Barra V, Ryvlin P and Mauguire F 2003 Extraction of epileptogenic foci from PET and SPECT images by fuzzy modeling and data fusion *Neuroimage* **19** 645–54
- Caldwell C B, Mah K, Skinner M and Danjoux C E 2003 Can PET provide the 3D extent of tumor motion for individualized internal target volumes? A phantom study of the limitations of CT and the promise of PET *Int. J. Radiat. Oncol. Biol. Phys.* **55** 1381–93
- Du Y, Tsui B M W and Frey E C 2005 Partial volume effect compensation for quantitative brain SPECT imaging *IEEE Trans. Med. Imaging* **24** 969–76
- Dutilleul P 1987 An implementation of the “algorithme à trous” to compute the wavelet transform *Congrès ondelettes et méthodes temps-fréquence et espace des phases (Marseille, France, 14–18 September)* pp 298–304
- Feuardent J, Soret M, de Dreuille O, Foehrenbach H and Buvat I 2003 Reliability of SUV estimates in FDG PET as a function of acquisition and processing protocols *IEEE Nuclear Science Symposium Conference Record* pp 2877–81
- Frouin V, Comtat C, Reilhac A and Gregoire M C 2002 Correction of partial-volume effect for PET striatal imaging: fast implementation and study of robustness *J. Nucl. Med.* **43** 1715–26
- Hickey K T *et al* 2004 Assessment of cardiac wall motion and ejection fraction with gated PET using N-13 ammonia *Clin. Nucl. Med.* **29** 243–8
- Holdschneider R, Kronland-Martinet R, Morlet J and Tchamitchian P 1989 A real time algorithm for signal analysis with the help of the wavelet transform *Wavelets* (Berlin: Springer)
- Hutton B F and Lau Y H 1998 Application of distance-dependent resolution compensation and post-reconstruction filtering for myocardial SPECT *Phys. Med. Biol.* **43** 1679–93

- IEC Publication 61675-1 1998 *Radionuclide Imaging Devices—Characteristics and Test Conditions: Part 1. Positron Emission Tomographs*
- Kennedy J A, Azhari H, Frenkel A, Bar-Shalom R and Israel O 2005 PET/CT image data fusion in hybrid scanners *Abstract in Proc. 52nd SNM Annual Meeting (Toronto, 18–22 June)* p 167
- Kusano M L, Caldwell C B, Lanctot K L and Chow T W 2005 Evaluation and refinement of an MR-based brain PET partial volume correction method using the Zubal brain phantom *Abstract in Proc. 52nd SNM Annual Meeting (Toronto, 18–22 June)* p 452
- Lamare F, Turzo A, Bizais Y, Cheze Le Rest C and Visvikis D 2006 Validation of a Monte Carlo simulation of the Philips Allegro/Gemini PET systems using GATE *Phys. Med. Biol.* **at press**
- Li S and Shawe-Taylor J 2005 Comparison and fusion of multiresolution features for texture classification *Pattern Recognit. Lett.* **26** 633–8
- Luo R C, Yih C C and Su K L 2002 Multisensor fusion and integration: approaches, applications, and future research directions *IEEE Sensors J.* **2** 107–19
- Mallat S 1989 A theory for multiresolution signal decomposition: the wavelet representation *IEEE Trans. Pattern Anal. Mach. Intell.* **11** 674–93
- Matsuda H, Ohnishi T, Asada T, Li Z J, Kanetaka H, Imabayashi E, Tanaka F and Nakano S 2003 Correction for partial-volume effects on brain perfusion SPECT in healthy men *J. Nucl. Med.* **44** 1243–52
- Meltzer C C, Kinahan P E, Greer P J, Nichols T E, Comtat C, Cantwell M N, Lin M P and Price J C 1999 Comparative evaluation of MR-based partial-volume correction schemes for PET *J. Nucl. Med.* **40** 2053–65
- Meltzer C C, Leal J P, Mayberg H S, Wagner H N Jr and Frost J J 1990 Correction of PET data for partial volume effects in human cerebral cortex by MR imaging *J. Comput. Assist. Tomogr.* **14** 561–70
- Muller-Gartner H W, Links J M, Prince J L, Bryan R N, McVeigh E, Leal J P, Davatzikos C and Frost J J 1992 Measurement of radiotracer concentration in brain gray matter using positron emission tomography: MRI-based correction for partial volume effects *J. Cereb. Blood Flow Metab.* **12** 571–83
- Pretorius P H and King M A 2004 Spillover and partial volume effect compensation to enhance the diagnostic accuracy of cardiac SPECT perfusion imaging *Abstract in Proc. 51st SNM Annual Meeting (Philadelphia, 18–23 June)* p 399
- Quarantelli M *et al* 2004 Integrated software for the analysis of brain PET/SPECT studies with partial-volume-effect correction *J. Nucl. Med.* **45** 192–201
- Ranchin T and Wald L 2000 Fusion of high spatial and spectral resolution images: the ARSIS concept and its implementation *Photogramm. Eng. Remote Sens.* **66** 49–61
- Reader A J *et al* 2002 One-pass list-mode EM algorithm for high-resolution 3D PET image reconstruction into large arrays *IEEE Trans. Nucl. Sci.* **49** 693–9
- Rota Kops E and Krause B J 2005 The influence of filtered back-projection and iterative reconstruction on partial volume correction in PET *Nuklearmedizin* **44** 99–106
- Rousset O G, Deep P, Kuwabara H, Evans A C, Gjedde A H and Cumming P 2000 Effect of partial volume correction on estimates of the influx and cerebral metabolism of 6-[(18)F]fluoro-L-dopa studied with PET in normal control and Parkinson's disease subjects *Synapse* **37** 81–9
- Rousset O G, Ma Y and Evans A C 1998 Correction for partial volume effects in PET: principle and validation *J. Nucl. Med.* **39** 904–11
- Som S, Hutton B F and Braun M 1998 Properties of minimum cross-entropy reconstruction of emission tomography with anatomically based prior *IEEE Trans. Nucl. Sci.* **45** 3014–21
- Somayajula S, Asma E and Leahy R 2005 PET image reconstruction using anatomical information through mutual information based priors *IEEE Nuclear Science Symposium Conference Record* pp 2722–6
- Soret M, Riddell C, Hapdey S and Buvat I 2001 Biases affecting tumor uptake measurements in FDG-PET *IEEE Nuclear Science Symposium Conference Record* pp 2119–23
- Starck J L, Murtagh F and Bijaoui A 1998 *Image Processing and Data Analysis: The Multiscale Approach* (Cambridge: Cambridge University Press)
- Vogel W V, Oyen W J, Barentsz J O, Kaanders J H and Corstens F H 2004 PET/CT: panacea, redundancy, or something in between? *J. Nucl. Med. Suppl.* **1** 45 15S–24S
- Wells W M III, Viola P, Atsumi H, Nakajima S and Kikinis R 1996 Multi-modal volume registration by maximization of mutual information *Med. Image Anal.* **1** 35–51
- Wen C Y and Chen J K 2004 Multi-resolution image fusion technique and its application to forensic science *Forensic Sci. Int.* **140** 217–32
- Zaidi H, Ruest T, Schoenahl F and Montandon M L 2005 Comparative assessment of brain MR image segmentation algorithms and their impact on partial volume effect correction in PET *Abstract in Proc. 52nd SNM Annual Meeting (Toronto, 18–22 June)* p 454

See endnote 1

Endnotes

(1) Author: Please update reference 'Lamare *et al* (2006)'.

Reference linking to the original articles

References with a volume and page number in blue have a clickable link to the original article created from data deposited by its publisher at CrossRef. Any anomalously unlinked references should be checked for accuracy. Pale purple is used for links to e-prints at arXiv.