Improved 3D Correction for Partial Volume Effects in Brain SPECT

Laurent Itti,[1]* Linda Chang,[1] Thomas Ernst,[1,2] and Fred Mishkin[3]

[1]*Department of Neurology, Harbor-UCLA Medical Center, Torrance, California 90509*
[2]*Department of Radiology, Harbor-UCLA Medical Center, Torrance, California 90509*
[3]*Department of Radiology, Division of Nuclear Medicine, Harbor-UCLA Medical Center,*
Torrance, California 90509

Abstract: An improved method for correction of partial volume effects (PVE) in brain SPECT is proposed. It is fully three-dimensional, does not require particular patient positioning, and works with scans only partially covering the brain. The location of functionally inactive brain regions (primarily cerebrospinal fluid) is extracted from high-resolution MRI. An automatic 3D registration algorithm then determines the geometric transformation between MRI and SPECT. Correction consists of: 1) counting the volumetric active/inactive ratio in each volume element of the functional scan using the measured SPECT point spread function; 2) correcting the functional measures according to these ratios; 3) fusing functional and anatomical information at the resolution of MRI. Quantitative validation was performed using a phantom containing a test region in which multiple parallel acrylic plates thinner than SPECT resolution created high PVE, as well as a large reference region not suffering from PVE. Reference activity was recovered in the test region with an accuracy of 1–3%. The method was applied to clinical images demonstrating a combination of hypoperfusion and cortical atrophy. The composite anatomical-functional corrected images, in which the main sulci are visible, yield better differentiation between decreased function and focal atrophy. *Hum. Brain Mapping 5:379–388, 1997.*   © 1997 Wiley-Liss, Inc.

Key words: partial volume effects; MRI; SPECT; brain atrophy; image processing

INTRODUCTION

Functional neuroimaging techniques such as positron emission tomography (PET) and single photon emission computed tomography (SPECT) have modest spatial resolutions of typically 5 to 15 mm; therefore, they suffer greatly from partial volume effects (PVE). These techniques image a local spatial average of brain function (e.g. metabolism, blood flow). Consequently, they may not accurately represent neuroanatomical structures smaller than approximately 2.5 times the imaging resolution, as has been demonstrated using hot spheres in a cold background [Hoffman et al., 1979; Mazziota et al., 1981]. As a result, poor differentiation is obtained between functional deficits (e.g. hypoperfused brain tissue) and morphological deficits (e.g. focal cerebral atrophy, in which brain cells are lost and replaced by non-perfused cerebrospinal fluid). This ambiguity in the functional images is particularly problematic when studying diseases in which both atrophy and functional deficits occur, such

as neurodegenerative brain disorders and focal brain infarctions. Also, PVE decrease the sensitivity of functional imaging for focal activation studies essential to brain mapping [Mazziotta et al., 1981].

Several methods have been proposed to correct *a posteriori* for PVE in PET imaging [Herscovitch et al., 1986; Labbe et al., 1996; Meltzer et al., 1990a; Müller-Gärtner et al., 1992; Videen et al., 1988]. These methods use high-resolution anatomical information extracted from computed tomography (CT) or magnetic resonance imaging (MRI) scans to estimate PVE. The assumption is made that only trace amounts of radiotracer are present in the cerebrospinal fluid (CSF)[Chawluk et al., 1987; Herscovitch et al., 1986]. A binary map of the CSF is extracted from the anatomical images, low-pass filtered, and down-sampled to simulate a PET scan. Correction of the initially measured PET activity is then possible by simple arithmetic operations using the low-resolution CSF map. Some methods not only assume zero PET activity in the CSF, but also uniform PET activity in the white matter (WM), and separately consider the gray matter (GM), WM and CSF compartments [Labbe et al., 1996; Müller-Gärtner et al., 1992]. Validation in these studies was performed with relatively simple phantoms, like spheres or large, homogeneous objects [Meltzer et al., 1990a; Müller-Gärtner et al., 1992]. One important limitation of these methods originally was to require that MRI and PET acquisitions be performed with similar head orientations, using special head positioning or marking devices [Meltzer et al., 1990b; Müller-Gärtner et al., 1992]. Subsequently, image registration between MRI and PET was applied prior to correction [Meltzer et al., 1995], which presents the drawback of typically introducing unnecessary losses of accuracy due to image resampling.

Based on these methodologies, we propose an improved PVE correction method, which more extensively utilizes the high-resolution anatomical information and explicitly addresses the problem of 3D image registration. It was designed to also work with scans only partially covering the brain, and was validated with a specifically designed phantom demonstrating large PVE. The proposed approach consists of explicitly reproducing the diffusion process that results from the camera's low-pass characteristics, and accounts for possibly different orientations of the MRI and SPECT image lattices.

MATERIALS AND METHODS

Image acquisition and processing hardware

MRI studies were performed with a 1.5 Tesla General Electric SIGNA imaging system (General Electric Medical Systems, Milwaukee, WI), in axial, sagittal or coronal orientations. No constraint was imposed on head position, head orientation or scanning volume geometry. All MR sequences routinely used at our institution could be used for PVE correction. These included various T_1-weighted (typically TE/TR = 14/600, 5 mm slice thickness), T_2-weighted (typically TE/TR = 102/4000, 5 mm slice thickness), T_1-weighted with contrast agent, inversion recovery (typically TE/TI/TR = 32/120/4000, 3.5 mm slice thickness), and multi-echo sequences.

SPECT studies were performed with a Shimadzu Headtome II dedicated brain unit (Shimadzu, Kyoto, Japan). It provided eleven non-uniformly spaced slices (four bed shifts with three fixed ring detectors, one of the 12 resulting slices being redundant), with a spatial resolution of approximately 1 cm in plane and a 2 cm slice thickness [Kanno et al., 1980]. Spatial oversampling performed during image reconstruction yielded $3.3 \times 3.3 \times 20$ mm³ voxels. HMPAO (Technetium-99m hexamethyl propyleneamine oxime; Ceretec, Amersham, Arlington Heights, IL) was used to measure relative cerebral blood flow [Masdeu et al., 1994]. Inhaled ¹³³Xe gas (Xenon-133; Dupont, Billerica, MA) was used for absolute blood flow measurements and calibration of HMPAO images to absolute values (in mL/100 g tissue/min)[Kanno and Lassen, 1979].

The method was developed as an extension of a 3D multimodality registration method [Itti et al., 1996, 1997]. The corresponding programs were written on a Sun SPARCStation 10 (Sun Microsystems, Mountain View, CA) and on a DEC Alpha workstation (Digital Equipment Corporation, Nashua, NH), using the AVS 5 graphical interface and object-oriented development tool (Application Visualization System, Waltham, MA) and specifically written (portable language C) AVS modules and networks of modules.

Model assumptions

The brain is modeled as two separate compartments: one active compartment eliciting radioactive tracer uptake (brain tissue, muscles [Costa et al., 1993]; see also Ceretec package insert) and one inactive compartment composed of CSF [Chawluk et al., 1987; Herscovitch et al., 1986], bone and air. The relationship between the number of γ-photons emitted by brain tissues and the activity assigned to a given voxel is globally modeled by the 3D point-spread function (PSF) h of the machine (a function of 3D spatial coordinates *x, y* and *z*). The PSFs of current cameras differ from the ideal Dirac distribution because of problems including scatter, finite detector size, and

relatively small number of detectors. The value of a SPECT voxel is influenced by all radioactivity distributed over the 3D extent of *h* centered at that voxel. The goal of PVE correction is to estimate *a posteriori* the concentration of tracer per unit volume of actual brain tissue, from the observed concentration per total unit volume.

Following the formulation introduced by Müller-Gärtner et al. [1992], the observed tracer concentration I_{obs} (which is a function of the 3D *x, y* and *z* image coordinates) may be written as a spatially weighted average of the tracer concentration from the active compartment I_{act} and the inactive compartment I_{inact} (also functions of 3D image coordinates):

$$I_{obs} = \frac{(I_{act}X_{act}) \otimes h + (I_{inact}X_{inact}) \otimes h}{(X_{act} + X_{inact}) \otimes h} \qquad (1)$$

where $\otimes$ is the 3D convolution operator. X_{act} and X_{inact} are characteristic functions computed from MRI for the active and inactive compartments (i.e. they are 1 inside the corresponding compartment and 0 outside; while previous methods used binary characteristic functions, a continuous range between zero and unity may be used). $X_{act} = 1 - X_{inact}$ in our two-compartment model. If *h* has unity integral, the denominator of eq. (1) is ideally unity (and was not present in Müller-Gärtner et al.'s formula). However, the denominator actually computed from MRI is not unity near the edges of the MRI 3D image matrix, where only partial information about X_{act} and X_{inact} is available: the exact convolution of $X_{act} + X_{inact}$ by *h* can be computed only at the locations where *h* and $X_{act} + X_{inact}$ fully overlap. As detailed below, this is particularly problematic for fully 3D processing if only partial, and possibly different, brain coverages are available in MRI and SPECT (e.g. partial sagittal MRI and partial axial SPECT); zero-padding of the images to obtain full overlap would underestimate the amount of correction needed near the volume edges. Since it is assumed in the model that the inactive compartment does not emit any radioactivity, the second term in the numerator of eq. (1) is zero. Finally, following Müller-Gärtner et al.'s remark that, although not true in general, the equality

$$(I_{act}X_{act}) \otimes h = I_{act}(X_{act} \otimes h) \qquad (2)$$

may yield a good estimate of the activity in the active compartment, we rewrite eq. (1) as

$$I_{act} = \frac{(X_{act} + X_{inact}) \otimes h}{X_{act} \otimes h} I_{obs} \qquad (3)$$

This equation reduces to the less general result obtained by Müller-Gärtner et al. as long as we remain farther than half the support size of *h* from the edges of the MRI volume. The next sections detail how this result may be used in a fully three-dimensional correction algorithm with specific processing at volume edges and integrated 3D registration.

Estimation of the SPECT point-spread function *h*

The in-plane 2D PSF of the SPECT camera was measured with a phantom modeled after existing prototypes [Hoffman et al., 1979]. Five acrylic tubes of known internal diameters were enclosed in a larger acrylic cylinder simulating a head. The five tubes were filled with a mixture of water, ^{99m}Tc as pertechnetate and copper sulfate while the large cylinder was filled with water. SPECT was acquired with the phantom's tubes orthogonal to the image planes. MRI was acquired with a T_1-weighted sequence showing high contrast between the copper sulfate solution inside the five tubes and the surrounding water. For simplicity, the PSF is assumed to be uniform across the SPECT field of view, although the method could accommodate for non-uniform PSF with no conceptual modification.

Recovering the PSF from these two imaging measurements is possible by using the fact that spatial convolution is equivalent to point-by-point multiplication in the Fourier domain. Consequently, the 2D in-plane SPECT PSF was determined by dividing (point-by-point) 2D Fourier-transformed sections of a given tube in SPECT by the corresponding Fourier-transformed sections in MRI, and fitting 2D Gaussian functions to the envelopes of the inverse Fourier-transformed results (although unstable in general with respect to image noise, this method yielded satisfactory results for this particular purpose). The 2D PSF could be expressed as a pair of orthogonal 1D Gaussian functions, *xPSF* and *yPSF*, both with a full width at half maximum (FWHM) of 12.3 mm. For the PSF orthogonal to the slice planes, *zPSF*, the manufacturer's specifications for the ^{99m}Tc collimator yielded a 1D Gaussian model with 20 mm FWHM [Kanno et al., 1980].

Extraction of anatomical information

In order to accurately account for all inactive areas, and not only for CSF, a double-window thresholding method is applied to the MR images. For our inversion-recovery MR images, a low intensity window extracts voxels which belong to bone, air-filled cavities and air surrounding the subject's head; a high intensity win-

dow extracts voxels which belong to CSF and fluids in the orbits. The low intensity window is defined as values between 0 and a manually selected low threshold and the high window as values between a manually selected high threshold and the maximum image intensity. Voxels with intensities in neither window are assigned to the active compartment (and typically include brain and muscles). Median filtering is applied in each window to eliminate noisy isolated pixels. In order to derive fractional inactivity values (between 0 and 1) for each MR voxel, we used simple linear interpolations between three intensity levels: pure fluid (the maximum intensity in the images), pure air/bone (intensity zero), and average tissue (computed as the average intensity over a large area of brain tissue). Voxels in the high intensity window represent a mixture of fluid and tissue; their inactivity ratio is linearly interpolated by comparing their intensity to the pure fluid (inactivity ratio of 1) and average tissue (inactivity ratio of 0) levels. Voxels in the low intensity window contain a mixture of air or bone and tissue; their inactivity ratio is linearly derived from the air/bone and average tissue levels. Finally, voxels in neither window may contain either tissue or a mixture of fluid and air/bone. To determine if one such voxels belongs to the second category, its 3×3 in-plane neighborhood is examined. When at least two of the eight neighbors belong to the low window and two belong to the high window, the considered voxel is assumed to lie at the interface between fluid and bone or air; its inactivity ratio is set to 1. Voxels classified as tissue receive an inactivity ratio of zero. The final map of inactive regions, X_{inact}, encodes the partial volume of inactive space in every MR voxel using values between 0 and 1.

Image registration

The geometric transformation between the MRI and SPECT volumes (3D translation and 3D rotation) is automatically obtained from a 3D registration method [Itti et al., 1996, 1997]. The 3D surface of the brain is automatically extracted from MRI and SPECT using mathematical morphology. Both surfaces are then automatically matched, using an iterative minimization algorithm and surface-fit criterion specifically designed to allow registration of scans only partially covering the brain. The accuracy of the registration method was found to be better than imaging resolution (average 0.07 deg rotation error and 0.19 mm translation error in registering an MRI with 0.9375 × 0.9375 × 5 mm³ voxels to 11 rotated and translated versions of itself [Itti et al., 1997]; 1.33 deg and 1.62 mm

respectively after low-pass filtering, subsampling to SPECT resolution and adding strong Poisson noise to the rotated and translated volumes). No reslicing is performed at this stage, in order to avoid losses of information typically associated with image resampling to an image lattice with different 3D axis orientations (unless resampling were performed to a prohibitively high resolution satisfying a 3D Nyquist criterion). The registration parameters describing the 3D geometric transformation between the original MRI and SPECT datasets are used directly by the correction method with the original datasets.

Correction for partial volume effects

Because the original MRI and SPECT volumes are not scanned in the same orientation, it would be incorrect to convolve the MR images by the 3D PSF of SPECT, which is non-isotropic in space. The correction method can be summarized as follows. First, the amount of active and inactive volume is determined for each SPECT voxel. Second, the measured SPECT values are corrected for partial volume effects using the active/inactive ratios. Third, the SPECT volume is reconstructed with the orientation and matrix size of the MRI volume, using the inactivity map from MRI to constrain the localization of the corrected SPECT activity.

The first step in the algorithm estimates the partial volume of active and inactive regions within every SPECT voxel, using two maps, $\hat{X}_{act}$ and $\hat{X}_{inact}$, with the matrix size and orientation of the SPECT volume. The coordinates P of each MR voxel are transformed into (real-number) SPECT coordinates using the geometric registration transformation T. Around $T(P)$, a subvolume of the SPECT volume is examined, consisting of the four closest SPECT slices (two above and two below) and the 8×8 pixels in-plane neighborhood. The SPECT PSF, centered around $T(P)$, is negligible outside this sub-volume (Fig.1a). For each SPECT voxel S in the sub-volume, real-number coordinate differences dx, dy and dz between $T(P)$ and the center point of S are computed. $\hat{X}_{act}$ (S) and $\hat{X}_{inact}$ (S) are incremented by the contribution of P:

$$\hat{X}_{inact}(S) \leftarrow \hat{X}_{inact}(S) + xPSF(dx)\,yPSF(dy)\,zPSF(dz)\,X_{inact}(P)$$

$$\hat{X}_{act}(S) \leftarrow \hat{X}_{act}(S) + xPSF(dx)\,yPSF(dy)\,zPSF(dz)\,(1 - X_{inact}(P))$$

After all the voxels from the MRI volume have been examined, and $\hat{X}_{act}$ and $\hat{X}_{inact}$ hence respectively contain the effective active and inactive volumes influenc-
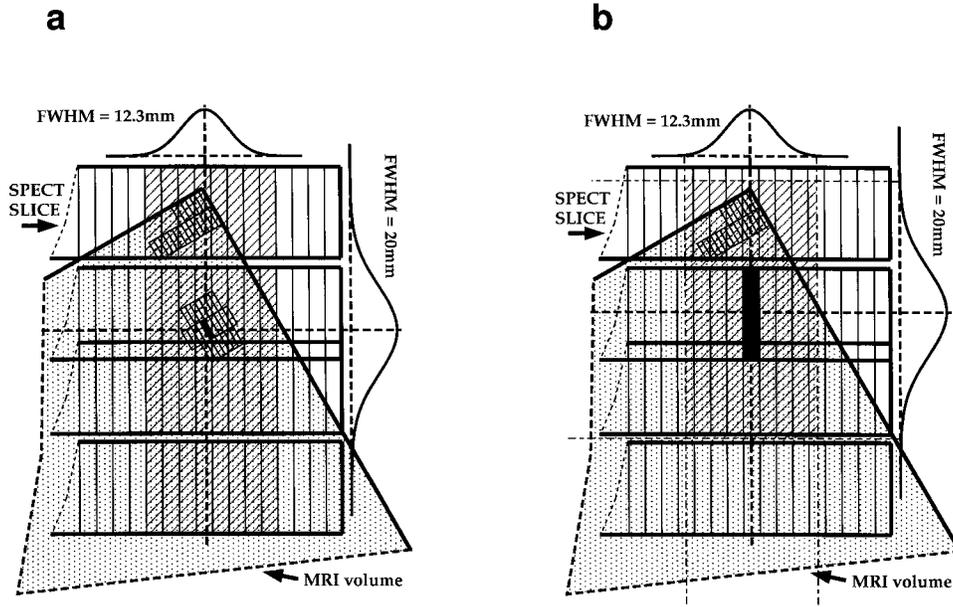
a

b

Figure 1.

a: Sub-volume of the SPECT volume explored by our algorithm (striped area) for a given MR voxel (black). The active/inactive ratio of each SPECT voxel in this $8 \times 8 \times 4$ sub-volume receives a contribution from the MR voxel of interest, weighted by the SPECT point spread functions *xPSF, yPSF* and *zPSF*. Note the non-uniform spacing of our SPECT slices. **b:** Sub-volume (striped area) for a given SPECT voxel (black) for which *xPSF, yPSF* and *zPSF* are not negligible. The value of the SPECT voxel of interest is influenced by local active/inactive distributions over this whole sub-volume. Our algorithm is designed to obtain accurate correction when the MRI only partially overlaps with this sub-volume, as shown here.

ing (accounting for diffusive effects) every SPECT voxel, i.e. $X_{act} \otimes h$ and $X_{inact} \otimes h$ far from the scanning edges.

The second step of the algorithm consists of correcting the measured SPECT activity, using the information in $\hat{X}_{act}$ and $\hat{X}_{inact}$. For each SPECT voxel S, its activity *SPECT(S)* is multiplied by the total/active ratio estimated from MRI according to eq. (3), yielding the corrected SPECT activity, *correctedSPECT(S)*:

$$correctedSPECT(S) \leftarrow \frac{\hat{X}_{act}(S) + \hat{X}_{inact}(S)}{\hat{X}_{act}(S)} SPECT(S)$$

If $\hat{X}_{act}$ is zero, the corrected value of S is set to zero (no active volume in S yields no activity). The reason for keeping two apparently complementary maps $\hat{X}_{act}$ and $\hat{X}_{inact}$ is to obtain accurate corrections at the edges of the scanned volumes when the SPECT and MRI volumes overlap only partially. Far from the MRI edges, the ratio by which each SPECT voxel is multiplied corrects for the difference between the total volume influencing a SPECT voxel and the actual volume of brain tissue influencing that voxel. Close to

the MRI edges, this difference is estimated from the partial information available from MRI (Fig.1b; this case typically occurs with sagittal MRI partially covering the brain), and the correction ratio becomes that of the volume influencing the SPECT voxel and seen in MRI by the volume of brain tissue influencing the same SPECT voxel and seen in MRI. This avoids underestimation of the correction ratio that would be obtained by simply assuming that MRI values were zero outside the scanned volume.

The third and last step of the algorithm consists of reslicing the corrected SPECT volume to the matrix size and orientation of the MRI volume. For each MR voxel P, the corrected SPECT activity at the location of $T(P)$ is obtained with trilinear interpolation from *correctedSPECT*. This activity is multiplied by the coefficient of activity in the MR voxel P, yielding the final composite resliced SPECT volume *reslicedSPECT*:

$$reslicedSPECT(P) \leftarrow X_{act}(P)\,correctedSPECT(T(P))$$

Regions in which only active brain tissue is present will directly reflect the corrected SPECT values at the

resolution of SPECT. However, in all the regions influenced by the presence of inactivity, the reconstructed SPECT values will be modulated, at the resolution of MRI, by the active/inactive ratio of each MR voxel. As a result, sulcal patterns appear in the final composite corrected SPECT images. The final display is masked using the brain contours automatically extracted from MRI, which eliminates extracranial perfused regions.

Validation phantom

A phantom was designed to assess the accuracy of the method when active and inactive regions are intimately interlaced, and when the assumption of two separate and uniform compartments is well verified. The phantom consists of a parallelepiped of acrylic filled with a mixture of water and ^{99m}Tc as pertechnetate, and is subdivided into two main regions (Fig. 2). At the top, a large control region allows for the measurement of SPECT activity when no PVE is present. Measuring the SPECT activity in the center of this region yields a pure signal reference to which the corrected values can be compared. In the test region at the lower portion of the phantom, we placed 1/4-in-thick plates of acrylic spaced by 1/4 in (6.35 mm). The space between the plates is filled with the same solution as the control region. The spacing between two of these plates is slightly larger (+3 mm) to provide a locally increased active/inactive ratio.

The phantom was imaged using SPECT and our standard 5 mm slice thickness T_2-weighted MR sequence (TE/TR = 102/4000; for this two-compartment object, any imaging sequence showing strong signal in the water solution and no signal in the acrylic or air would yield similar results). SPECT and MRI were automatically registered using the convex hull of the outer surface of the phantom. SPECT activity was measured for several slices inside entire cross-sections of the phantom and inside smaller, manually drawn regions of interest (ROIs). The same ROI shape was used for five consecutive slices in each region of the phantom.

VALIDATION AND RESULTS

Results obtained with the validation phantom are presented in Table I and Figure 3. The reference activity measured inside the control ROI before correction was assigned the value 100 and used to calibrate all other measures. PVE correction did not significantly alter the activity in the control ROI (corrected
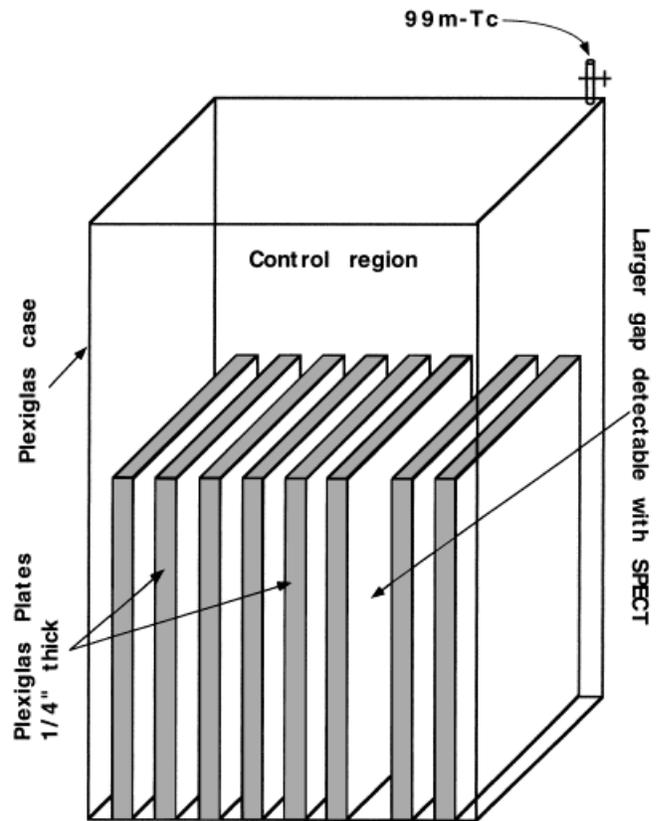

Figure 2.

Schematic drawing of the validation phantom. The large control region (top) allows for the acquisition of a reference SPECT count. The parallel Plexiglax plates (bottom), which are thinner than SPECT resolution, introduce large partial volume effects. The phantom was filled with water, 99m-Tc, and copper sulfate for SPECT and MRI.

value 100.5), indicating that this ROI was far enough from the phantom edges and was influenced minimally by PVE. The initially measured activity in the test region of the phantom was approximately 55% the reference activity, both inside the test ROI and inside

TABLE I. Quantitative validation of the correction method using a dedicated phantom demonstrating high PVE

	Initial activity	Corrected activity
Control ROI	100 ± 4	100.5 ± 4
Test ROI	58.5 ± 7	97.1 ± 8
Test full slice	55.8 ± 11	100.1 ± 13

All measurements (mean ± standard deviation) were normalized to the initial activity in the control ROI.

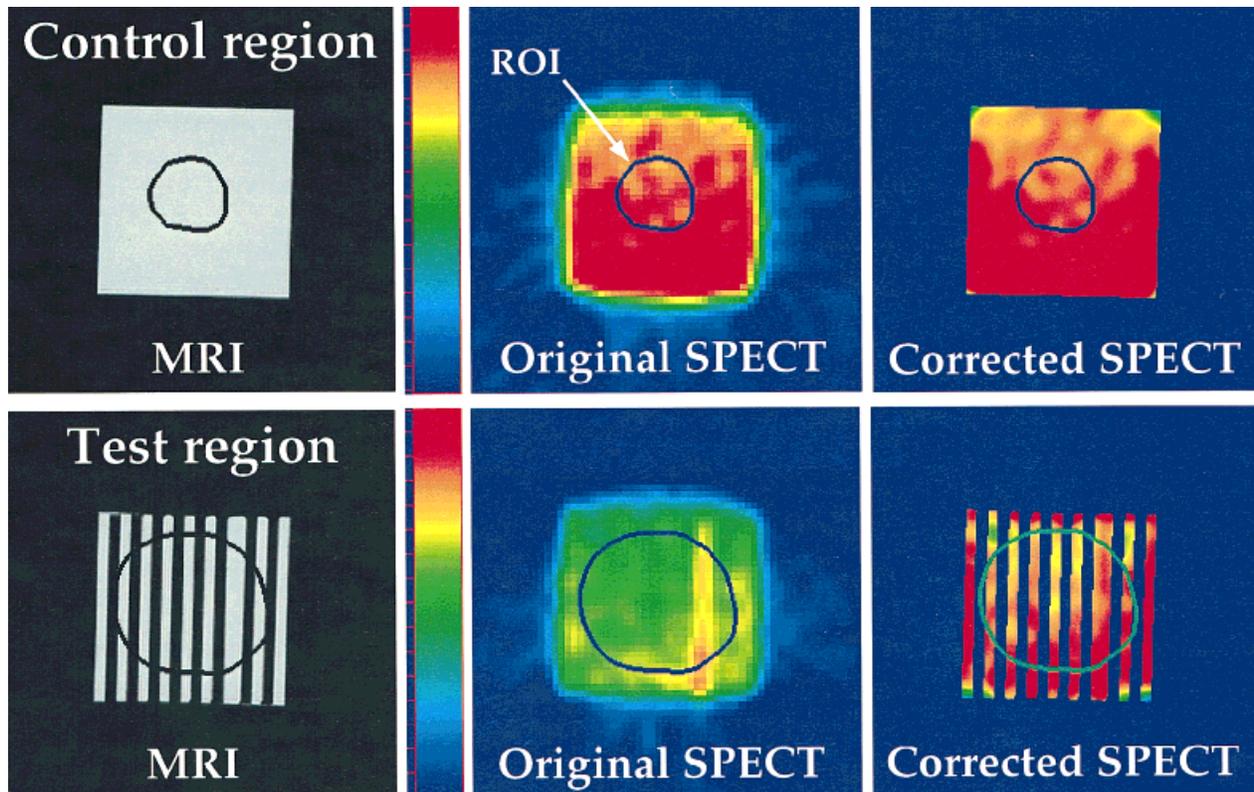**Figure 3.**

Transaxial MRI ($0.9875 \times 0.9875 \times 5$ mm³ voxels) and SPECT images of the validation phantom. SPECT activity inside the control ROI (top row) was virtually unaffected by the correction process. SPECT activity measured in the test region (bottom row) was initially underestimated by a factor of two due to partial volume effects. After correction, activity in the test region matched the reference activity within 3% inside the test ROI and less than 1% for the full slice.

the sections of the whole extent of the phantom. This corresponds to the ratio of active/inactive volume in the test region, which also represents approximately 55% due to the larger gap between two of the plates. Inside the test ROI, the reference activity was recovered with an accuracy within 3%. When the full cross-sectional area of the phantom was used, the corrected activity matched the reference activity within less than 1%, despite relatively large variances in these measurements. In general, corrected values appear underestimated by a few percents deep inside the phantom, and overestimated by a few percents at the edges of the phantom (Fig. 3).

Clinical applications

Examples of clinical applications of the method are presented in Figure 4. Locally, brain regions containing large sulci yielded uncorrected SPECT activity that resembled the activity in truly hypoperfused cortex.

This ambiguity in the interpretation of the SPECT images was eliminated after correction for PVE, because sulci appeared in the composite anatomical-functional corrected images. In high transaxial slices, the global presence of large volumes of CSF throughout an image would appear as global hypoperfusion from the initial SPECT data. After correction, however, it can be seen that only minimal regional hypoperfusion may be present. Total correction and reslicing time was approximately 7 min using a $256 \times 256 \times 40$ MRI dataset on a Sun SPARCStation 10, and approximately 2.5 min on a DEC Alpha workstation. This is only approximately five times more than the time necessary to perform a simple reslicing of a SPECT volume.

DISCUSSION

The basis of our method is similar to that previously suggested: a spatial map of CSF is extracted from a structural scan, which is then used to correct *a posteriori*
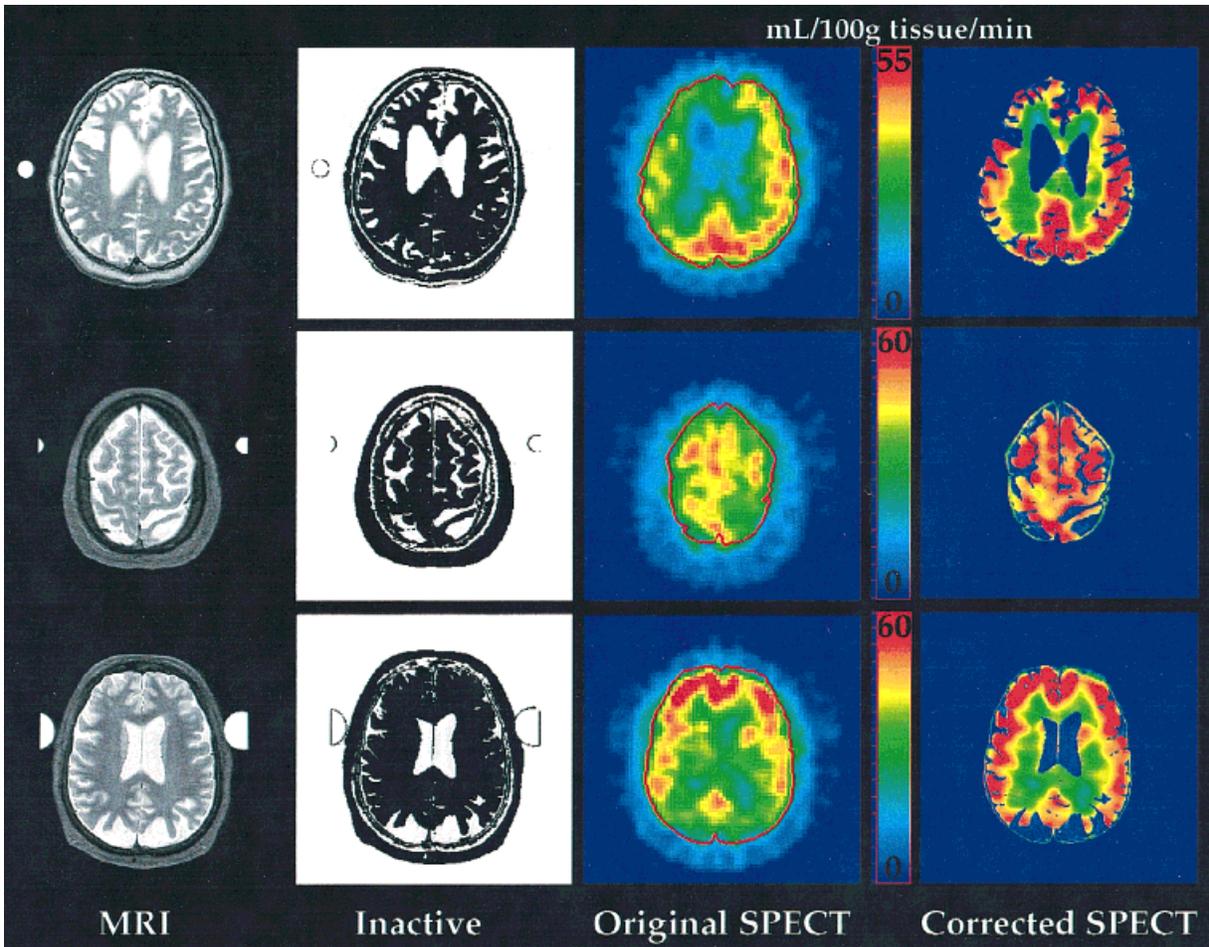
Figure 4.

Examples of clinical applications of the correction method: top row shows images from a patient with fronto-temporal dementia. Note marked atrophy on MRI and apparent severe hypoperfusion on the original SPECT in the frontal lobes; however, the corrected SPECT shows only minimal hypoperfusion in the remaining cortex. Middle row demonstrates that, at a higher transaxial slice, the presence of numerous large sulci may be misinterpreted for global hypoperfusion. The bottom row shows severe atrophy in MRI and apparent hypoperfusion in SPECT in the parietal regions of a patient with Alzheimer's disease; the corrected SPECT shows both the atrophy and mild hypoperfusion in these regions.

for PVE in the functional images. We discuss in this section some of the major differences and advances in our method.

First, the accuracy of the method was improved by not only accounting for the presence of CSF, but also for all other regions supposedly without any radioactivity. Local quantitative correction errors were created in a previous version of the method that only used CSF (not shown). Such local errors have been reported by Meltzer et al. [1990b], who then eliminated all PET activity outside the brain in a preprocessing step. This however presents the danger of introducing correction artifacts by discarding peripheral active areas (e.g. glands, muscles) close to the brain. It was important to address the problem of peripheral tracer uptake in our method designed for HMPAO SPECT, which shows significant uptake in muscle [Costa et al., 1993], while a variety of PET tracers do not show much muscle uptake and suffer less from it. We also account for partial volumes of inactivity within individual MR voxels, for more accurate local correction than when binary approximations are performed to create CSF maps. More sophisticated methods have been developed to extract CSF from MR images [Paley et al., 1994]. While our simple thresholding method yielded satisfactory results with our high-contrast MR sequences (e.g. inversion-recovery; Fig. 4), it may not be adequate for other sequences.

Second, our method is integrated into the necessary 3D registration process. This eliminates the need for specific patient positioning protocols during the anatomical and functional acquisitions as was previously reported [Meltzer et al., 1990a; Müller-Gärtner et al., 1992] or the need for prior registration. The explicit formulation adopted does not introduce unnecessary rounding errors, at the cost of increased computational expense when compared to convolutional methods. In our optimized implementation, execution time of only a few minutes, however, did not appear to be a limiting factor.

Third, our method explicitly addresses the problem that the region influencing a given SPECT voxel may be imaged only partially in MRI. This problem is solved by creating two separate maps $\hat{X}_{act}$ and $\hat{X}_{inact}$ for the partial active and inactive volumes within each SPECT voxel. This is particularly useful for the application of the method to three-dimensional sagittal and coronal MRIs, while previous methods were restricted to precisely positioned transaxial scans. Also, more accurate correction results are obtained in the extremal slices of transaxial volumes.

The composite anatomical-functional corrected images proved particularly informative for medical evaluation, because anatomical landmarks such as major sulci or the ventricules are visible. Corrected images contain slightly more information than corrected blood flow values; for example, although the SPECT camera was unable to resolve the individual acrylic plates of the validation phantom (Fig. 3, bottom row), the edges of these plates are resolved precisely in the corrected images.

The results obtained with the plate-filled phantom suggest a high quantitative accuracy of the correction method. The slight overestimation at the edges of the test region can be explained by minor spatial variations in the SPECT PSF between the center and the periphery of the imaged volume [Kim et al., 1992]. However, these errors seem to be comparable or below the limited detection accuracy (mainly due to poor statistical quality) of SPECT, as demonstrated by the relatively large variances of the initial SPECT measurements. Therefore, the quantitative accuracy of the method, with corrected values 1–3% from the values measured in the absence of PVE for the same radiotracer concentration, seems appropriate for SPECT. The phantom study further emphasized the large extent of PVE: although the control ROI was small and located deep inside the large control region (Fig. 3, top row), SPECT measurements inside this ROI were still slightly affected by PVE, which originated 3–4 cm away from the ROI, at the outer edges of the phantom.

Therefore, SPECT activity measured in most of the deep brain structures is likely to be affected by cortical atrophy.

An important additional source of degradation in quantitative SPECT measurements is photon attenuation. Several techniques have been proposed to correct for the effects of attenuation, and should be used in conjunction with our technique. These include pre-processing of the projection data, modifications of the back-projection process, iterative reconstruction techniques (see [Blokland et al., 1992] for a review) and the acquisition of a transmission map unique to each individual using radioactive line sources [Hashimoto et al., 1997; Tan et al., 1989]. Attenuation correction methods are particularly important for quantitative activity measurements in deep brain structures. With this respect, these methods are desirable complements to PVE correction methods which largely affect cortical areas but may leave the measured activity in small, deep brain structures unchanged.

Finally, our method could be extended to explicitly differentiate the gray and white matter compartments, as has been proposed for PET. The gain in accuracy obtained with such three-compartment models is not negligible, although it was found that a two-compartment model already corrects for most of the PVE [Meltzer et al., 1995]. However, we have not included any distinction between gray and white matter in our model because the distribution of radiotracer in the white matter seems unlikely to be uniform with the HMPAO tracer, while this is the central assumption for three-compartment models in PET. In HMPAO SPECT indeed, the lipophilic tracer entrapped in the brain tissue after the first passage of the intravenous injection [Costa et al., 1993] is likely to diffuse in the white matter fibers during the 30–45 min between injection and imaging.

CONCLUSIONS

Our PVE correction method addresses some of the problems encountered by previous methods. Higher accuracy is obtained by, first, a more precise definition of the inactive regions extracted from MRI; second, the integration of the correction into the 3D image registration procedure; and third, an approximate solution to the problem of incomplete overlap between the anatomical and functional scans. In addition, our method creates informative corrected functional images which also contain high-resolution anatomical details. The method has been applied to SPECT and MRI at our institution, but can be applied directly to other modalities such as PET and CT. Finally, our dedicated phan-

tom allowed us to demonstrate the high quantitative accuracy of the method. We believe that this method will aid with the differentiation between brain atrophy and focal functional deficits. Such differentiation may be particularly helpful for the study of HIV brain disease and degenerative brain disorders such as Alzheimer's disease, in which a combination of atrophy and functional deficits is present.

ACKNOWLEDGMENTS

REFERENCES

Blokland J, Reiber J, Pauwels E (1992): Quantitative analysis in single photon emission tomography. Eur J Nucl Med 19:47–61.

Chawluk J, Alavi A, Dann R, Hurtig H, Bais S, Kushner M, Zimmerman R, Reivich M (1987): Positron emission tomography in aging and dementia: Effect of cerebral atrophy. J Nucl Med 28:431–437.

Costa D, Morgan G, Lassen N (1993): New Trends in Nuclear Neurology and Psychiatry. London: John Libbey.

Hashimoto J, Kubo A, Ogawa K. Amano T, Fukuuchi Y, Motomura N, Ichihara T (1997): Scatter and attenuation correction in technetium-99m brain SPECT. J Nucl Med 38:157–162.

herscovitch P, Auchus A, Gado M, Chi D, Raichle M (1986): Correction of positron emission tomography data for cerebral atrophy. J Cereb Blood Flow Metab 6:120–124.

Hoffman E, Huang S, Phelps ME (1979): Quantitation in positron emission computed tomography: 1. Effect of object size. J Comput Assist Tomogr 3:299–308.

Itti L, Chang L, Ernst T (1996): Robust multimodality registration for neuroimaging. In: Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM), New York 1:35.

Itti L, Chang L. Mangin JF, Darcourt J, Ernst T (1997): Robust multimodality registration for brain mapping. Hum Brain Mapping 5:3–17.

Kanno I, Lassen NA (1979): Two methods for calculating regional cerebral blood flow from emission computed tomography of inert gas concentration. J Comput Assist Tomogr 3:71–76.

Kanno I, Uemura K, Miura Y, Miura S (1980): Headtome: A hybrid emission tomograph for brain. Med Radionucl Imag 1:153–154.

Kim H, Zeeberg BR, Fahey FH, Hoffman EJ, Reba RC (1992): 3-D SPECT simulations of a complex 3-D mathematical brain model: Effects of 3-D geometric detector response, attenuation, scatter, and statistical noise. IEEE Trans Med Imag 11:176–184.

Labbe C, Froment J, Kennedy A, Ashburner J, Cinotti L (1996): Positron emission tomography metabolic data corrected for cortical atrophy using magnetic resonance imaging. PhD thesis, Department of Neurology, CERMEP, Lyon.

Masdeu J, Brass L, Holman B, Kushner M (1994): Brain single-photon emission computed tomography. Neurology 44:1970–1977.

Mazziotta JC, Phelps ME, Plummer D, Kuhl DE (1981): Quantitation in positron emission computed tomography: 5. physical-anatomical effects. J Comput Assist Tomogr 5:734–743.

Meltzer CC, Leal J, Mayberg H, Wagner HJ, Frost J (1990a): Correction of PET data for partial volume effects in human cerebral cortex by MR imaging. J Comput Assist Tomogr 14:561–570.

Meltzer CC, Bryan R, Holcomb H, Kimball A, Mayberg H, Sadzot B, Leal J, Wagner H, Frost J (1990b): Anatomical localization for PET using MR imaging. J Comput Assist Tomogr 14:418–426.

Meltzer CC, Nichols TE, Mintun MA (1995): Effect of resolution on recovery of gray matter activity in FDG PET. J Nucl Med 36:114P.

Müller-Gártner HW, Links JM, Prince JL, Bryan RN, McVeigh E, Leal JP, Davatzikos C, Frost JJ (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 12:571–583.

Paley M, Chong W, Wilkinson I, Shepherd J, Clews A, Sweeney B, Hall-Craggs M, Kendall B, Newman S, harrison M (1994): Cerebrospinal fluid-intracranial volume ratio measurements in patients with HIV infection: CLASS image analysis technique. Radiology 190:879–886.

Tan P, Bailey D, Hutton B, al e (1989): A moving line source for simultaneous transmission/emission SPECT. J Nucl Med 30:964.

videen T, Perlmutter J, Mintun M, Raichle M (1988): Regional correction of positron emission tomography data for the effects of cerebral atrophy. J Cereb Blood Flow Metab 8:662–670.