Published in final edited form as:

Mol Imaging Biol. 2008 ; 10(1): 40–47. doi:10.1007/s11307-007-0110-7.

Compressed Voxels for High-Resolution Phantom Simulations

Richard Taschereau and Arion F. Chatziioannou

Department of Molecular and Medical Pharmacology, Crump Institute for Molecular Imaging, The David Geffen School of Medicine at UCLA, Los Angeles, CA, USA

Abstract

Purpose—We report here on a technique to implement high-resolution objects with voxels having variable dimensions (compressed) for the reduction of memory and central processing unit (CPU) requirements in Monte Carlo simulations. The technique, which was implemented in GATE, the GEANT4 application for positron emission tomography/single photon emission computed tomography (PET/SPECT) imaging simulations, was developed in response to our need for realistic high-resolution phantoms for dosimetry calculations.

Procedures—A compression algorithm similar to run-length encoding for one-dimensional data streams, was used to fuse together adjacent voxels with identical physical properties. The algorithm was verified by conducting dosimetric calculations and imaging experiments on compressed and uncompressed phantoms.

Results—Depending on the initial phantom size and composition, compression ratios of up to 99.9% were achieved allowing memory and CPU reductions of up to 85% and 70%, respectively. The output of the simulations was consistent with respect to the goals for each type of simulation performed (dosimetry and imaging).

Conclusions—The implementation of compressed voxels in GATE allows for significant memory and CPU reduction and is suitable for dosimetry as well as for imaging experiments.

Keywords

GATE; Simulations

Introduction

Monte Carlo simulation of radiation transport through matter is a powerful tool used for many diagnostic imaging applications such as positron emission tomography (PET) [1–6], single photon emission computed tomography (SPECT) [7–9], and X-ray computed tomography (CT) [10]. The same tool is also used in therapeutic applications such as dosimetry and radiation therapy planning.

The simulations include a description of all the relevant elements involved in the experiment, including the tomographic gantry, the detectors, the subject, the bed and other support, as well as any shields. This description can be achieved either with mathematical formulations that range from elementary geometrical shapes like spheres, cylinders, and boxes, to non-uniform rational B-splines [11,12], or with voxelized phantoms. Elementary shapes are appropriate to describe tomographs and simple phantoms, but they cannot adequately represent living subjects. As a result, geometrical phantoms can only provide estimates of imaging parameters like spatial resolution, partial volume recovery coefficients,

scatter, and radiation dose, but cannot truly reproduce realistic imaging situations. Mathematical formulations achieve much better approximation of living subjects, but the ultimate flexibility and ease of use is provided from voxelized phantoms that can represent arbitrary distributions.

Realistic description is essential in the evaluation of imaging systems and their performance at both the design stages as well as at the optimization of the imaging protocols [13,14]. It also is critical in dosimetry applications [15,16] as well as for more complicated imaging systems [17]. In these cases, the accurate geometry description of animals for small animal PET or of larger patients for clinical PET is important.

Anatomically detailed digital phantoms have been developed [18–20] that are described in the form of either a matrix of voxels (voxelized form) or in mathematical terms (e.g., non-uniform rational B-splines surfaces). For the purpose of simulations though, voxelized realizations of phantoms are used by most computer codes. In this case, depending on the overall size and desired spatial resolution, one can rapidly reach memory and computational speed limitations as phantoms become very large (10^8 voxels). However, in practice high spatial resolution is often not required throughout the entire phantom, and therefore high voxelation phantoms often lead to significant waste of computer resources. For example, in dosimetry applications, one is concerned with target volumes and sensitive organs, which typically constitute only a fraction of the entire phantom. In imaging applications, high spatial resolution is only needed to maintain smooth boundaries between phantom sections. Fortunately, voxelized phantoms have large groups of adjacent voxels belonging in regions with identical properties. This observation suggested that some form of compression of voxelized descriptions could be used to reduce memory requirements. To satisfy our need for high spatial resolution phantoms, we have developed such a compression scheme.

In this paper, we report on our method to implement variable-size compressed voxels for phantoms. This method reduces significantly the memory and central processing unit (CPU) requirements in GEANT4 application for tomographic emission (GATE) Monte Carlo simulations.

Methods

Algorithm

We define compressed voxels as rectangular voxels having variable size in every dimension. They are the result of the fusion of elementary equal-sized voxels from an initial phantom, sharing physical content properties such as material composition and density. Compressed voxels V have the following format: $\{\bar{X}, d\bar{X}, \bar{n}\}$, where for a 3D space, $\bar{X} = \{X_1, X_2, X_3\}$ is the position vector, $d\bar{X} = \{dX_1, dX_2, dX_3\}$ is the voxel size vector, and $\bar{n}$ is a vector representative of the voxel contents. This vector represents the material or other phantom property of interest. The compression method developed for three-dimensional (3-D) voxelized objects is similar to the run length encoding (RLE) compression algorithm [21] for one-dimensional streams of data. Essentially, the method fuses adjacent voxels of the same dimensions and material to yield a set of rectangular voxels of various sizes. More specifically, three successive searches—one along each dimension—are performed on the voxelized object. In each search, voxels are fused if they are adjacent and of identical material or other relevant property.

As an example, consider the $4 \times 8 \times 8$ synthetic phantom shown in Fig. 1, made with three “materials” identified by white, light gray, and dark gray. The first search fuses voxels along lines (the X_3 dimension), leaving voxels of dimensions $d\bar{X} = \{1, 1, dX_3\}$. Here, the entire first row of eight voxels is fused because all voxels are made out of “light gray” material.

This produces one compressed voxel with dimensions $\{1, 1, 8\}$. The second search fuses voxels in the X_2 dimension creating voxels with dimensions $d\vec{X} = \{1, dX_2; dX_3\}$. Similarly, for the third dimension the resulting voxels are $d\vec{X} = \{dX_1; dX_2; dX_3\}$. The same approach can be extended to phantoms containing more than three dimensions as, for example, temporal information for gated data acquisitions.

A single compressed voxel holds more information than an uncompressed voxel: position and size in addition to other voxel properties; therefore, it uses more memory. As with one-dimensional RLE, there is a minimum number of voxels (threshold) before an overall memory allocation gain can be achieved. This threshold is clearly implementation dependent. However, even in the case when this criterion is not met and no substantial memory reduction is achieved, there will still be a reduction in the number of voxels. This can lead to decreased CPU time as there will be fewer geometrical boundaries crossed and less checks to determine voxel properties during particle transport.

Implementation

The algorithm has been implemented in GATE [22], the GEANT4 application for tomographic emission, used for Monte Carlo simulations of PET and SPECT experiments. The implementation exploits the G4PVParameterized class of GEANT4 [23] to define parameterized volumes. In this case, a parameterized volume is a generic volume representing a collection of repeated volumes in which each copy can be different in size, shape, and material. The actual properties of an elementary volume are given by a parameterization function. This function must provide the geometrical transformation to be applied (rotation and translation), the dimensions, shape, and material. This GEANT4 class is especially well suited for implementing the variable-size voxel scheme. One only needs to provide a parameterization function that will consult a list of compressed voxels containing all the required information.

In the implementation described here, 2-byte unsigned integers are used to hold each element of the compressed voxel. The position vector $\vec{X} = \{X_1, X_2, X_3\}$ uses three 2-byte integers; the same holds for the dimension vector $d\vec{X} = \{dX_1, dX_2, dX_3\}$, and finally 2 bytes are allocated for the voxel contents value n . The result is a total of 14 bytes per voxel. The voxel value n is an index (a scalar) into a material pointer table. As there are 2^{16} different values that can be represented by a 2-byte unsigned integer, the maximum phantom dimensions capable of being represented by this scheme is 2^{16} along each axis, with no more than 2^{16} different materials.

As a convenience, an option has also been implemented to allow regions to be excluded from the compression mechanism. In those excluded regions, although voxels are not compressed, they are still converted to the compressed format. The exclusion mechanism still allows some memory (and CPU) reduction while retaining the benefits of high resolution when required. Regions to be excluded are identified by material n with an “exclude” statement in the GATE input script. More than one region can be excluded in the same simulation run.

Experiments

The algorithm was verified by conducting two types of simulated experiments. First, radiation dosimetry experiments were selected because they are most likely to be affected by changes in phantom representation as they record information (dose) at the voxel level. Second, we have selected PET imaging experiments. These simulation experiments should be less sensitive to the phantom voxel representation as long as the activity distribution is not affected.

Dosimetry

For the radiation dosimetry experiments, different phantom sizes and shapes were used to assess memory and CPU time reduction and also to assess any effects compression could have on dose calculations. Dose was calculated using an optional feature of GATE, which produces dose matrices having the same dimensions as the input phantom.

In the first simulation experiment, a voxelized version of the Hoffman brain phantom [19] was used. The phantom is a rectangular matrix of $128 \times 128 \times 55 \approx 10^6$ voxels defining five regions (materials): exterior of the phantom (air), the polymethyl methacrylate cylinder, and three internal regions representing gray, white matter, and ventricles. The phantom was used in three ways: uncompressed form, fully compressed form, and in compressed form with gray matter excluded from compression. One slice of the phantom (Fig. 2a) is shown without and with compression in Fig. 2b and c, respectively. After full compression, the number of voxels was reduced from 901,120 to about 40,000, giving a compression ratio of 95.6%, whereas memory use decreased by about 30%. Activity was disseminated in regions corresponding to gray and white matter with concentrations of 906 and 430 Bq/mL, respectively, and a total of 10^6 positrons were tracked for each experiment. Dose-volume histograms were subsequently calculated by combining dose from dose matrices and organ information from the phantom.

It is important to point out that the spatial distribution of activity in the phantom remains unaffected by the compression mechanism and is solely determined by the original, uncompressed phantom. It is therefore the exact same number of particles followed for all three brain phantom simulation experiments, regardless of the compression status of the phantom.

Additional experiments with larger phantoms were also performed. A realistic mathematical mouse phantom—the MOBY phantom [18]—was realized with two voxel sizes, 200 and 100 μm , yielding phantoms with 10^7 and 10^8 voxels, respectively. Dosimetry experiments were performed by subjecting these phantoms to a polyenergetic x-ray beam mimicking a micro-CT scanner [15,24]. For the purpose of this work, 10^9 photon histories were followed.

Imaging

The imaging experiment was inspired from current work at our institution aimed at determining the minimum detectable activity in typical mouse experiments. Whereas the actual simulations, image reconstruction algorithms, and data analysis for this goal is beyond the scope of the work presented here, the voxel compression has made this research possible. Here, therefore, we only present the feasibility of simulating a realistic activity distribution in a computational platform. To achieve this, a modified version of the MOBY phantom was used in which a spherical tumor (7 mm diameter) was added in the axillary area. The phantom was realized with a voxel size of $(400 \mu\text{m})^3$. Fluorine-18 activity was distributed in the phantom as follows: uniform background of 5 nCi/ mm^3 was put in major organs (guts, muscle mass, liver, kidneys, spleen, brain, and testes), 10 nCi/ mm^3 of activity was put in the tumor and the myocardium to obtain a 2:1 activity ratio, and 25 nCi/ mm^3 of activity was put in the bladder, for a total of 110 μCi . A 10-minute acquisition scan was divided into 600 one-second simulations to be run on a 27-dual-cpu computer cluster. The simulated tomograph was based on the microPET Focus 220 (Siemens, Preclinical Solutions, Knoxville, TN). The image was reconstructed from an attenuation-corrected sinogram using two iterations of OSEM-3D (six subsets) followed by 18 iterations of maximum *a posteriori* (MAP) algorithm with a regularization (beta) parameter value of 0.1. Here again, the experiment was performed once with the uncompressed phantom and once with the compressed version.

Results and Discussion

Effect of Compression on Dose-Volume Histograms

Dose distributions for different organs calculated with compressed phantoms had a narrower width when compared to those calculated with uncompressed phantoms. This effect can be best explained by the definition of dose, which is the ratio of energy deposited over mass. Let $E_{i,k}$ be the i^{th} energy deposit from particle interactions in elementary (uncompressed) voxel k , N_k be the number of such energy deposits, and m_k be the voxel mass. The dose D_k deposited in elementary voxel k is given by:

$$D_k = \frac{\sum_{i=1}^{N_k} E_{i,k}}{m_k}. \quad (1)$$

If a number P of those voxels are fused together, then dose D in the fused larger voxel is given by:

$$D = \frac{\sum_{k=1}^P \sum_{i=1}^{N_k} E_{i,k}}{\sum_{k=1}^P m_k}. \quad (2)$$

As media in fused voxels is homogeneous and all elementary voxels have the same size, they also have the same mass. Hence, the sum in the denominator can be replaced by a product:

$$D = \frac{\sum_{k=1}^P \sum_{i=1}^{N_k} E_{i,k}}{P \cdot m_k} = \frac{1}{P} \sum_{k=1}^P \frac{\sum_{i=1}^{N_k} E_{i,k}}{m_k} = \frac{1}{P} \sum_{k=1}^P D_k, \quad (3)$$

which is the average of dose values (1) in elementary voxels. Therefore, larger voxels have an averaging effect, i.e., they dampen out fluctuations. This will be reflected in dose-volume histograms: compressed voxels with large volumes will score dose values closer to the average (averaging effect). The dose distribution will then be closer to the average, and the distribution is narrower. This narrowing effect will be more pronounced with higher compression rates that produce larger parametric voxels. In the limit case of 100% compression rate: only one voxel would be left and the dose distribution would be a single value at the (average) dose of that voxel, for any single Monte Carlo simulation.

Brain Phantom

The dose-volume histograms for the PET simulation with the Hoffman brain phantom are shown in Fig. 3. Dose is normalized to the total activity present in the phantom ($\mu\text{Gy/GBq}$) and curves show results from the uncompressed (solid thin line), partially compressed—gray matter excluded (solid thick gray line) and fully compressed (dotted line) versions of the phantom.

The dose averages for white matter (Fig. 3a) were 35.2 $\mu\text{Gy/GBq}$ (standard deviation ± 5.4 $\mu\text{Gy/GBq}$) and 35.4 $\mu\text{Gy/GBq}$ (± 13.0 $\mu\text{Gy/GBq}$) for the fully compressed and uncompressed versions, respectively. Whereas the average dose is the same, the distribution is narrower for the compressed version as can be seen in Fig. 3 and from the standard deviation (SD) values. For the partially compressed phantom, dose was 35.3 (± 5.4) $\mu\text{Gy/GBq}$ and the distribution was identical to the fully compressed phantom, as white matter was compressed in both versions.

For gray matter (Fig. 3b), the dose was 53.6 (± 6.0) $\mu\text{Gy/GBq}$, 53.7 (± 16.3) $\mu\text{Gy/GBq}$, and 53.8 (± 16.3) $\mu\text{Gy/GBq}$ for the fully compressed, uncompressed, and partially compressed phantoms, respectively. Here, the distributions for the uncompressed and partially compressed phantoms are identical as gray matter was excluded from compression. As with white matter, the distribution from the fully compressed phantom is narrower than the others.

Mouse Phantom

In Fig. 4 are shown dose-volume histograms for the liver from an x-ray CT simulation with the MOBY phantom (200 μm voxel size) using uncompressed and compressed versions. Dose averages were 6.3 (± 2.2) μGy and 6.3 (± 1.2) μGy for the uncompressed and compressed phantoms, respectively. Again, with a high compression rate ($>99\%$), a narrower dose distribution was obtained without changing the average.

Finally, in Table 1 are shown performance results obtained as a function of phantom size. As the phantom size grows larger, the corresponding memory reduction becomes more significant. Although this technique was primarily implemented to reduce memory requirements, important reduction in CPU time has also been achieved. In the case of the high-resolution mouse phantom (MOBY 100 μm), a comparison with the uncompressed version was not possible because the phantom could not be accommodated in the ~ 3 GB of memory available in the user address space.

PET Imaging

A comparison of the activity distribution in the phantom and the reconstructed PET image (both compressed and uncompressed cases) are shown in Fig. 5. From a qualitative point of view, the PET images look similar; there is no apparent artifact from the voxel compression process and they both render regions with high, medium, and no activity. The images do not look exactly the same because they are two different realizations of a random process.

The Student t test has been used to evaluate the difference in means between regions of interest (ROI) in the two images. ROIs were drawn around the tumor, bladder, the heart, and the mid section of the body. The average value and variance of each region was calculated for each image and a t statistic was calculated with the following formula:

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2 + s_2^2}}, \quad (4)$$

where $\bar{x}_k$, and s_k^2 are the average pixel value, and variance, in a ROI of image k , respectively. P values, representing the probability that the t statistic would be less than the observed value assuming that the two means are the same, were calculated using:

$$A(t|\nu) = \frac{1}{\sqrt{\nu} B\left(\frac{1}{2}, \frac{\nu}{2}\right)} \int_{-t}^t \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}} dx, \quad (5)$$

where ν is the number of degrees of freedom and B is the Beta function. All calculated p values were below 10^{-4} for all regions.

In Fig. 6, which shows profiles obtained across the tumor from the compressed and uncompressed phantom, the profiles are similar for both phantoms (allowing for statistical fluctuations) and the tumor-to-background ratio of recovered activity is close to the nominal ratio of 2:1. The total cpu time spent for simulation with the compressed and uncompressed phantoms was 2,400 and 4,200 hours, respectively.

Remarks

In theory, because the compression algorithm does not alter the information content of the phantom, the use of compressed voxels should not have any effect on the outcome of the simulation. In practice, however, the performance ultimately depends on the Monte Carlo code and the physics models used. For codes using a condensed history technique for charged particles transport (such as GEANT4), artifacts can be introduced by altering the particle transportation step size [25]. As geometrical boundaries are step-limiting factors, small uncompressed voxels will force shorter steps as opposed to larger ones with compressed voxels, potentially leading to different results. As a general rule, shorter steps (and small, uncompressed voxels) yield more accurate results. Users would be well advised to verify that compression is appropriate for their application (geometry and energies involved) by performing simple experiments, e.g., calculating dose-volume histograms. Any observed differences in results could be overcome as most codes offer parameters to limit the electron step size.

Applications other than PET imaging and dosimetry should be able to take advantage of compressed voxels without incurring any significant changes. This includes other imaging applications (x-ray, CT, SPECT), systems performance analysis, and scatter analysis among others. The voxelized object does not have to be a synthetic phantom like the MOBY mouse; more realistic subjects could be used. For example, compressed voxelized objects could be generated from segmented CT images of a small animal and used in any type of imaging experiment. It is even possible to consider compressing images obtained from clinical scans. The size of a clinical scan is of the order of 10^8 voxels, similar to the high-resolution MOBY phantom used in this study. Using the compression method described here, the image would fit into memory and CPU reduction would be achieved. However, enough computing power would still be required to perform the simulations in a reasonable amount of time.

The compression exclusion mechanism should prove very useful in dosimetry applications when one needs to investigate dose distribution in only one or a few organs. By excluding these organs from compression, high-resolution dose-volume histograms can be calculated with improved performance without loss in precision. In our opinion, compression should be the rule rather than the exception in these types of simulations.

Conclusion

A voxel compression algorithm has been developed and was implemented in GATE. Dosimetry calculations and PET imaging experiments were performed with various

phantoms with and without compression. Depending on phantom size and shape, memory reduction of up to 85% and CPU reduction of up to 70% have been observed. Dose distributions in compressed phantoms were narrower (smaller standard deviation) than in uncompressed phantoms; however, dose averages were similar. No artifact was noticed in PET images because of the voxel-compression process. In addition to memory and CPU time reduction, the technique makes it possible to use very large phantoms that would not otherwise fit into memory.

Acknowledgments

The authors wish to thank the members of the Open-GATE collaboration for letting them be part of it and for accepting this work as a contribution to the development of GATE. This work was supported in part by the U.S. Department of Energy under Contract DE-FC03 02ER63420 and by the National Institutes of Health under Grant R24 CA92865.

References

1. Vandenberghe S, Daube-Witherspoon ME, Lewitt RM, Karp JS. Fast reconstruction of 3D time-of-flight PET data by axial rebinning and transverse mashing. *Phys Med Biol* 2006;51:1603–1621. [PubMed: 16510966]
2. Qi JY, Leahy RM. Resolution and noise properties of MAP reconstruction for fully 3-D PET. *IEEE Trans Med Imag* 2000;19:493–506.
3. Huang SC, et al. An Internet-based “kinetic imaging system” (KIS) for MicroPET. *Mol Imaging Biol* 2005;7:330–341. [PubMed: 16132473]
4. Kaplan MS, Harrison RL, Vannoy SD. Coherent scatter implementation for SimSET. *IEEE Trans Nucl Sci* 1998;45:3064–3068.
5. Lewitt RM, Muehllehner G, Karp JS. 3-dimensional image-reconstruction for PET by multislice rebinning and axial image filtering. *Phys Med Biol* 1994;39:321–339. [PubMed: 15551583]
6. Thompson CJ, Morenecantu J, Picard Y. PETSIM: Monte-Carlo simulation of all sensitivity and resolution parameters of cylindrical positron imaging-systems. *Phys Med Biol* 1992;37:731–749. [PubMed: 1565700]
7. Beekman FJ, de Jong HWAM, van Geloven S. Efficient fully 3-D iterative SPECT reconstruction with Monte Carlo-based scatter compensation. *IEEE Trans Med Imag* 2002;21:867–877.
8. Yanch JC, Dobrzeniecki AB. Monte-Carlo simulation in SPECT - Complete 3D-modeling of source, collimator and tomographics data acquisition. *IEEE Trans Nucl Sci* 1993;40:198–203.
9. Yanch JC, Dobrzeniecki AB, Ramanathan C, Behrman R. Physically realistic Monte-Carlo simulation of source, collimator and tomographic data acquisition for emission computed-tomography. *Phys Med Biol* 1992;37:853–870.
10. Ay MR, Zaidi H. Development and validation of MCNP4C-based Monte Carlo simulator for fan- and cone-beam x-ray CT. *Phys Med Biol* 2005;50:4863–4885. [PubMed: 16204878]
11. Segars WP, et al. Development and application of the new dynamic Nurbs-based Cardiac-Torso (NCAT) phantom. *J Nucl Med* 2001;42:7P–7P.
12. Ward RC, et al. Creating a human phantom for the virtual human program. *Stud Health Technol Inform* 2000;70:368–374. [PubMed: 10977575]
13. Barret O, Carpenter TA, Clark JC, Ansorge RE, Fryer TD. Monte Carlo simulation and scatter correction of the GE advance PET scanner with SimSET and Geant4. *Phys Med Biol* 2005;50:4823–4840. [PubMed: 16204875]
14. Baete K, et al. Evaluation of anatomy based reconstruction for partial volume correction in brain FDG-PET. *Neuroimage* 2004;23:305–317. [PubMed: 15325378]
15. Taschereau R, Chow PL, Chatzioannou AF. Monte Carlo simulations of dose from microCT imaging procedures in a realistic mouse phantom. *Med Phys* 2006;33:216–224. [PubMed: 16485428]
16. Taschereau R, Chatzioannou A. Monte Carlo simulations of absorbed dose in a mouse phantom from 18-fluorine compounds. *Med Phys* 2007;34:1026–1036. [PubMed: 17441249]

17. Rannou FR, Kohli V, Prout DL, Chatziioannou AF. Investigation of OPET performance using GATE, a Geant4-based simulation software. *IEEE Trans Nucl Sci* 2004;51:2713–2717. [PubMed: 16429604]
18. Segars WP, Tsui BMW, Frey EC, Johnson GA, Berr SS. Development of a 4-D digital mouse phantom for molecular imaging research. *Mol Imaging Biol* 2004;6:149–159. [PubMed: 15193249]
19. Hoffman EJ, Cutler PD, Digby WM, Mazziotta JC. 3-D phantom to simulate cerebral blood-flow and metabolic images for PET. *IEEE Trans Nucl Sci* 1990;37:616–620.
20. Zubal IG, et al. Computerized 3-dimensional segmented human anatomy. *Med Phys* 1994;21:299–302. [PubMed: 8177164]
21. Salomon, D. Data compression: the complete reference. New York: Springer; 2000.
22. Jan S, et al. GATE: a simulation toolkit for PET and SPECT. *Phys Med Biol* 2004;49:4543–4561. [PubMed: 15552416]
23. Agostinelli S, et al. Geant4—a simulation toolkit. *Nucl Instrum Methods Phys Res Sect A Accel Spectrom Detect Assoc Equip* 2003;506:250–303.
24. Taschereau R, Chow PL, Cho JS, Chatziioannou A. A microCT x-ray head model for spectra generation with Monte Carlo simulations. *Nucl Instrum Methods Phys Res Sect A Accel Spectrom Detect Assoc Equip* 2006;569:373–377.
25. Bielajew, A.; Rogers, DWO. Monte Carlo transport of electrons and photons. Jenkins, TM.; Nelson, WR.; Rindi, A., editors. Plenum Press; New York: 1988. p. 115-137.

Fig. 1. The three steps of compression illustrated on this synthetic phantom: **a** original; **b** compression along X_3 , **c** compression along X_2 and **d** compression along X_1 .

Fig. 2.

a A coronal slice of the voxelized Hoffman brain phantom. Close up of the outlined section in uncompressed **b** and compressed **c** formats.

Fig. 3.

Dose-volume histograms for the Hoffman brain phantom: **a** white matter, **b** gray matter. *Solid thin lines* represent uncompressed phantoms, *solid thick gray lines* represent partially compressed phantoms, and *dotted line* fully compressed phantoms. The average dose is the same in all cases, only the distribution changes (narrower when the region is compressed).

Fig. 4. Dose-volume histograms for the liver in the MOBY phantom at 200 μm resolution. The *thick gray line* is for the uncompressed phantom and the *thin black line* is for the compressed version. The average dose is almost identical at 6.3 μGy and the distribution is narrower with the compressed phantom (± 1.2 μGy) compared to the uncompressed one (± 2.2 μGy).

Fig. 5. Activity distribution in the voxelized mouse phantom in **a** a transverse plane and **d** a coronal plane; and from the MAP-reconstructed PET image in a transverse plane, **b** uncompressed phantom, **c** compressed phantom, and coronal plane, **e** uncompressed phantom, and **f** compressed phantom.

Fig. 6. Comparison of profiles across the tumor obtained from coronal images with compressed (*solid line*) and uncompressed (*dotted line*) phantoms. Curves are normalized to the maximum (in the tumor). The *horizontal dashed line* at 0.5 represents the level where the background should be to recover the nominal 2:1 tumor-to-background ratio.

Table 1

Performance achieved with various phantoms

Phantom name	No. of voxels	Compression (%)	Memory diff.	CPU diff.
Hoffman fully compressed	10 ⁶	97	−30%	−45%
Hoffman partially compressed	10 ⁶	88	−28%	−45%
MOBY 200 μm	10 ⁷	98	−85%	−70%
MOBY 100 μm	10 ⁸	99	n/a	n/a