Published in final edited form as:

AIP Conf Proc. 2010 January 5; 1204: 126–132.

Monte Carlo Simulation Of Emission Tomography And Other Medical Imaging Techniques

Robert L. Harrison

Department of Radiology, University of Washington Medical Center, Box 357987, 1959 Pacific NET – RR215, Seattle, Washington 98195, USA

Abstract

An introduction to Monte Carlo simulation of emission tomography. This paper reviews the history and principles of Monte Carlo simulation, then applies these principles to emission tomography using the public domain simulation package SimSET (a Simulation System for Emission Tomography) as an example. Finally, the paper discusses how the methods are modified for X-ray computed tomography and radiotherapy simulations.

Keywords

Simulation; emission tomography; x-ray computed tomography; radiotherapy

Introduction

Emission tomography is a branch of medical imaging in which radioactively-labeled compounds are injected (or otherwise introduced) into a patient. Detectors are deployed around the patient to detect the photons created by the decay of the radioactive labels. A three-dimensional estimate of the distribution of the radiolabel is created from the detected events.

Monte Carlo simulation of emission tomography has been used extensively in the development and testing of data correction and image reconstruction algorithms (e.g. [1]), for prototyping tomographs (e.g. [2]), and to optimize patient studies (e.g. [3]). In general these simulations track individual photons through the patient and tomograph, simulating the physical processes involved in photon transport through the patient and tomograph, then use more approximate models of the detection process.

In this paper, we discuss the models and approximations typically used in emission tomography simulation, using the public domain simulation package SimSET (a Simulation System for Emission Tomography, [4]) as an example.

In addition, we give an overview of some of the differences between emission tomography simulations and simulations of X-ray computed tomography and radiotherapy.

What is Emission Tomography?

Emission tomography differs from X-ray computed tomography (CT) in that the radiation source is inside the patient rather than transmitted through the patient (Fig 1). X-ray CT measures the density of the tissues inside the patient, resulting in high resolution images of patient anatomy; emission tomography measures the concentration of a radioactively labeled compound in the patient, and gives information about the metabolic state of the patient's tissues. The information is dependent on the compound that is labeled: for instance, Sestamibi labeled with Technetium-99m is used for myocardial perfusion imaging in single photon emission

computed tomography (SPECT); in positron emission tomography (PET), fluorine-18 labeled fluorodeoxyglucose (FDG), an analogue of glucose, is often used to image glucose uptake. Emission tomography images are much lower resolution than X-ray CT images, but the metabolic information in them is often a useful complement. For instance, cancers often have the same density as the tissue surrounding them and are hard to see in X-ray CT images, but cancers tend to use much more glucose than surrounding tissues and show up as hot spots in FDG PET images (Fig. 2).

Simulating Emission Tomography

It is impossible and unnecessary to exactly simulate an emission tomography study; for instance, we can not know the exact movement of a patient's heart during a study, but its exact position is not going to affect images of the prostate. When designing a simulation we need to determine what factors we want/need to simulate and how exactly we want to simulate them. In the case of the heart, we may decide to ignore motion some simulations and simulate a moving heart in others.

Some of the decisions are driven by the appraisal that a factor is too small to worry about: few emission tomography simulations worry about the backscatter of photons off the walls of the imaging suite. Other decisions are made based on convenience: modeling the motion of a human heart is very difficult to do well, so it has often been ignored or modeled quite simplistically (though there is now a digital phantom that incorporates a realistic model [8]).

There are many emission tomography simulation programs available, each incorporating different choices about exactness, ease-of-use, and efficiency. In general, as simulations provide greater flexibility and accuracy, they become slower and more difficult to use (Fig. 3). For example, some of the general-purpose simulation software developed by national labs for high-energy physics can also be used for emission tomography. They incorporate the most up-to-date, precise modeling of all the physics involved in particle transport, track each particle separately, and allow anything to be simulated, but they are very slow and can be extremely difficult to set up for emission tomography simulations (this latter issue has been addressed, to some extent, with the GATE package [9], which is a shell that makes the Géant general-purpose simulation easier to use for emission tomography). In contrast, many researchers have used much simpler, quicker analytic simulations (e.g. ASIM [10]), that model decay and photon attenuation (nearly) exactly as integral equations, but either ignore or approximate other factors (e.g. photon scatter).

The remainder of this section discusses the modeling decisions necessary when designing an emission tomography simulation, focusing on the choices we made designing SimSET, particularly the PET options. Our goals were to create software that could rival general-purpose packages in accuracy for most emission tomography simulations, yet would be considerably easier-to-use and faster.

What Needs Simulating?

In emission tomography the scanner surrounds (most PET scanners, e.g. Fig. 2) or rotates around (most SPECT scanners) some portion of a patient. In general, an elliptical cylinder can be placed inside the tomograph and outside the patient. SimSET assumes that all radioactivity is inside the cylinder and the entire detection process occurs outside the cylinder. (Knowledge of the basic geometry of the simulation enables considerable efficiency improvements.)

The Patient/Phantom/Object—For the most part, the general-purpose simulations model everything in the simulation as a collection of unions and intersections of geometrical primitives like cylinders, rectangular boxes, and ellipsoids. However, it is difficult to model

humans with such primitives, and tracking particles through large collections of geometric primitives can be quite slow: tracking a particle may require computing intersections with a large fraction of the primitives. For SimSET we chose to represent the patient, or object as we call it, as a voxelated volume. Emission tomography researchers often use this format for their data, and there are fast methods for tracking through a voxelated volume.

To define the object volume requires two different views: the activity object and the attenuation object. The activity object specifies the concentration of radioactive isotope in each voxel and the isotope used. Some radioactive isotopes emit a varying cascade of several types of particles as they decay. SimSET currently simulates isotopes that emit a positron or a single photon; isotopes with complex decay schemes are not supported, though to support them would not significantly slow program execution. For the attenuation object, SimSET has a selection of supported materials that includes materials used in scanners or phantoms (e.g., lead, scintillation crystals, water, polyurethane) and tissue types (e.g., brain, blood, bone). The selection is far from complete, but new materials can be added as necessary.

The Scanner—SimSET splits the scanner simulation into two modules, the collimator and detector modules. SimSET provides a few general models for each: only tomographs fitting one of these models can be simulated. This is considerably less flexibility than general-purpose simulations. However, many emission tomographs fit into SimSET's models reasonably well, and the prior knowledge of the tomograph geometry allows the code to be considerably faster.

Many modern PET scanners have no collimators within the field-of-view, but most of those that do use cylindrical annuli of lead or tungsten called septa to help define the image slices axially (Fig. 4). SimSET models these cylindrical annuli, allowing for material changes both radially and axially. In tomographs, these septa need some support structure to hold them in place: SimSET's model does not include this.

SimSET has two PET detector models: detector arrays represented as a collection of cylindrical annuli (cylindrical detectors) or as a collection of rectangular boxes (block detectors), the latter allowing more exact representation of typical PET scanners (Fig. 5).

In addition to the above elements, tomographs use complicated electronics to read and process the detected events. SimSET does not try and model this process exactly, using a simple approximation for the detected photon position. For each photon SimSET does provide a call to a user-modifiable function where more complicated modeling may be added.

The Output—Typically PET scanners output a histogram of how many coincidences were detected by each pair of detectors. Some tomographs also track the detection time difference for the two photons, and split the line between the detectors into a series of time-of-flight bins. A simulation should be able to replicate these capabilities, but frequently we also want a simulation to output quantities not available from the tomograph, for instance the energy of the photons in a coincidence, whether one or both of the photons has scattered, whether the coincidence is a random coincidence. To support this, SimSET allows the user to select from a range of spatial, energy, time-of-flight, and coincidence-type histogramming options. SimSET also supports list-mode files, files which are a list of the above information for each coincidence.

The Simulation Processes

For the most part, even the most detailed simulation can be broken down to a decay-by-decay process, with all the particles and secondary particles (e.g., a photon from the capture of an electron excited by the original particle) from tracked independently. The exception to this is in the processing that occurs as tracked particles are converted to electrical signals, usually by

photomultiplier tubes connected to scintillation crystals. At this point particles from different decays can interact with each other to create random coincidences, detector deadtime and pileup.

Modeling Decays—The first step when generating a decay is to determine where and when it will occur. The number of decays that will occur in a given time interval from a given amount of radioactivity is a Poisson random variable. One can use the relationship between the Poisson and exponential distributions to generate the time to the next decay, and keep generating decays until the elapsed time is greater than the scan time. The voxel for each decay is chosen randomly, with a voxel's probability proportional to the activity in it. This may be the best method if factors like deadtime and pileup are being studied, where signals from several events can interact with one another. However, the study of these factors is not a major focus of SimSET. We chose a slightly faster method: we step through the voxels, generating all the decays for a voxel before moving to the next, assigning each decay a random time within the scan interval. (The correct decay order can be recovered by saving a list-mode file and sorting it by decay time.) Each decay is simulated at a random position within its voxel.

Next the type of decay is chosen. For some isotopes there are multiple possibilities, for instance some decays just produce a photon, others produce both a photon and a positron. For simplicity SimSET only simulates decays that produce a single photon or decays that produce a single positron.

When a radioactive decay produces a positron, the positron leaves the atom's nucleus with a lot of energy and follows a very torturous path until it annihilates with an electron, producing two 511 keV photons that are almost exactly anti-parallel. Simulating the positron path is very expensive computationally [12]. For SimSET we have chosen instead to sample from a relatively accurate approximation of the distribution of the distances positrons travel [13] and project a positron in a random direction from its decay [14]. Two 511 keV photons are created at this point and a random direction chosen for one of them. The direction for the second is chosen nearly anti-parallel to the first, with the difference sampled from a Gaussian with mean 0°, standard deviation 0.5° [15].

Photon Tracking—After annihilation the two photons are separately tracked. The photon tracking process remains the same throughout the object, collimator, and detectors.

First a dimensionless mean-free-paths-to-travel, P , is sampled from the exponential distribution with mean 1. The photon is then tracked from voxel to voxel (or through collimator/detector elements) until

$$\sum \mu_i d_i = P \quad (1)$$

where the summation is over the voxels the photon travels through, μ_i is the attenuation coefficient for the i^{th} voxel at the photon's energy, and d_i is the distance traveled through the i^{th} voxel. An interaction is simulated at this position.

SimSET models three interaction types: photoelectric absorption, Compton scatter, and coherent scatter. Interaction type is randomly sampled, with the probability for each type determined by the material at the interaction location and the photon's energy. If the photon is absorbed we discard it (except in detectors, as discussed below). For Compton scatter, we sample a scatter angle and compute a new photon energy using the Klein-Nishina formulas [16].

For coherent scatter we have computed tables using the Lawrence Livermore National Laboratories Evaluated Electron Density Library (LLNL EPDL) database [17] that approximate the coherent scatter distribution as a function of material and photon energy. The details of our algorithm can be found in [18]. (We have also developed, but not yet distributed, a similar algorithm for Compton scatter. It will be a bit faster and more accurate for low energy photons in high density materials: the Klein-Nishina formulas use a free-electron approximation, while the LLNL EPDL includes the bound electron correction. This correction is larger at lower energies/in denser materials.)

For Compton and coherent scatter, the photon is deflected by the scatter angle (the azimuth of the deflection is sampled from 0° to 360°) and the new energy computed if the scatter is Compton. A new mean-free-paths is sampled and tracking continues.

Photon tracking ends when the photon either leaves the user-defined object/collimator/detector space, is absorbed, or falls below a user-set energy threshold.

When energy is deposited during photon tracking (energy lost in Compton scattering or the photon energy in photoelectric absorption), an electron is left in an unbound or excited state. As it returns to a lower energy state photons are emitted. SimSET does not simulate these secondary photons – there are far more of them than primary photons, so the additional computational cost would be large. In general these photons have little impact in emission tomography. One exception is K-shell fluorescence from the collimator (“lead x-rays”) which can have an impact on scans using thallium-201 (lead x-rays are about the same energy as thallium's main emission) and on deadtime. Another exception is scintillation photons.

Detection and Signal Processing—Photon tracking continues using the same algorithm through the detector elements. However, in detectors we keep track of where energy is deposited. Most scanners use scintillating crystals for detectors. In these, energy deposited by photoelectric absorption or Compton scatter is emitted isotropically as visible light photons called scintillation photons. These photons are much lower energy than incoming gamma rays: hundreds are emitted per absorbed gamma ray.

Scintillation photons are then converted to current by light detectors, typically arrays of photomultiplier tubes (PMTs), or a single position-sensitive PMT. The detection electronics process these signals and to estimate the detected position of the photon. There are many algorithms for doing this, but most scanners still use algorithms similar to that proposed by Anger [19].

SimSET approximates this process using Anger logic, estimating the detected position of the photon as the energy-weighted centroid, X , of the interactions

$$X = \sum_{i=1}^N \frac{X_i E_i}{N} \quad (2)$$

where N is the number of gamma ray interactions in the crystal(s), X_i is the position of the i^{th} interaction, and E_i is the energy deposited at the i^{th} interaction. Scanners also estimate the total energy the gamma ray deposited in the scintillation crystal and, sometimes, the time-of-flight difference between the arrival of the coincidence photons (for PET). SimSET can output these values as well, Gaussian-blurred by the scanner's energy or time-of-flight resolution if desired.

Importance Sampling—SimSET has options to use importance sampling (IS) or variance reduction, statistical methods that speed convergence of the output estimate to its mean [20]. To simulate a typical patient scan requires simulating many billions of decays, and can require days or even months of computer time. Using IS techniques can reduce the number of decays/computer time required for a given precision by a factor of 3-100, depending on the problem being simulated.

Simulating X-Ray Tomography and Radiotherapy

In one sense the problems of simulating X-ray CT and radiotherapy are same as that of simulating emission tomography. The radiation source is outside the body, but the physics of photon transport is the same. One could simulate both using photon-tracking.

In practice this is never done. Simulating a typical X-ray CT scan would require tracking 1000s of times more photons than an emission tomography simulation, a radiotherapy treatment far more than even that. As mentioned above, a simulation of an emission tomography scan can days or months, so it would take far too long to do a photon-tracking scan for X-ray CT or radiotherapy. Instead analytic simulation techniques are used.

X-ray CT

In X-ray CT we measure how much of an incident beam of X-rays escapes the patient (Fig. 6). This transmitted fraction is a function of the densities (or, more specifically, electron densities) of the tissues between the source and detector,

$$T=1-\exp\left(-\int_{\text{line}}\mu(x,y)ds\right) \quad (3)$$

where the integral is over the line between the source and detector and μ gives the attenuation coefficient of the patient as a function of position.

A few corrections need to be made to this estimate. Firstly, the incoming beam is not monoenergetic, but a range of energies. As lower energies are more easily attenuated than higher energies, the average energy of the beam increases as it passes through the patient (Fig. 7). In a simulation this can be compensated for by repeating the simulation with a range of beam energies, then combining the simulations weighted according to the distribution of incident beam energies.

Simulations based on equation 3 ignore scatter. Fortunately, collimation at the beam source combined with the collimation implicit in the limited axial extent of the detector greatly reduces scatter. Scatter is a problem in some systems with large area detectors.

The process of simulation is to repeat the calculation in equation 3 for each line in the X-ray CT data set. It should be noted that this is not a Monte Carlo simulation: there is no random component.

Radiotherapy

Radiotherapy is used to treat cancer and other diseases. The patient is irradiated with huge doses – often 20 or more times what would be a lethal instantaneous whole-body dose. These huge doses can be tolerated because they are localized, because they are split over several incoming beams so that only the cancer (and not normal tissue) gets the maximum dose (Fig. 8), and because the dose is split up into many smaller doses over a period of weeks or months

(fractionated), which gives normal tissue a chance to partially heal between fractions. Even with these caveats, patients are often irradiated near to lethal limits, particularly for harder-to-cure cancers.

In some ways radiotherapy simulation is exactly the same as X-ray CT simulation – except it is completely different.

First of all, the output we are interested in is entirely different. In X-ray CT we are interested in the fraction of the beam that is transmitted; in radiotherapy we are interested in how much energy the beam deposits in the patient, organ by organ (Fig. 9), so the energy deposited by interactions as the beam crosses each voxel must be computed. Secondly, in X-ray CT the incident beam is the same strength on every line; in radiotherapy the tumor is irradiated from a limited number of angles. Thirdly, the beam energy is much higher, around 6 MeV as opposed to 70 keV. Finally, the radiotherapy simulator is used as a hands-on tool: a physicist will arrange the beam shape, direction, and intensity to increase dose to the tumor while keeping dose to normal tissues within safe limits (and each organ has its own limits). The simulation needs to be fast – there's a patient waiting to start therapy.

Until recently, software simulations were too slow and required too much user interaction/manipulation to be useful. Instead many treatment centers rely on “physical simulation”: a special X-ray CT scanner/simulator with the same positioning and beam-shaping capabilities as a radiation therapy treatment device, but with normal CT energies (which are better for imaging) and with a suite of programs designed for therapy planning. The patient is positioned in exactly the same position as for the treatment and gets a CT scan. A doctor and physicist identify the tumor volume(s) and key organs near the treatment volume, set the number of beams to use and the intensity and shape of each beam. The simulator is then set up to irradiate the patient at the selected angles, with the selected beam intensities and shapes, and the transmitted beams (still at CT energies) imaged. From these images, the simulator software computes the dose to the tumor and organs the beams intersect [24].

At some centers, virtual simulation is beginning to replace physical simulation. Much of the process is the same: in particular, a CT is still needed with the patient in treatment position (often requiring a CT scanner with a large patient port), and a doctor and physicist must still identify tumor and organs. However, no physical simulation of the beam shape/positioning is carried out. This is done, instead, entirely as a software simulation.

Further Reading

The book containing reference [4] is a good overview of Monte Carlo simulations of emission tomography, including the basic principles and physics of scintillation cameras, photon transport, importance sampling techniques, and some sample applications. For information about SimSET, see our website, http://depts.washington.edu/simset/html/simset_main.html. Aird and Conway [24] give a good overview of radiotherapy treatment planning simulation.

Acknowledgments

This work was supported in part by PHS grants CA42593 and CA126593.

References

1. Ollinger J. *Phys Med Biol* 2006;41:153–176. [PubMed: 8685253]
2. Shokouhi, S. Ph D Thesis. Stony Brook University; 2005. Image Reconstruction and System performance Simulations of Rat Conscious Animal PET (RatCAP).
3. Chicco P, et al. *Eur J Nucl Med* 2001;28:150–154. [PubMed: 11303884]

4. Lewellen, et al. The SimSET Program. In: Ljungberg, M., et al., editors. Monte Carlo Calculations in Nuclear Medicine. Bristol: Inst of Phys Publ.; 1998. p. 77-92.
5. Greenville, SC: Imaginis; Sep 13. 2007 Imaginis – How Does CT Work.
http://www.imaginis.com/ct-scan/how_ct.asp
6. San Francisco: Wikimedia Foundation, Inc.; Aug 21. 2009 Positron Emission Tomography – Wikipedia. http://en.wikipedia.org/wiki/Positron_emission_tomography
7. San Francisco: Wikimedia Foundation, Inc.; Aug 22. 2009 PET/CT – Wikipedia.
<http://en.wikipedia.org/wiki/PET-CT>
8. Segars, WP. Ph D Thesis. University of North Carolina; 2001. Development and Application of the New Dynamic NURBS-based Cardiac-Torso (NCAT) Phantom.
9. Jan S, et al. GATE: a simulation toolkit for PET and SPECT. *Phys Med Biol* 2004;49:4543–4561. [PubMed: 15552416]
10. Comtat, C., et al. Proceedings of: IEEE Nuclear Science Symposium and Medical Imaging Conference 1999; Seattle, WA. 1999. p. 1260-1264.
11. Suetens, P. Fundamentals of medical imaging. Cambridge: Cambridge University Press; 2002. p. 204
12. Levin CS, Hoffman EJ. *Phys Med Biol* 1999;44:781–799. [PubMed: 10211810]
13. Palmer MR, Brownell GL. *IEEE Trans Med Im* 1992;11:373–378.
14. Harrison, RL., et al. Proceedings of: Nuclear Science Symposium and Medical Imaging Conference 1999; Seattle, WA. 1999. p. 1265-1268.
15. Evans, RD. The Atomic Nucleus. New York: McGraw Hill; 1955. p. 629-630.
16. Ibid., pp. 674-676.
17. Perkins S, Cullen D. ENDL Type Formats for the LLNL Evaluated Atomic Data Library, EADL, for the Evaluated Electron Data Library, EEDL, and for the Evaluated Photon Data Library, EPDL. UCRL-ID 117796 1994:1–27.
18. Kaplan MS, Harrison RL, Vannoy SD. *IEEE Trans Nucl Sci* 1998;45:3064–8.
19. Anger HO. *J Nucl Med* 1964;65:515–31. [PubMed: 14216630]
20. Haynor DR, Harrison RL, Lewellen TK. *Med Phys* 1991;18:990–1001. [PubMed: 1961165]
21. United States National Library of Medicine; Aug 28. 2009 The Visible Human Project®.
http://www.nlm.nih.gov/research/visible/visible_human.html
22. Bushberg, JT. The essential physics of medical imaging. Philadelphia: Lippincott Williams & Wilkins; 2002. p. 370
23. Stieber, VW., et al. Central Nervous System Tumors. In: Levitt, SH., et al., editors. Technical Basis of Radiation Therapy: Practical Clinical Applications. New York: Springer; 2008.
24. Aird EGA, Conway J. *Brit J Radiol* 2002;75:937–949. [PubMed: 12515702]

FIGURE 1.

(A) Schematic of X-ray Computed Tomography (CT) [5]; (B) Schematic of Positron Emission Tomography (PET) [6]. In X-ray CT, a beam of photons is transmitted through the patient and the detectors measure the attenuation of the beam; in PET, a radioactively labeled pharmaceutical is introduced into the patient and the tomograph counts positron-electron annihilations by looking for the (nearly) simultaneous detection of the two photons from the annihilation.

FIGURE 2.

(A) X-ray CT, (B) PET, and (C) Fused PET/CT Images of the Pelvis. Metastatic colorectal cancer shows no contrast from the surrounding tissue in X-ray CT, but shows up as high-contrast hot spots in an FDG PET scan [7].

FIGURE 3.

Simulation Trade-offs. In general, flexibility and accuracy comes at the cost of ease-of-use and computational speed.

FIGURE 4. Cutaway View of Half a PET Scanner With and Without Septa [11]. (Fundamentals of Medical Imaging, 2nd Edition, by Paul Suetens, © 2009 P. Suetens. Reprinted with the permission of Cambridge University Press.)

FIGURE 5.
SimSET Offers Two Detector Models For PET, Cylindrical Detectors And Block Detectors.

Beam of 70 keV photons

% detected?

FIGURE 6.

In X-ray CT We Measure the Fraction of an Incident Beam of X-rays that Reaches a Detector on the Opposite Side of the Patient [21].

FIGURE 7. Beam Hardening. The average energy of the beam becomes higher as it passes through the patient [22].

FIGURE 8.

Beam Arrangement for a Tumor in the Thalamic Region [23]. The radiation dose is split into several different beams focusing on the tumor so that the tumor gets a considerably larger dose than the surrounding tissue.

FIGURE 9.
In Radiotherapy Treatment Planning We Are Interested in Dose Deposited in the Patient, Organ by Organ.