
Medical Imaging

Charles L. Epstein

1 Introduction

Over the past fifty years the processes and techniques of medical imaging have undergone a veritable explosion, calling into service increasingly sophisticated mathematical tools. Mathematics provides a language to describe the measurement processes that lead, eventually, to algorithms for turning the raw data into high-quality images. There are four principal modalities in wide application today: x-ray computed tomography (x-ray CT), ultrasound, magnetic resonance imaging (MRI), and emission tomography (positron emission tomography (PET) and single-photon emission computed tomography (SPECT)). Each modality uses a different physical process to produce image contrast: x-ray CT produces a map of the x-ray attenuation coefficient, which is strongly correlated with density; ultrasound images are produced by mapping absorption and reflection of acoustic waves; in their simplest form, magnetic resonance images show the density of water protons, but the subtlety of the underlying physics provides many avenues for producing clinically meaningful contrasts in this modality; PET and SPECT give spatial maps of the chemical activity of metabolites, which are bound to radioactive elements. It has recently been found useful to merge different modalities. For example, a fused MRI/PET image shows metabolic activity produced by PET, at a fairly low spatial resolution, against the background of a detailed anatomic image produced by MRI. Figure 1 shows a PET image, a PET image fused with a CT image, and the CT image as well.

In this article we consider mathematical aspects of PET, whose underlying physics we briefly explain. Positron emission is a mode of radioactive decay stemming from the reaction



Two isotopes, of clinical importance, that undergo this type of decay are F^{18} and C^{11} . The positron, which is the positively charged antiparticle of the electron, is typically very short-lived as it is annihilated, along with the first electron it encounters, producing a pair of 0.511 MeV photons. This usually happens within a millimeter or two of the site of the radioactive decay. Due to conservation of momentum, these two photons travel in nearly opposite directions along a straight line

(see figure 2). The phenomenon of pair annihilation underlies the operation of a PET scanner.

A short-lived isotope that undergoes the reaction in (1) is incorporated into a metabolite, e.g., fluorodeoxyglucose, which is then injected into the patient. This metabolite is taken up differentially by various structures in the body. For example, many types of cancerous tumors have a very rapid metabolism and quickly take up available fluorodeoxyglucose. The detector in a PET scanner is a ring of scintillation crystals that surrounds some portion of the patient. The high-energy photon interacts with the crystal to produce a flash of light. These flashes are fed into photomultiplier tubes with electronics that localize, to some extent, where the flash of light occurred and measure the energy of the photon that produced it. Finally, different arrival times are compared to determine which events are likely to be “coincidences,” caused by a single pair annihilation. Two photons detected within a time window of about 10 nanoseconds are assumed to be the result of a single annihilation event. The measured locations of a pair of coincident photons then determines a line. If the photons simply exited the patient’s body without further interactions, then the annihilation event must have occurred somewhere along this line (see figure 2). It is not difficult to imagine that sufficiently many such measurements could be used to reconstruct an approximation for the distribution of sources, a goal which is facilitated by a more quantitative model.

2 A Quantitative Model

Radioactive decay is usually modeled as a Poisson random process. Recall that Y is a Poisson random variable of intensity λ if

$$\text{Prob}(Y = k) = \frac{\lambda^k e^{-\lambda}}{k!}. \quad (2)$$

A simple calculation shows that $E[Y] = \lambda$ and $\text{Var}[Y] = \lambda$ as well. Let H denote the region within the scanner that is occupied by the patient, and, for $p \in H$, let $\rho(p)$ denote the concentration of radioactive metabolite as a function of position. If ρ is measured in the correct units, then the probability of k decay events originating from a small volume dV centered at p , in a time interval of unit length, is

$$\text{Prob}(k; p) = \frac{\rho(p)^k e^{-\rho(p)}}{k!} dV. \quad (3)$$

Decays originating at different spatial locations are regarded as independent events.




Figure 1 (a) A PET “heat map” image; (b) the image in (a) fused with the CT scan of the same section shown in (c). From the fused image it is apparent that the increased uptake of fluorodeoxyglucose, indicated by the yellow arrow, is in the gall bladder and is not the result of bowel activity. Images courtesy of Dr. Joel Karp, Hospital of the University of Pennsylvania.

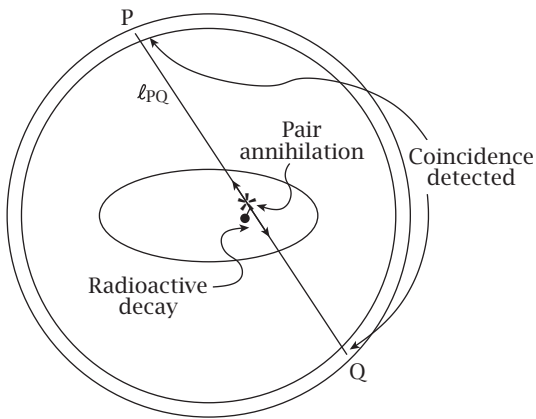


Figure 2 A radioactive decay leading to a positron–electron annihilation, exiting along ℓ_{PQ} , which is detected as a coincidence event at P and Q in the detector ring.

Assume, for the moment, that

- (i) there are many decay events, so that we are justified in replacing this probabilistic law by its mean, $\rho(p)$,
- (ii) the high-energy photons simply exit the patient without interaction, and
- (iii) we are equally likely to detect a given decay event on any line passing through the source point.

Let ℓ_{PQ} be the line joining the two detector positions P and Q where photons are simultaneously detected. With these assumptions we see that by counting up the coincidences observed at P and Q we are finding an approximation to the line integral

$$X\rho(\ell_{PQ}) = \int_{\ell_{PQ}} \rho(p) dl,$$

here dl is the arc length along the line ℓ_{PQ} . This is nothing other than a sample of the three-dimensional x-ray




Figure 3 The lines in the plane π_0 are labeled by θ , the angle the normal makes with the x -axis, and s , the distance from the line to the origin.

transform of ρ , which, if it could be approximately measured with sufficient accuracy, for a sufficiently dense set of lines, could then be inverted to produce a good approximate value for ρ . This is essentially what is done in X-RAY CT [??].

For the moment, we restrict our attention to lines that lie in a plane π_0 intersecting the patient and choose coordinates (x, y, z) so that $\pi_0 = \{z = z_0\}$. The lines in this plane are parametrized by an angle $\theta \in [0, \pi]$ and a real number s , with

$$\ell_{\theta,s} = \{(s \cos \theta, s \sin \theta, z_0) + t(-\sin \theta, \cos \theta, 0) : t \in (-\infty, \infty)\}$$

(see figure 3). In terms of s and θ , the two-dimensional x-ray transform is given by the integral

$$X\rho(s, \theta, z_0) = \int_{-\infty}^{\infty} \rho((s \cos \theta, s \sin \theta, z_0) + t(-\sin \theta, \cos \theta, 0)) dt.$$

The inverse of this transform is usually represented as a composition of two operations: a filter acting on $X\rho(s, \theta, z_0)$ in the s variable, followed by the back-projection operator. If $g(s, \theta)$ is a function on the space

of lines in a plane, then the filter operation can be represented by $\mathcal{F}g(s, \theta) = \partial_s \mathcal{H}g(s, \theta)$, where \mathcal{H} is a constant multiple of the Hilbert transform acting in the s variable. The back-projection operator X^*g defines a function of $(x, y) \in \pi_0$ that is the average of g over all lines passing through (x, y) :

$$X^*g(x, y) = \frac{1}{\pi} \int_0^\pi g(x \cos \theta + y \sin \theta, \theta) d\theta.$$

Putting together the pieces we get the filtered back-projection (FBP) operator, which inverts the two-dimensional x-ray transform: $\rho(x, y, z_0) = [X^* \circ \mathcal{F}] \cdot X\rho$. By using this approach for a collection of parallel planes, the function ρ could be reconstructed in a volume. This provides a possible method for reconstruction of PET images, and indeed the discrete implementations of this method have been extensively studied. In the early days of PET imaging this approach was widely used, and it remains in use today. Note, however, that using only data from lines lying in a set of parallel planes is very wasteful and leads to images with low signal to noise ratio.

Assumption (i) implies that our measurement is a good approximation to the x-ray transform of ρ , $X\rho(s, \theta, z_0)$. Because of the very high energies involved in positron emission radioactivity, only very small amounts of short-lived isotopes can be used. The measured count rates are therefore low, which leads to measurements dominated by Poisson noise that are not a good approximation to the mean. Because the FBP algorithm involves a derivative in s , the data must be significantly smoothed before this approach to image reconstruction can be applied. This produces low-resolution images that contain a variety of artifacts due to systematic measurement errors, which we describe below.

At this point it is useful to have a more accurate description of the scanner and the measured data. We model the detector as a cylindrical ring surrounding the patient, which is partitioned into a finite set of regions $\{d_1, \dots, d_n\}$. The scanner can localize a scintillation event as having occurred in one of these regions, which we heretofore refer to as detectors. This instrument design suggests that we divide the volume inside the detector ring into a collection of tubes, $\{T_{ij}\}$, with each tube defined as the union of lines joining points in d_i to points in d_j (see figure 4). A measurement n_{ij} is the number of coincidence events observed by the pair of detectors (d_i, d_j) . The simplest interpretation of n_{ij} is as a sample of a Poisson random variable with mean




Figure 4 The detector ring is divided into finitely many detectors of finite size. Each pair (d_i, d_j) defines a tube T_{ij} in the region occupied by the patient. This region is divided into boxes $\{b_k\}$.

proportional to

$$\int_{\ell_{PQ} \subset T_{ij}} X\rho(\ell_{PQ}). \quad (4)$$

Below we will see that this interpretation requires several adjustments.

Assumption (ii) fails as the photons tend to interact quite a lot with the bulk of the patient's body. Large fractions of the photons are absorbed, or scattered, with each member of an annihilation pair meeting its fate independently of the other. This leads to three distinct types of measurement errors.

Randoms. These are coincidences that are observed by a pair of detectors but that do not correspond to a single annihilation event. These can account for 10–30% of the observed events (see figure 5(a)).

Scatter. If one or both photons is/are scattered and then both are detected, this may register as a coincidence at a pair of detectors (d_i, d_j) , but the annihilation event did not occur at a point lying near T_{ij} (see figure 5(b)).

Attenuation. Most photon pairs (often 95%) are simply absorbed, leading to a substantial underestimate of the number of events occurring along a given line.

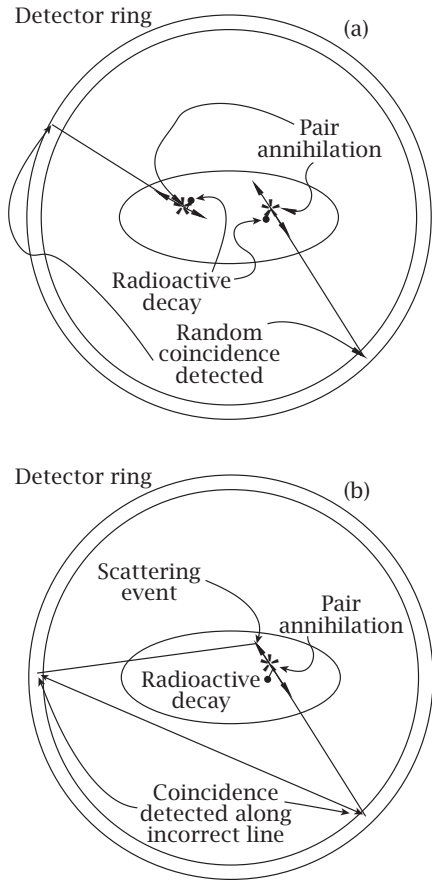


Figure 5 The measurement process in PET scanners is subject to a variety of systematic errors. (a) Randoms are detected coincidences that do not result from a single decay event. (b) Scatter is the result of one or both of the photons scattering off an object before being detected as a coincidence event, but along the wrong line.

Below we discuss how the effects of these sorts of measurement errors can be incorporated into the model and the reconstruction algorithm. To get quantitatively meaningful, artifact-free images, these errors must be corrected before application of any image-reconstruction method.

Assumption (iii) is false in that the detector array, which is usually a ring of scintillation counters, only encloses part of the patient. Many lines through the patient will therefore be disjoint from the detector, or only intersect it at one end. This problem can, to some extent, be mitigated by only using observations coming from lines that lie in planes that intersect the detector in a closed curve. If the detector is a section of a

cylinder, then each point p lies in a collection of such planes $\{\pi_{\psi,\phi}\}$ whose normal vectors $\{v_{\psi,\phi}\}$ fill a disk D_p lying on the unit sphere. If $\rho^{\psi,\phi}(p)$ denotes the approximate value for $\rho(p)$ determined using the FBP algorithm in the plane $\pi_{\psi,\phi}$, then an approximate value with improved signal to noise ratio is obtained as the average:

$$\bar{\rho}(p) = \frac{1}{|D_p|} \int_{D_p} \rho^{\psi,\phi}(p) dS(\psi, \phi),$$

where $dS(\psi, \phi)$ is the spherical areal measure. A particular implementation of this idea that is often used in PET scanners goes under the name of the ‘‘Colsher filter.’’ Other methods use a collection of parallel two-dimensional planes to reconstruct an approximate image from which the missing data for the three-dimensional x-ray transform can then be approximately computed.

In addition to these inherent physical limitations on the measurement process, there are a wide range of instrumentation problems connected to the detection and spatial localization of high-energy photons, as well as the discrimination of coincidence events. Effective solutions to these problems are central to the success of a PET scanner, but they are beyond the scope of this article.

3 Correcting Measurement Errors

To reconstruct images that are quantitatively meaningful and reasonably free of artifacts, the measured data $\{n_{ij}\}$ must *first* be corrected for randoms, scatter, and attenuation (see figure 5). This requires both additional measurements and models for the processes that lead to these errors.

3.1 Randoms

We first discuss how to correct for randoms. Let R_{ij} denote the number of coincidences detected on the pair (d_i, d_j) that are not caused by a decay event in T_{ij} . In practice, coincidences are considered to be two events that are observed within a certain time window τ (usually about 10 nanoseconds). In addition to coincidences between two detectors, the numbers of single counts, $\{n_i\}$, observed at $\{d_i\}$ are recorded. In fact, the number of ‘‘singles’’ is usually one or two *orders of magnitude* larger than the number of coincidences. From the measured number of singles observed over a known period of time we can infer rates of singles events $\{r_i\}$ for each detector. Assuming that each of these singles processes is independent, a reasonable estimate for the number

of coincidences observed on the detector pair (d_i, d_j) over the course of T units of time that are actually randoms is $R_{ij} \approx \tau Tr_i r_j$. A somewhat more accurate estimate is obtained if one accounts for the decay of the radioactive source.

There are other measurement techniques for estimating R_{ij} , though these estimates tend to be rather noisy. Simply subtracting R_{ij} from n_{ij} can increase the noise in the measurements and also change their basic statistical properties. There is a useful technique for replacing R_{ij} with a lower-variance estimate. Let A be a collection of contiguous detectors including d_i that are joined to d_j , and let B be a similar collection, including d_j , that are joined to d_i . Suppose that \tilde{R}_{mn} are estimates for the randoms detected in the pairs $\{(d_n, d_m) : n \in B, m \in A\}$. The expression

$$\hat{R}_{ij} = \frac{[\sum_{m \in A} \tilde{R}_{im}][\sum_{n \in B} \tilde{R}_{nj}]}{\sum_{m \in A, n \in B} \tilde{R}_{mn}}$$

provides an estimate for R_{ij} with reduced noise variance.

3.2 Scatter

The next source of error we consider is scatter, which results from one or both photons in the annihilation pair scattering off some matter in the patient before being recorded as a coincidence at a pair of detectors (d_i, d_j) . If the scattering angle is not small, then the annihilation event will not have occurred in expected tube T_{ij} . Some part of n_{ij} , which we denote S_{ij} , therefore corresponds to radioactive decays that did not occur in T_{ij} . Scattered photons tend to lose energy, so many approaches to estimating the amount of scatter are connected to measurement of the energies of detected photons. Depending on the design of the scanner, scatter can account for 15–40% of the observed coincidences. There are many methods for estimating the contribution of scatter but most of them are related to the specific design of the PET scanner and are therefore beyond the purview of this article.

3.3 Attenuation

Once contributions from randoms and scatter have been removed to obtain corrected observations $\hat{n}_{ij} = n_{ij} - R_{ij} - S_{ij}$, we still need to account for the fact that many photon pairs are simply absorbed. This process is described by Beer's law, which is the basis for x-ray CT. Suppose that an annihilation event takes place at a point p_0 within the patient, and that the photons travel

along the rays ℓ_+ and ℓ_- originating at p_0 . The attenuation coefficient is a function $\mu(p)$ defined throughout the patient's body such that the probabilities of detecting photons traveling along ℓ_{\pm} are

$$P_{\pm} = \exp\left(-\int_{\ell_{\pm}} \mu dl\right).$$

As we only "count" coincidences, and the two photons are independent, the probability of observing the coincidence is simply the product

$$P_+ P_- = \exp\left(-\int_{\ell} \mu dl\right),$$

where $\ell = \ell_+ \cup \ell_-$. In other words, the attenuation of coincidence counts due to photons traveling along ℓ does not depend on the location of the annihilation event along this line! This factor can therefore be measured by observing the fraction of photons, of the given energy, emitted *outside* the patient's body that pass through the patient along ℓ and are detected on the opposite side.

For each pair (d_i, d_j) we can therefore determine an attenuation coefficient q_{ij} . The extent of the intersection of T_{ij} with the patient's body has a marked effect on the size of q_{ij} , which can range from approximately 1 for tubes lying close to the skin to approximately 0.1 for those passing through a significant part of the body. The corrected data, which is passed to a reconstruction algorithm, is therefore

$$N_{ij} = \frac{\hat{n}_{ij}}{q_{ij}} = \frac{n_{ij} - R_{ij} - S_{ij}}{q_{ij}}.$$

In addition to the corrections described above, there are a variety of adjustments that are needed to account for measurement errors attributable to the details of the behavior of the detector and the operation of the electronics. Applying an FBP algorithm to the corrected data, we can obtain a discrete approximation $\rho^{\text{FBP}}(p)$ to $\rho(p)$. In the next section we describe iterative algorithms for PET image reconstruction. While the FBP algorithm is linear, and efficient, iterative algorithms allow for incorporation of more information about the measurement process and are better suited to low signal to noise ratio data.

4 Iterative Reconstruction Algorithms

While filtered back projection provides a good starting point for image reconstruction in PET, the varying statistical properties of different measurements cannot be easily incorporated into this algorithm. A variety of approaches have been developed that allow for

the exploitation of such information. To describe these algorithms we need to provide a discrete measurement model that is somewhat different from that discussed above. The underlying idea is that we are developing a statistical estimator for the strengths of the Poisson processes that produce the observed measurements. Note that these measurements must still be corrected as described in the previous section.

In the previous discussion we described the region, H , occupied by the patient as a continuum, with $\rho(p)$ the strength of the radioactive decay processes, a continuous function of a continuous variable $p \in H$. We now divide the measurement volume into a finite collection of boxes, $\{b_1, \dots, b_B\}$. The radioactive decay of the tracer in box b_k is modeled as a Poisson process of strength λ_k . For each point $p \in H$ and each detector pair (d_i, d_j) , we let $c(p; i, j)$ denote the probability that a decay event at p is detected as a coincidence in this detector pair. The patient's body will produce scatter and attenuation that will in turn alter the values of $c(p; i, j)$ from what they would be in its absence, i.e., the area fraction of a small sphere centered at p intercepted by lines joining points in d_i to points in d_j .

In the simplest case, the measurements $\{n_{ij}\}$ would be interpreted as samples of Poisson random variables with means

$$\sum_{k=1}^B p(k; i, j) \lambda_k,$$

where

$$p(k; i, j) = \frac{1}{V(b_k)} \int_{b_k} c(p; i, j) dp$$

is the probability that a decay event in b_k is detected in the pair (d_i, d_j) . Here, $V(b_k)$ is the volume of b_k . Assuming that the attenuation coefficient does not vary rapidly within the tube T_{ij} , we can incorporate attenuation into this model, as above, by replacing $p(k; i, j)$ with $q_{ij} p(k; i, j) = p^a(k; i, j)$. Ignoring scatter and randoms, the expected value of n_{ij} would then satisfy

$$E[n_{ij}] = \sum_{k=1}^B p^a(k; i, j) \lambda_k = \bar{n}_{ij}.$$

In this model, scatter and randoms are regarded as independent Poisson processes, with means $\lambda^s(i, j)$ and $\lambda^r(i, j)$, respectively. Including these effects, we see that the measurement n_{ij} is then a sample of a Poisson random variable with mean $\bar{n}_{ij} + \lambda^s(i, j) + \lambda^r(i, j)$. The reconstruction problem is then to infer estimates for

the intensities of the sources $\{\lambda_k\}$ from the observations $\{n_{ij}\}$. There are a variety of approaches to solving this problem.

First we consider the reconstruction problem ignoring the contributions of scatter and randoms. The measurement model suggests that we look for a solution, $(\lambda_1^*, \dots, \lambda_B^*)$, to the system of equations

$$n_{ij} = \sum_{k=1}^B p^a(k; i, j) \lambda_k^*.$$

If the array of detectors is three dimensional, there are likely to be many more detector pairs than boxes in the volume. The number of such pairs is quadratic in the number of detectors. This system of equations is therefore highly overdetermined, so a least-squares solution is a reasonable choice. That is, λ^* could be defined as

$$\lambda^* = \arg \min_{\{\gamma: 0 \leq \gamma_k\}} \sum_{i,j} \left(n_{ij} - \sum_{k=1}^B p^a(k; i, j) \gamma_k \right)^2.$$

Note that we constrain the variables $\{\gamma_k\}$ to be non-negative, as this is certainly true of the actual intensities. The least-squares solution can be interpreted as a maximum-likelihood (ML) estimate for λ , when the likelihood of observing n given the intensities γ is given by the product of Gaussians:

$$L_g(\gamma) = \prod_{i,j} \exp \left[- \left(n_{ij} - \sum_{k=1}^B p^a(k; i, j) \gamma_k \right)^2 \right].$$

It is assumed that the various observations are samples of independent processes. If we have estimates for the variances $\{\sigma_{ij}\}$ of these measurements, then we could instead consider a weighted least-squares solution and look for

$$\lambda_{g,\sigma}^* = \arg \min_{\{\gamma: 0 \leq \gamma_k\}} \sum_{i,j} \frac{1}{\sigma_{ij}} \left(n_{ij} - \sum_{k=1}^B p^a(k; i, j) \gamma_k \right)^2.$$

Because the data tends to be very noisy, in addition to the "data term" many algorithms include a regularization term, such as

$$R(\gamma) = \sum_{k=1}^B \sum_{k \in N(k)} |\gamma_k - \gamma_{k'}|^2,$$

where for each k , the $N(k)$ are the indices of the boxes contiguous to b_k . The β -regularized solution is then defined as

$$\lambda_{g,\sigma,\beta}^* = \arg \min_{\{\gamma: 0 \leq \gamma_k\}} \left[\sum_{i,j} \frac{1}{\sigma_{ij}} \left(n_{ij} - \sum_{k=1}^B p^a(k; i, j) \gamma_k \right)^2 + \beta R(\gamma) \right].$$

As noted above, these tend to be very large systems of equations and are therefore usually solved via iterative methods. Indeed, a great deal of the research effort in PET is connected with finding data structures and algorithms to enable solution of such optimization problems in a way that is fast enough and stable enough for real-time imaging applications.

Given the nature of radioactive decay it is perhaps more reasonable to consider an expression for the likelihood in terms of Poisson processes. With

$$\mu(i, j)(\gamma) = \sum_k^B p^a(k; i, j) \gamma_k,$$

we get the Poisson likelihood function

$$L_p(\gamma) = \prod_{i,j} \frac{e^{-\mu(i,j)(\gamma)} [\mu(i,j)(\gamma)]^{n_{ij}}}{n_{ij}!}.$$

The expectation—maximization (EM) algorithm provides a means to iteratively find the nonnegative vector that maximizes $\log L_p(\gamma)$. After choosing a nonnegative starting vector $\gamma^{(0)}$, the map from $\gamma^{(m)}$ to $\gamma^{(m+1)}$ is given by the formula

$$\gamma_k^{(m+1)} = \gamma_k^{(m)} \frac{1}{\sum_{i,j} p^a(k; i, j)} \sum_{i,j} \left[\frac{p^a(k; i, j) n_{ij}}{\sum_k p^a(k; i, j) \gamma_k^{(m)}} \right]. \quad (5)$$

This algorithm has several desirable features. Firstly, if the initial guess $\gamma^{(0)}$ is positive, then the positivity condition on the components of $\gamma^{(m)}$ is automatic. Secondly, before convergence, the likelihood is monotonically increasing; that is, $L(\gamma^{(m)}) < L(\gamma^{(m+1)})$. Note that if $n_{ij} = \sum_k p^a(k; i, j) \gamma_k^{(m)}$ for all (i, j) , then $\gamma^{(m+1)} = \gamma^{(m)}$. The algorithm defined in (5) converges too slowly to be practical in clinical applications. There are several methods to accelerate its convergence, which also include regularization terms.

We conclude this discussion by explaining how to include estimates for the contributions of scatter and randoms to n_{ij} in an ML reconstruction algorithm. We suppose that n_{ij} can be decomposed as a sum of three terms:

$$n_{ij} = \sum_{k=1}^B p^a(k; i, j) \lambda_k + R_{ij} + S_{ij}. \quad (6)$$

With this decomposition, it is clear how to modify the update rule in (5):

$$\gamma_k^{(m+1)} = \gamma_k^{(m)} \frac{1}{\sum_{i,j} p^a(k; i, j)} \times \sum_{i,j} \left[\frac{p^a(k; i, j) n_{ij}}{\sum_k p^a(k; i, j) \gamma_k^{(m)} + R_{ij} + S_{ij}} \right].$$




Figure 6 Two reconstructions from the same PET data illustrating the superior noise and artifact suppression attainable using iterative algorithms: (a) image reconstructed using the FBP algorithm and (b) image reconstructed using an iterative ML algorithm. Images courtesy of Dr. Joel Karp, Hospital of the University of Pennsylvania.

In addition to the ML-based algorithms, there are many other iterative approaches to solving these optimization problems that go under the general rubric of “algebraic reconstruction techniques,” or ART. Figure 6 shows two reconstructions of a PET image: part (a) is the result of using the FBP algorithm, while part (b) shows the output of an iterative ML-EM algorithm.

5 Outlook

PET imaging provides a method for directly visualizing and spatially localizing metabolic processes. As is clear from our discussion, the physics involved in interpreting the measurements and designing detectors is rather complicated. In this article we have only touched on some of the basic ideas used to model the measurements and develop reconstruction algorithms. The FBP algorithm gives the most direct method for reconstructing images, but the images tend to have low resolution, streaking artifacts, and noise. One can easily incorporate much more of the physics into iterative techniques based on probabilistic models, and this should lead to much better images. Because of the large number of detector pairs for three-dimensional volumes, naive implementations of iterative algorithms require vast computational resources. At the time of writing, both reconstruction techniques and the modality as a whole are rapidly evolving in response to the development of better detectors and faster computers, and because of increased storage capabilities.

Further Reading

A comprehensive overview of PET is given in Bailey et al. (2005). A review article on instrumentation and clinical applications is Muehllehner and Karp (2006). An early article on three-dimensional reconstruction is Colsher (1980). An early paper on the use of ML algorithms in PET is Vardi et al. (1985). Review articles on reconstruction methods from projections and reconstruction methods in PET are Lewitt and Matej (2003) and Reader and Zaidi (2007), respectively.

Bailey, D. L., D. W. Townsend, P. E. Valk, and M. N. Maisey, eds. 2005. *Positron Emission Tomography*. London: Springer.

Colsher, J. G. 1980. Fully three-dimensional positron emission tomography. *Physics in Medicine and Biology* 25: 103–15.

Lewitt, R. M., and S. Matej. 2003. Overview of methods for image reconstruction from projections in emission computed tomography. *Proceedings of the IEEE* 91:1588–611.

Muehllehner, G., and J. Karp. 2006. Positron emission tomography. *Physics in Medicine and Biology* 51:R117–37.

Reader, A. J., and H. Zaidi. 2007. Advances in PET image reconstruction. *PET Clinics* 2:173–90.

Vardi, Y., L. A. Shepp, and L. Kaufman. 1985. A statistical model for positron emission tomography. *Journal of the American Statistical Association* 80:8–20.