

CHAPTER 13

IMAGE RECONSTRUCTION

J. NUYTS

Department of Nuclear Medicine and Medical Imaging Research Center,
Katholieke Universiteit Leuven,
Leuven, Belgium

S. MATEJ

Medical Image Processing Group,
Department of Radiology,
University of Pennsylvania,
Philadelphia, Pennsylvania,
United States of America

13.1. INTRODUCTION

This chapter discusses how 2-D or 3-D images of tracer distribution can be reconstructed from a series of so-called projection images acquired with a gamma camera or a positron emission tomography (PET) system [13.1]. This is often called an ‘inverse problem’. The reconstruction is the inverse of the acquisition. The reconstruction is called an inverse problem because making software to compute the true tracer distribution from the acquired data turns out to be more difficult than the ‘forward’ direction, i.e. making software to simulate the acquisition.

There are basically two approaches to image reconstruction: analytical reconstruction and iterative reconstruction. The analytical approach is based on mathematical inversion, yielding efficient, non-iterative reconstruction algorithms. In the iterative approach, the reconstruction problem is reduced to computing a finite number of image values from a finite number of measurements. That simplification enables the use of iterative instead of mathematical inversion. Iterative inversion tends to require more computer power, but it can cope with more complex (and hopefully more accurate) models of the acquisition process.

13.2. ANALYTICAL RECONSTRUCTION

The (n -dimensional) radon transform maps an image of dimension n to the set of all integrals over hyperplanes of dimension $(n - 1)$ [13.2]. Thus, in two dimensions, the radon transform of image Λ corresponds to all possible line integrals of Λ . In three dimensions, the radon transform contains all possible plane integrals.

The (n -dimensional) X ray transform maps an image of dimension n to the set of all possible line integrals. In all PET and in almost all single photon emission computed tomography (SPECT) applications, the measured projections can be well approximated as a subset of the (possibly attenuated) X ray transform, because the mechanical (SPECT) or electronic (PET) collimation is designed to acquire information along lines (the line of response (LOR), see Chapter 11). Consequently, reconstruction involves computing the unknown image Λ from (part of) its X ray transform. Figure 13.1 shows PET projections, which are often represented as a set of projections or a set of sinograms.




FIG. 13.1. The relation between projections and sinograms in parallel-beam projection. The parallel-beam (PET) acquisition is shown as a block with dimensions s , ϕ and z . A cross-section at fixed ϕ yields a projection; a cross-section at fixed z yields a sinogram.

An important theorem for analytical reconstruction is the central slice (or central section) theorem, which gives a relation between the Fourier transform of an image and the Fourier transforms of its parallel projections. Below, the central slice theorem for 2-D is found as Eq. (13.7) and the 3-D central section theorem as Eq. (13.29).

The direct Fourier method is a straightforward application of the central section theorem: it computes the Fourier transform of the projections, uses the central section theorem to obtain the Fourier transform of the image and applies the inverse Fourier transform to obtain the image. In practice, this method is rarely used; the closely related filtered back projection (FBP) algorithm is far more popular.

13.2.1. Two dimensional tomography

13.2.1.1. X ray transform: projection and back projection

In 2-D, the radon transform and X ray transform are identical. Mathematically, the 2-D X ray (or radon) transform of the image Λ can be written as follows:

$$\begin{aligned} Y(s, \phi) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Lambda(x, y) \delta_{s=x \cos \phi + y \sin \phi} dx dy \\ &= \int_{-\infty}^{\infty} \Lambda(s \cos \phi - t \sin \phi, s \sin \phi + t \cos \phi) dt \end{aligned} \quad (13.1)$$

where the δ function is unity for the points on the LOR (s, ϕ) and zero elsewhere. It should be noted that with the notation used here, $\phi = 0$ corresponds to projection along the y axis.

The radon transform describes the acquisition process in 2-D PET and in SPECT with parallel-hole collimation, if attenuation can be ignored. Assuming that $\Lambda(x, y)$ represents the tracer distribution at transaxial slice Z through the patient, then $Y(s, \phi)$ represents the corresponding sinogram, and contains the z -th row of the projections acquired at angles ϕ . Figure 13.1 illustrates the relation between the projection and the sinogram.

The X ray transform has an adjoint operation that appears in both analytical and iterative reconstruction. This operator is usually called the back projection operator, and can be written as:

$$\begin{aligned} B(x, y) &= \text{Backproj}(Y(s, \phi)) \\ &= \int_0^{\pi} d\phi \int_{-\infty}^{\infty} Y(s, \phi) \delta_{s=x \cos \phi + y \sin \phi} ds \\ &= \int_0^{\pi} Y(x \cos \phi + y \sin \phi, \phi) d\phi \end{aligned} \quad (13.2)$$

The back projection is not the inverse of the projection, $B(x, y) \neq \Lambda(x, y)$. Intuitively, the back projection sends the measured activity back into the image by distributing it uniformly along the projection lines. As illustrated in Fig. 13.2, projection followed by back projection produces a blurred version of the original image. This blurring corresponds to the convolution of the original image with the 2-D convolution kernel $1/\sqrt{x^2 + y^2}$.




FIG. 13.2. The image (left) is projected to produce a sinogram (centre), which in turn is back projected, yielding a smoothed version of the original image.

13.2.1.2. Central slice theorem

The central slice theorem gives a very useful relation between the 2-D Fourier transform of the image and the 1-D Fourier transform of its projections (along the detector axis). Consider the projection along the y axis, $\phi = 0$, and its 1-D Fourier transform:

$$Y(s,0) = \int_{-\infty}^{\infty} \Lambda(s,t) dt \quad (13.3)$$

$$\begin{aligned} (\mathcal{F}_1 Y)(v_s,0) &= \int_{-\infty}^{\infty} Y(s,0) e^{-i2\pi v_s s} ds \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Lambda(s,t) e^{-i2\pi v_s s} dt ds \end{aligned} \quad (13.4)$$

and compare this to the 2-D Fourier transform of the image $\Lambda(x, y)$:

$$(\mathcal{F}_2\Lambda)(v_x, v_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Lambda(x, y) e^{-i2\pi(v_x x + v_y y)} dx dy \quad (13.5)$$

Both expressions are equal if we set $v_y = 0$:

$$(\mathcal{F}_1 Y)(v_s, 0) = (\mathcal{F}_2\Lambda)(v_x, 0) \quad (13.6)$$

$(\mathcal{F}_1 Y)(v_s, 0)$ is the 1-D Fourier transform of the projection along the y axis and $(\mathcal{F}_2\Lambda)(v_x, 0)$ is a ‘central slice’ along the v_x axis through the 2-D Fourier transform of the image. Equation (13.6) is the central slice theorem for the special case of projection along the y axis. This result would still hold if the object had been rotated or equivalently, the x and y axes. Consequently, it holds for any angle ϕ :

$$(\mathcal{F}_1 Y)(v_s, \phi) = (\mathcal{F}_2\Lambda)(v_s \cos \phi, v_s \sin \phi) \quad (13.7)$$

13.2.1.3. Two dimensional filtered back projection

The central slice theorem (Eq. (13.7)) can be directly applied to reconstruct an unknown image $\Lambda(x, y)$ from its known projections $Y(s, \phi)$. The 1-D Fourier transform of the projections provides all possible central slices through $(\mathcal{F}_2\Lambda)(v_x, v_y)$ if $Y(s, \phi)$ is known for all ϕ in an interval with a length of at least π (Tuy’s condition). Consequently, $(\mathcal{F}_2\Lambda)(v_x, v_y)$ can be constructed from the 1-D Fourier transform of $Y(s, \phi)$. Inverse 2-D Fourier transform then provides $\Lambda(x, y)$.

However, a basic Fourier method implementation with a simple interpolation in Fourier space does not work well. In contrast, in the case of the FBP algorithm derived below, a basic real-space implementation with a simple convolution and a simple interpolation in the back projection works well. Inverse Fourier transform of Eq. (13.5) yields:

$$\Lambda(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (\mathcal{F}_2\Lambda)(v_x, v_y) e^{i2\pi(v_x x + v_y y)} dv_x dv_y \quad (13.8)$$

This can be rewritten with polar coordinates as:

$$\Lambda(x, y) = \int_{-\infty}^{\infty} dv \int_0^{\pi} (\mathcal{F}_2\Lambda)(v \cos \phi, v \sin \phi) e^{i2\pi(xv \cos \phi + yv \sin \phi)} |v| d\phi \quad (13.9)$$

Application of the central slice theorem (Eq. (13.7)) and reversing the order of integration finally results in:

$$\Lambda(x, y) = \int_0^\pi d\phi \int_{-\infty}^\infty (\mathcal{F}_1 Y)(v, \phi) |v| e^{i2\pi v(x \cos \phi + y \sin \phi)} dv \quad (13.10)$$

which is the FBP algorithm. This algorithm involves the following steps:

- (a) Apply 1-D Fourier transform to $Y(s, \phi)$ to obtain $(\mathcal{F}_1 Y)(v, \phi)$;
- (b) Filter $(\mathcal{F}_1 Y)(v, \phi)$ with the so-called ramp filter $|v|$;
- (c) Apply the 1-D inverse Fourier transform to obtain the ramp filtered projections $\hat{Y}(s, \phi) = \int (\mathcal{F}_1 \Lambda)(v, \phi) |v| e^{i2\pi vs} dv$;
- (d) Apply the back-projection operator Eq. (13.2) to $\hat{Y}(s, \phi)$ to obtain the desired image $\Lambda(x, y)$.

It should be noted that the ramp filter sets the DC component (i.e. the amplitude of the zero frequency) of the image to zero, while the mean value of the reconstructed image should definitely be positive. As a result, straightforward discretization of FBP causes significant negative bias. The problem is reduced with ‘zero padding’ before computing the Fourier transform with fast Fourier transform (FFT). Zero padding involves extending the sinogram rows with zeros at both sides. This increases the sampling in the frequency domain and results in a better discrete approximation of the ramp filter. However, a huge amount of zero padding is required to effectively eliminate the bias completely. The next paragraph shows how this need for zero padding can be easily avoided. It should be noted that after inverse Fourier transform, the extended region may be discarded, so the size of the filtered sinogram remains unchanged.

Instead of filtering in the Fourier domain, the ramp filtering can also be implemented as a 1-D convolution in the spatial domain. For this, the inverse Fourier transform of $|v|$ is required. This inverse transform actually does not exist, but approximating it as the limit for $\varepsilon \rightarrow 0$ of the well behaved function $|v|e^{-\varepsilon|v|}$ gives [13.3, 13.4]:

$$\mathcal{F}^{-1}(|v|e^{-\varepsilon|v|}) = \frac{\varepsilon^2 - (2\pi s)^2}{(\varepsilon^2 + (2\pi s)^2)^2} \quad (13.11)$$

$$\approx -\frac{1}{(2\pi s)^2} \quad \text{for } |s| \gg \varepsilon \quad (13.12)$$

In practice, band limited functions are always worked with, implying that the ramp filter has to be truncated at the frequencies $v = \pm 1/(2\tau)$, where τ represents

the sampling distance. The corresponding convolution kernel h then equals [13.3]:

$$h(s) = \mathcal{F}^{-1}(|v|b(v)) = \frac{1}{2\tau^2} \frac{\sin(\pi s / \tau)}{\pi s / \tau} - \frac{1}{4\tau^2} \left(\frac{\sin(\pi s / (2\tau))}{\pi s / (2\tau)} \right)^2 \quad (13.13)$$

$$\begin{aligned} \text{with } b(v) &= 1 \text{ if } |v| \leq 1/(2\tau) \\ &= 0 \text{ if } |v| > 1/(2\tau) \end{aligned}$$

The kernel is normally only needed for samples $s = n\tau$: $h(n\tau) = 1/(4\tau^2)$ if $n = 0$, $h(n\tau) = 0$ if n is even and $h(n\tau) = -1/(n\pi\tau)^2$ if n is odd. The filter can either be implemented as a convolution or the Fourier transform can be used to obtain a digital version of the ramp filter. Interestingly, this way of computing the ramp filter also reduces the negative bias mentioned above. The reason is that this approach yields a non-zero value for the DC component [13.3]. When the filtering is done in the frequency domain, some zero padding before FFT is still recommended because of the circular convolution effects, but far less is needed than with straightforward discretization of $|v|$.

Although this is not obvious from the equations above, an algorithm equivalent to FBP is obtained by first back projecting $Y(s, \phi)$ and then applying a 2-D ramp filter to the back projected image $B(x, y)$ [13.4]:

$$B(x, y) = \int_0^\pi Y(x \cos \phi + y \sin \phi, \phi) d\phi \quad (13.14)$$

$$(\mathcal{F}_2 \Lambda)(v_x, v_y) = \sqrt{v_x^2 + v_y^2} (\mathcal{F}_2 B)(v_x, v_y) \quad (13.15)$$

This algorithm is often referred to as the ‘back project-then-filter’ algorithm.

FBP assumes that the projections $Y(s, \phi)$ are line integrals. As discussed in Chapter 11, PET and SPECT data are not line integrals because of attenuation, detector non-uniformities, the contribution of scattered photons and/or random coincidences, etc. It follows that one has to recover (good estimates of) the line integrals by pre-correcting the data for these effects. However, a particular problem is posed by the attenuation in SPECT because, different from PET, the attenuation depends on the position along the projection line, precluding

straightforward pre-correction. A detailed discussion of this problem is beyond the scope of this contribution, but three solutions are briefly mentioned here:

- (a) If it can be assumed that the attenuation is constant inside a convex body contour, then FBP can be modified to correct for the attenuation. Algorithms have been proposed by Bellini, Tretiak, Metz and others; an algorithm is presented in Ref. [13.3].
- (b) If the attenuation is not constant, an approximate correction algorithm proposed by Chang can be applied [13.5]. It is a post-correction method, applied to the image obtained without any attenuation correction. To improve the approximation, the attenuated projection of the first reconstruction can be computed, and the method can be applied again to the difference of the measurement and this computed projection.
- (c) Finally, a modified FBP algorithm, compensating for non-uniform attenuation in SPECT, was found by Novikov in 2000. An equivalent algorithm was derived by Natterer [13.6]. However, because this algorithm was only found after the successful introduction of iterative reconstruction in clinical practice, it has not received much attention in the nuclear medicine community.

13.2.2. Frequency–distance relation

Several very interesting methods in image reconstruction, including Fourier rebinning, are based on the so-called frequency–distance relation, proposed by Edholm, Lewitt and Lindholm, and described in detail in Ref. [13.7]. This is an approximate relation between the orthogonal distance to the detector and the direction of the frequency in the sinogram. The relation can be intuitively understood as follows.

Consider the PET acquisition of a point source, as illustrated in Fig 13.3. Usually, the acquisition is described by rotating the projection lines while keeping the object stationary. However, here the equivalent description is considered, where projection is always along the y axis, and tomographic acquisition is obtained by rotating the object around the origin. Suppose that the point is located on the x axis when $\phi = 0$. When acquiring the parallel projections for angle ϕ , the point has polar coordinates (r, ϕ) , with r the distance from the centre of the field of view (FOV) and ϕ the angle with the x axis. The distance to the x axis is $d = r \sin \phi$. The corresponding sinogram $Y(s, \phi)$ is zero everywhere, except on the curve $s = r \cos \phi$ (Fig. 13.3). The complete sinogram is obtained by rotating the point over 360° $\phi = -\pi \dots \pi$. Consider a small portion of this curve, which can be well approximated as a tangential line segment near a particular point (s, ϕ) ,




FIG. 13.3. The frequency–distance principle. Left: sinogram; right: vertical projection of a point located at polar coordinates (r, ϕ) .

as illustrated in Fig. 13.3. In the 2-D Fourier transform of the sinogram, this line segment contributes mostly frequencies in the direction orthogonal to the line segment. This direction is represented by the angle α , given by:

$$\tan \alpha = \frac{\partial}{\partial \phi}(r \cos \phi) = -r \sin \phi = -d \tag{13.16}$$

Thus, in the 2-D Fourier transform $(\mathcal{F}g)(v_s, v_\phi)$, the value at a particular point (v_s, v_ϕ) carries mostly information about points located at a distance $d = -\tan \alpha = -v_\phi/v_s$ from the line through the centre, parallel to the detector. This relation can be exploited to apply distance dependent operations to the sinogram. One example is distance dependent deconvolution, to compensate for the distance dependent blurring in SPECT. Another example is Fourier rebinning, where data from oblique sinograms are rebinned into direct sinograms.

13.2.3. Fully 3-D tomography

13.2.3.1. Filtered back projection

Owing to the use of electronic collimation, the PET scanner can simultaneously acquire information in a 4-D space of line integrals. These are

the so-called LORs, where each pair of detectors in coincidence defines a single LOR. In this section, the discrete nature of the detection is ignored, since the analytical approach is more easily described assuming continuous data. Consider the X ray transform in 3-D, which can be written as:

$$Y(\hat{u}, s) = \int_{-\infty}^{\infty} \Lambda(s + t\hat{u}) dt \tag{13.17}$$

where the LOR is defined as the line parallel to \hat{u} and through the point s . The vector \hat{u} is a unit vector, and the vector s is restricted to the plane orthogonal to \hat{u} , hence (\hat{u}, s) is 4-D.

Most PET systems are either constructed as a cylindrical array of detectors or as a rotating set of planar detector arrays and, therefore, have cylindrical symmetry. For this reason, the inversion of Eq. (13.17) is studied for the case where \hat{u} is restricted to the band Ω_{θ_0} on the unit sphere, defined by $|u_z| \leq \sin\theta_0$, as illustrated in Fig. 13.4. It should be noted that only half of the sphere is actually needed because $Y(\hat{u}, s) = Y(-\hat{u}, s)$, but working with the complete sphere is more convenient.




FIG. 13.4. Each point on the unit sphere corresponds to the direction of a parallel projection. An ideal rotating gamma camera with a parallel-hole collimator only travels through the points on the equator. An idealized 3-D PET system would also acquire projections along oblique lines; it collects projections for all points of the set Ω . The set Ω , defined by θ_0 , is the non-shaded portion of the unit sphere. To recover a particular frequency v (of the Fourier transform of the object), at least one point on the circle C_v is required.

With $\theta_0 = 0$, the problem reduces to 2-D parallel projection (for multiple slices), which was shown to have a unique solution. It follows that with

$|\theta_0| > 0$, the problem becomes overdetermined, and there are infinitely many ways to compute the solution. This can be seen as follows. Each point of Ω corresponds to a parallel projection. According to the central slice theorem, this provides a central plane perpendicular to $\hat{\mathbf{u}}$ of the 3-D Fourier transform $\mathcal{L}(\mathbf{v})$ of $\Lambda(\mathbf{x})$. Thus, the set Ω_0 (i.e. all points on the equator of the unit sphere in Fig. 13.4) provides all planes intersecting the \mathbf{v}_z axis, which is sufficient to recover the entire image $\Lambda(\mathbf{x})$ via inverse Fourier transform. The set Ω_0 with $\theta_0 > 0$ provides additional (oblique) planes through $\mathcal{L}(\mathbf{v})$, which are obviously redundant. A simple solution would be to select a sufficient subset from the data. However, if the data are noisy, a more stable solution is obtained by using all of the measurements. This is achieved by computing $\mathcal{L}(\mathbf{v})$ from a linear combination of all available planes:

$$\mathcal{L}(\mathbf{v}) = \int_{\Omega_{\theta_0}} \mathcal{Y}(\hat{\mathbf{u}}, \mathbf{v}) H(\hat{\mathbf{u}}, \mathbf{v}) \delta(\hat{\mathbf{u}}, \mathbf{v}) d\hat{\mathbf{u}} \quad (13.18)$$

Here, $\mathcal{Y}(\hat{\mathbf{u}}, \mathbf{v})$ is the 2-D Fourier transform with respect to \mathbf{s} of the projection $Y(\hat{\mathbf{u}}, \mathbf{s})$. The Dirac function $\delta(\hat{\mathbf{u}}, \mathbf{v})$ selects the parallel projections $\hat{\mathbf{u}}$ which are perpendicular to \mathbf{v} (i.e. the points on the circle C_v in Fig. 13.4). Finally, the filter $H(\hat{\mathbf{u}}, \mathbf{v})$ assigns a particular weight to each of the available datasets $\mathcal{Y}(\hat{\mathbf{u}}, \mathbf{v})$. The combined weight for each frequency should equal unity, leading to the filter equation:

$$\int_{\Omega_{\theta_0}} H(\hat{\mathbf{u}}, \mathbf{v}) \delta(\hat{\mathbf{u}}, \mathbf{v}) d\hat{\mathbf{u}} = 1 \quad (13.19)$$

A solution equivalent to that of unweighted least squares (LS) is obtained by assigning the same weight to all available data [13.8]. This results in the Colsher filter which can be written as:

$$\begin{aligned} H_C(\hat{\mathbf{u}}, \mathbf{v}) &= |\mathbf{v}| / (2\pi) && \text{if } \sin\psi \leq \sin\theta_0 \\ &= |\mathbf{v}| / (4\arcsin(\sin\theta_0 / \sin\psi)) && \text{if } \sin\psi > \sin\theta_0 \end{aligned} \quad (13.20)$$

where ψ is the angle between \mathbf{v} and the z axis: $\mathbf{v}_z / |\mathbf{v}| = \cos\psi$. The direct Fourier reconstruction method can be applied here, by straightforward inverse Fourier transform of Eq. (13.18). However, an FBP approach is usually preferred, which can be written as:

$$\Lambda(\mathbf{x}) = \int_{\Omega_{\theta_0}} d\hat{\mathbf{u}} Y^F(\hat{\mathbf{u}}, \mathbf{x} - (\mathbf{x} \cdot \hat{\mathbf{u}})\hat{\mathbf{u}}) \quad (13.21)$$

Here, Y^F is obtained by filtering Y with the Colsher filter (or another filter satisfying Eq. (13.19): $Y^F(\hat{u}, s) = \mathcal{F}^{-1}(H_C(\hat{u}, \nu)Y(\hat{u}, \nu))$. The coordinate $s = \mathbf{x} - (\mathbf{x} \cdot \hat{u})\hat{u}$ is the projection of the point \mathbf{x} on the plane perpendicular to \hat{u} ; it selects the LOR through \mathbf{x} in the parallel projection \hat{u} .

13.2.3.2. The reprojection algorithm

The previous analysis assumed that the acceptance angle θ_0 was a constant, independent of \mathbf{x} . As illustrated in Fig. 13.5, this is not the case in practice. The acceptance angle is maximum for the centre of the FOV, becomes smaller with increasing distance to the centre and vanishes near the axial edges of the FOV. In other words, the projections are complete for \hat{u} orthogonal to the z axis (these are the 2-D multislice parallel-beam projections) and are truncated for the oblique parallel projections. The truncation becomes more severe for more oblique projections (Fig. 13.5).




FIG. 13.5. An axial cross-section through a cylindrical PET system, illustrating that the acceptance angle is position dependent (a). Oblique projections are truncated (b). In the reprojection algorithm, the missing oblique projections (dashed lines) are computed from a temporary multislice 2-D reconstruction (c).

As the acceptance angle is position dependent, the required filtering is position dependent as well, and cannot be implemented as a shift-invariant convolution (or Fourier filter). Several strategies for dealing with this truncation have been developed. One approach is to subdivide the image into a set of regions, and then optimize a shift-invariant filter in each of the regions. The filter is determined by the smallest acceptance angle of the region, so some of the data

will not be used. A good compromise between minimum data loss and practical implementation must be sought [13.9].

Another approach is to start with a first reconstruction, using the smallest acceptable angle over all positions \mathbf{x} in the FOV. This usually means that only the parallel projections orthogonal to the z axis are used. The missing oblique projection values are computed from this first reconstruction (Fig. 13.4) and used to complete the measured oblique projections. This eliminates the truncation, and the 3-D FBP method of the previous section can be applied. This method [13.10] was the standard 3-D PET reconstruction method for several years, until it was replaced by the faster Fourier rebinning approach (see below).

13.2.3.3. Rebinning techniques

The complexity (estimated as the number of LORs) increases linearly with the axial extent for 2-D PET, but quadratically for 3-D PET. To keep the processing time acceptable, researchers have sought ways to reduce the size of the data as much as possible, while minimizing the loss of information induced by this reduction.

Most PET systems have a cylindrical detector surface: the detectors are located on rings with radius R , and the rings are combined in a cylinder along the z axis. The data are usually organized in sinograms which can be written as:

$$Y_P(s, \phi, z, \Delta_z) = \int_{-\infty}^{\infty} dt \Lambda(s \cos \phi + t \hat{u}_x, s \sin \phi + t \hat{u}_y, z + t \hat{u}_z) \quad (13.22)$$

where $\hat{\mathbf{u}}$ is a unit vector in the direction of the LOR:

$$\hat{\mathbf{u}} = \mathbf{u} / \|\mathbf{u}\| \quad \text{with} \quad \mathbf{u} = (-\sin \phi, \cos \phi, \Delta_z / (2\sqrt{R^2 - s^2}))$$

The parameter s is the distance between the LOR and the z axis. The LOR corresponds to a coincidence between detector points with axial positions $z - \Delta_z/2$ and $z + \Delta_z/2$. Finally, ϕ is the angle between the y axis and the projection of the LOR on the xy plane. The coordinates (s, ϕ, z) are identical to those often used in 2-D tomography. It should be noted that, in practice, $s \ll R$ and, as a result, the direction of the LOR, the vector $\hat{\mathbf{u}}$, is virtually independent of s . In other words, a set of LORs with fixed Δ_z can then be treated as a parallel projection with good approximation. LORs with $\Delta_z = 0$ are often called ‘direct’ LORs, while LORs with $\Delta_z \neq 0$ are called ‘oblique’.

The basic idea of rebinning algorithms is to compute estimates of the direct sinograms from the oblique sinograms. If the rebinning algorithm is good, most of the information from the oblique sinograms will go into these estimates.

As a result, the data have been reduced from a complex 3-D geometry into a much simpler 2-D geometry without discarding measured signal. The final reconstruction can then be done with 2-D algorithms, which tend to be much faster than fully 3-D algorithms. A popular approach is to use Fourier rebinning, followed by maximum-likelihood reconstruction.

13.2.3.4. Single slice and multislice rebinning

The simplest way to rebin the data is to treat oblique LORs as direct LORs [13.11]. This corresponds to the approximation:

$$Y_p(s, \phi, z, \Delta_z) \approx Y_p(s, \phi, z, 0) \quad (13.23)$$

The approximation is only exact if the object consists of points located on the z axis, and it introduces mis-positioning errors that increase with increasing distance to the z axis and increasing Δ_z . Consequently, single slice rebinning is applicable when the object is small and positioned centrally in the scanner or when Δ_z is small. The axial extent of most current PET systems is too large to rebin an entire 3-D dataset with Eq. (13.23). However, single slice rebinning is used on all PET systems to reduce the sampling of the Δ_z dimension in the 3-D data, by combining sinograms with similar Δ_z . This typically reduces the data size by a factor of about ten, when compared to the finest possible sampling.

Application of Eq. (13.23) obviously causes blurring in the z direction, to a degree proportional to the distance from the z axis. However, it may also cause severe inconsistencies in the sinograms, producing blurring artefacts in the xy planes of the reconstructed images as well. Lewitt et al. [13.12] proposed distributing the oblique LOR values $Y_p(s, \phi, z, \Delta_z)$ over all LORs with $z \in [z - \Delta_z R_f / (2R), z + \Delta_z R_f / (2R)]$, i.e. over all slices intersected by the LOR, and within an FOV with radius R_f . This so-called multislice rebinning reduces the inconsistencies in the sinograms, eliminating most of the xy blurring artefacts in the reconstruction. Unfortunately, the improvement comes at the cost of strong axial blurring. This blurring depends strongly on z , but it is found to be approximately independent of x and y . A z -dependent 1-D axial filter is applied to reduce this axial blurring [13.12]. Multislice rebinning is superior to single slice rebinning, but the noise characteristics are not optimal.

13.2.3.5. *Fourier rebinning*

Fourier rebinning [13.13] is based on the frequency–distance principle, which was explained previously. The Fourier rebinning method is most simply formulated when the projection is written as follows:

$$Y(s, \phi, z, \delta) = \int_{-\infty}^{\infty} dt \Lambda(s \cos \phi - t \sin \phi, s \sin \phi + t \cos \phi, z + t\delta) \quad (13.24)$$

where

$$\delta = \tan \theta;$$

θ is the angle between the LOR and the xy plane;

and the integration variable t is the distance between the position on the LOR and the z axis.

It follows that:

$$Y(s, \phi, z, \delta) = \frac{Y_P(s, \phi, z, \Delta_z = 2\delta\sqrt{R^2 - s^2})}{\sqrt{1 + \delta^2}} \quad (13.25)$$

$$\approx \frac{Y_P(s, \phi, z, \Delta_z = 2\delta R)}{\sqrt{1 + \delta^2}} \quad (13.26)$$

where the approximation is valid whenever $s \ll R$. In this case, no interpolation is needed; it is sufficient to scale the PET data Y_p with the weight factor $\sqrt{1 + \delta^2}$.

Fourier rebinning uses the frequency–distance principle to find the distance d corresponding to a particular portion of the oblique sinogram. As illustrated in Fig. 13.6, distance is used to locate the direct sinogram to which this portion should be assigned. Denoting the 2-D Fourier transform of Y with respect to s and ϕ as \mathcal{Y} , this can be written as:

$$\mathcal{Y}(v_s, v_\phi, z, \delta) \approx \mathcal{Y}(v_s, v_\phi, z - \delta \frac{v_\phi}{v_s}, 0) \quad (13.27)$$

This equation explains how to distribute frequency components from a particular oblique sinogram into different direct sinograms. Frequencies located on the line




FIG. 13.6. Fourier rebinning: the distance from the rotation axis is obtained via the frequency–distance principle. This distance is used to identify the appropriate direct sinogram.

$v_\phi = v_s$ in the oblique sinogram z can be assigned to that same line in the direct sinogram $z + d\delta$.

The final rebinning algorithm (often called ‘FORE’) is obtained by averaging all of the available estimates of the direct sinogram:

$$\begin{aligned}
 \mathcal{Y}(v_s, v_\phi, z, 0) &\approx \frac{1}{\delta_{\max}} \int_0^{\delta_{\max}} d\delta \mathcal{Y}(v_s, v_\phi, z + \delta \frac{v_\phi}{v_s}, \delta) && \text{if } |v_s| \gg 0 \\
 &\approx \mathcal{Y}(0, 0, z, 0) && \text{if } v_s \approx 0, v_\phi \approx 0 \\
 &\approx 0 && \text{if } |v_\phi / v_s| > R_f
 \end{aligned} \tag{13.28}$$

It should be noted that the rebinning expression is only valid for large v_s . In the low frequency range, only the direct sinogram is used. The last line of Eq. (13.28) holds because the image $\Lambda(x, y, z)$ is assumed to be zero outside the FOV $\sqrt{x^2 + y^2} > R_f$.

A more rigorous mathematical derivation of the frequency–distance relation is given in Ref. [13.14]. Alternative derivations based on exact rebinning expressions are given in Ref. [13.13].

After Fourier rebinning, the resulting 2-D dataset can be reconstructed with any 2-D reconstruction algorithm. A popular method is the combination of Fourier rebinning with a 2-D statistical reconstruction algorithm.

13.2.3.6. Exact rebinning methods

Fourier rebinning is an approximate method, but was found to be sufficiently accurate for apertures up to $\theta_0 = 25^\circ$, and it is, therefore, largely sufficient for most current PET systems. However, there is a tendency towards still larger acceptance angles, and a more exact Fourier rebinning algorithm may be needed in the future. An example of an ‘exact’ rebinning algorithm is FOREX [13.13]. It is exact in the sense that the rebinning expression is exact for the continuous 3-D X ray transform.

According to the central section theorem, the 2-D Fourier transform of a projection $Y(s, \phi, z, \delta)$ equals a cross-section through the 3-D Fourier transform of the image $\Lambda(x, y, z)$:

$$\mathcal{Y}_{13}(v_s, \phi, v_z, \delta) = \mathcal{L}(v_s \cos \phi + v_z \delta \sin \phi, v_s \sin \phi - v_z \delta \cos \phi, v_z) \quad (13.29)$$

The subscript of \mathcal{Y}_{13} denotes a Fourier transform with respect to s and z . Defining:

$$\sigma = \arctan(\delta v_z / v_s)$$

$$v'_s = v_s \sqrt{1 + \delta^2 v_z^2 / v_s^2}$$

Equation (13.29) can be rewritten as:

$$\mathcal{Y}_{13}(v_s, \phi, v_z, \delta) = \mathcal{L}(v'_s \cos(\phi - \sigma), v'_s \sin(\phi - \sigma), v_z) \quad (13.30)$$

Taking the 1-D Fourier transform with respect to ϕ yields:

$$\mathcal{Y}_{123}(v_s, v_\phi, v_z, \delta) = e^{-iv_\phi \sigma} \int_0^{2\pi} e^{-iv_\phi \phi} \mathcal{L}(v'_s \cos \phi, v'_s \sin \phi, v_z) d\phi \quad (13.31)$$

By comparing the expressions for $\mathcal{Y}_{123}(v_s, v_\phi, v_z, \delta)$ and $\mathcal{Y}_{123}(v_s, v_\phi, v_z, 0)$, one finally obtains:

$$\mathcal{Y}_{123}(v_s, v_\phi, v_z, \delta) = e^{-iv_\phi \sigma} \mathcal{Y}_{123}(v'_s, v_\phi, v_z, 0) \quad (13.32)$$

A problem of FOREX is that it needs the 1-D Fourier transform along z , which cannot be computed for truncated projections. Similar as with 3-D filtered back

projection, the problem can be avoided by completing the truncated projections with synthetic data. Fortunately, Eq. (13.32) can be used in both ways, and allows estimation of (missing) oblique sinograms from the available direct sinograms. The resulting algorithm is slower than FORE, but still considerably faster than 3-D FBP with reprojection.

13.2.4. Time of flight PET

In time of flight (TOF) PET, the difference in arrival time of the two detected photons is used to estimate the position of their emission along the LOR. The uncertainty in the time estimation results in a similar uncertainty in the position estimation, which can be well modelled as a Gaussian distribution. As a result, the TOF projections correspond to Gaussian convolutions along lines, rather than to line integrals, as illustrated in Fig. 13.7.




FIG. 13.7. Time of flight projection can be well modelled as a 1-D Gaussian convolution in the direction of the line of response.

The corresponding TOF back projection corresponds to convolving the measured data with the same 1-D Gaussians, followed by summation over all angles.

Recall from Eq. (13.2) that the regular projection followed by the regular back projection corresponds to a convolution with a blurring filter:

$$B_{\text{nonTOF}}(x, y) = \frac{1}{\sqrt{x^2 + y^2}} \quad (13.33)$$

The Fourier transform of $1/\sqrt{x^2 + y^2}$ equals $1/\sqrt{v_x^2 + v_y^2}$. Consequently, this blurring can be undone by the ramp filter $\sqrt{v_x^2 + v_y^2}$, which can be applied either before or after back projection (see Section 13.2.1).

If σ_{TOF} is the standard deviation of the TOF-blurring kernel, then TOF projection followed by TOF back projection corresponds to convolution with the blurring kernel:

$$B_{\text{TOF}}(x, y) = \frac{\text{Gauss}(x, y, \sqrt{2}\sigma_{\text{TOF}})}{\sqrt{x^2 + y^2}} \quad (13.34)$$

$$= \frac{1}{\sqrt{x^2 + y^2}} \frac{1}{2\sqrt{\pi}\sigma_{\text{TOF}}} \exp\left(-\frac{x^2 + y^2}{4\sigma_{\text{TOF}}^2}\right) \quad (13.35)$$

It should be noted that the Gaussian in the equation above has a standard deviation of $\sqrt{2}\sigma_{\text{TOF}}$. This is because the Gaussian blurring is present in the projection and in the back-projection. The filter required in TOF PET FBP is derived by inverting the Fourier transform of B_{TOF} , and equals:

$$\text{TOF_recon_filter}(v) = \frac{1}{\exp(-2\pi^2\sigma_{\text{TOF}}^2v^2)I_0(2\pi^2\sigma_{\text{TOF}}^2v^2)} \quad (13.36)$$

where I_0 is the zero order modified Bessel function of the first kind.

This FBP expression is obtained by using the ‘natural’ TOF back projection, defined as the adjoint of the TOF projection. This back projection also appears in LS approaches, and it has been shown that with this back projection definition, FBP is optimal in an (unweighted) LS sense [13.15]. However, TOF PET data are redundant and different back projection definitions could be used; they would yield different expressions for $B_{\text{TOF}}(x, y)$ in Eq. (13.34) and, therefore, different TOF reconstruction filters.

Just as for non-TOF PET, exact and approximate rebinning algorithms for TOF PET have been derived to reduce the data size. As the TOF information limits the back projection to a small region, the errors from approximate rebinning are typically much smaller than in the non-TOF case.

13.3. ITERATIVE RECONSTRUCTION

13.3.1. Introduction

13.3.1.1. Discretization

In analytical reconstruction, it is initially assumed that the unknown object can be represented as a function $\Lambda(\vec{x})$ with $\vec{x} \in \mathbb{R}^3$, and that the acquired data can be represented as a function $Y(s, \theta)$ with $s \in \mathbb{R}^2$ and θ a unit vector in \mathbb{R}^2 or \mathbb{R}^3 . The reconstruction algorithm is then derived by mathematical inversion (assuming some convenient properties for Λ and Y), and finally the resulting algorithm is discretized to make it ready for software implementation. In iterative reconstruction, one usually starts by discretizing the problem. This reduces the reconstruction problem to finding a finite set of unknown values from a finite set of equations, a problem which can be solved with numerical inversion. The advantage of numerical inversion is that only a model for the acquisition process is needed, not for its inverse. That makes it easier (although it may still be non-trivial) to take into account some of the undesired but unavoidable effects that complicate the acquisition, such as photon attenuation, position dependent resolution, gaps between the detectors and patient motion.

After discretization, the unknown image values and the known measured values can be represented as column vectors λ and y . The PET or SPECT acquisition process is characterized by the system matrix A and an additive contribution \bar{b} , and n is the measurement noise:

$$y = A\lambda + \bar{b} + n \quad \text{or} \quad y_i = \sum_{j=1}^J A_{ij}\lambda_j + \bar{b}_i + n_i, \quad i = 1, \dots, I \quad (13.37)$$

The symbol y_i denotes the number of photons measured at LOR i , where the index i runs over all of the sinogram elements (merging the three or four sinogram dimensions into a single index). The index j runs over all of the image voxels, and A_{ij} is the probability that a unit of radioactivity in j gives rise to the detection of a photon (SPECT) or photon pair (PET) in LOR i . The estimate of the additive contribution is denoted as \bar{b} . This estimate is assumed to be noise-free and includes, for example, scatter and randoms in PET or cross-talk between different energy windows in multitracer SPECT studies. Finally, n_i represents the noise contribution in LOR i .

Image reconstruction now consists of finding λ , given A , y and \bar{b} , and a statistical model for n .

For further reading about this subject, the recent review paper on iterative reconstruction by Qi and Leahy [13.16] is an ideal starting point.

13.3.1.2. *Objective functions*

The presence of the noise precludes exact reconstruction. For this reason, the reconstruction is often treated as an optimization task: it is assumed that a useful clinical image can be obtained by maximizing a well chosen objective function. When the statistics of the noise are known, a Bayesian approach can be applied, searching for the image $\hat{\lambda}$ that maximizes the conditional probability on the data:

$$\begin{aligned} \hat{\lambda} &= \arg \max_{\lambda} p(\lambda | \mathbf{y}) \\ &= \arg \max_{\lambda} \frac{p(\mathbf{y} | \lambda)p(\lambda)}{p(\mathbf{y})} \end{aligned} \tag{13.38}$$

$$\begin{aligned} &= \arg \max_{\lambda} p(\mathbf{y} | \lambda)p(\lambda) \\ &= \arg \max_{\lambda} (\ln p(\mathbf{y} | \lambda) + \ln p(\lambda)) \end{aligned} \tag{13.39}$$

The second equation is Bayes' rule. The third equation holds because \mathbf{y} does not depend on λ , and the fourth equation is valid because computing the logarithm does not change the position of the maximum. The probability $p(\mathbf{y} | \lambda)$ gives the likelihood of measuring a particular sinogram \mathbf{y} , when the tracer distribution equals λ . This distribution is often simply called the likelihood. The probability $p(\lambda)$ represents the a priori knowledge about the tracer distribution, available before PET or SPECT acquisition. This probability is often called the prior distribution. The knowledge available after the measurements equals $p(\mathbf{y} | \lambda)p(\lambda)$ and is called the posterior distribution. To keep things simple, it is often assumed that no prior information is available, i.e. $p(\lambda | \mathbf{y}) \propto p(\mathbf{y} | \lambda)$. Finding the solution then reduces to maximizing the likelihood $p(\mathbf{y} | \lambda)$ (or its logarithm). In this section, maximum-likelihood algorithms are discussed. Maximum a posteriori (MAP) algorithms are discussed in Section 13.3.5, as a strategy to suppress noise propagation.

A popular approach to solve equations of the form of Eq. (13.37) is LS estimation. This is equivalent to a maximum-likelihood approach, if it is assumed that the noise is Gaussian with a zero mean and a fixed, position independent

standard deviation σ . The probability to measure the noisy value y_i when the expected value was $\sum_j A_{ij}\lambda_j + \bar{b}_i$ then equals:

$$p_{\text{LS}}(y_i | \sum_j A_{ij}\lambda_j + \bar{b}_i) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(y_i - (\sum_j A_{ij}\lambda_j + \bar{b}_i))^2}{2\sigma^2}\right] \quad (13.40)$$

As the noise in the sinogram is not correlated, the likelihood (i.e. the probability of measuring the entire noisy sinogram \mathbf{y}) equals:

$$p_{\text{LS}}(\mathbf{y} | \boldsymbol{\lambda}) = p_{\text{LS}}(\mathbf{y} | \mathbf{A}\boldsymbol{\lambda} + \bar{\mathbf{b}}) = \prod_i p_{\text{LS}}(y_i | \sum_j A_{ij}\lambda_j + \bar{b}_i) \quad (13.41)$$

It is more convenient to maximize the logarithm of p_{LS} ; dropping constants, the objective function L_{LS} is finally obtained:

$$L_{\text{LS}} = -\sum_i (y_i - (\sum_j A_{ij}\lambda_j + \bar{b}_i))^2 = -(\mathbf{y} - (\mathbf{A}\boldsymbol{\lambda} + \bar{\mathbf{b}}))'(\mathbf{y} - (\mathbf{A}\boldsymbol{\lambda} + \bar{\mathbf{b}})) \quad (13.42)$$

where the prime denotes matrix transpose. Setting the first derivatives with respect to λ_j to zero for all j gives:

$$\begin{aligned} \mathbf{A}'(\mathbf{y} - \mathbf{A}\boldsymbol{\lambda} - \bar{\mathbf{b}}) &= 0 \\ \boldsymbol{\lambda}(\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'(\mathbf{y} - \bar{\mathbf{b}}) \end{aligned} \quad (13.43)$$

provided that $\mathbf{A}'\mathbf{A}$ is non-singular. The operator \mathbf{A} is the discrete projection; its transpose \mathbf{A}' is the discrete back projection. Its analytical counterpart was given in Eq. (13.2) and illustrated in Fig. 13.1. The same figure shows that the operator $\mathbf{A}'\mathbf{A}$ behaves as a blurring filter.




FIG. 13.8. The image of point sources is projected and back projected again along ideal parallel beams. This yields a shift-invariant blurring.

Figure 13.8 is similar, but illustrates A and $A'A$ on an image of three point sources, using ideal parallel-beam projection. The figure shows the resulting point spread functions of $A'A$ for each of the point sources. They are identical: for ideal parallel-beam projection, $A'A$ is shift-invariant, equivalent to a convolution. It follows that $(A'A)^{-1}$ is the corresponding shift-invariant deconvolution, which is easily computed via the Fourier transform. In this situation, LS reconstruction (Eq. (13.43)) is the discrete equivalent of the ‘back project-then-filter’ algorithm (Eq. (13.15)), applied to the data after pre-correction for \bar{b} .

Figure 13.9 illustrates A and $A'A$ for a projector that models the position dependent blurring of a typical parallel-beam SPECT collimator. The blurring induced by $A'A$ is now shift-variant — it cannot be modelled as a convolution and its inverse cannot be computed with the Fourier transform. For real life problems, direct inversion of $A'A$ is not feasible. Instead, iterative optimization is applied to find the maximum of Eq. (13.42).




FIG. 13.9. The image of point sources is projected and back projected again with collimator blurring. This yields a shift-variant blurring.

It is known that the number of detected photons is subject to Poisson noise, not to uniform Gaussian noise. The Poisson distribution can be well approximated with a Gaussian distribution, where the variance of the Gaussian equals its mean. With this approximation, σ must be replaced by σ_i in Eq. (13.40) because now there is a different Gaussian distribution for every sinogram pixel i . Proceeding as before, the weighted least squares (WLS) objective function is:

$$\begin{aligned}
 L_{\text{WLS}} &= -\sum_i \frac{(y_i - (\sum_j A_{ij}\lambda_j + \bar{b}_i))^2}{\sigma_i^2} \\
 &= -(\mathbf{y} - (A\boldsymbol{\lambda} + \bar{\mathbf{b}}))' \mathbf{C}_y^{-1} (\mathbf{y} - (A\boldsymbol{\lambda} + \bar{\mathbf{b}}))
 \end{aligned}
 \tag{13.44}$$

where \mathbf{C}_y is the covariance matrix of the data.

For emission tomography, it is a diagonal matrix (all covariances are zero) with elements $C_y[i, i] = \sigma_i^2$. The corresponding WLS reconstruction can be written as:

$$\lambda = (A'C_y^{-1}A)^{-1}A'C_y^{-1}(y - \bar{y}) \tag{13.45}$$

The operator $A'C_y^{-1}A$ is always shift-variant, even for ideal parallel-beam tomography. This is illustrated in Fig. 13.10. The noise-free sinogram \bar{y} is computed for a particular activity distribution. Setting $C_y = \text{diag}(\bar{y})$, the operator $A'C_y^{-1}A$ can be analysed by applying it to the image of a few point sources, called x in the figure. The image x is projected, the sinogram Ax is divided by \bar{y} on a pixel basis and the result is back projected. Clearly, position dependent blurring is obtained. Consequently, iterative optimization must be used for WLS reconstruction.




FIG. 13.10. The operator $A'C_y^{-1}A$ is derived for a particular activity distribution (top left) and then applied to a few point sources x . Although ideal parallel-beam projection was used, shift-variant blurring is obtained.

In practice, because there is only a noisy sinogram y , the noise-free sinogram \bar{y} must be estimated to find C_y . There are basically two approaches. In the first approach, \bar{y} is estimated from y , e.g. by smoothing y to suppress the noise. In the second approach, \bar{y} is estimated as $A\lambda^{(k)} + \bar{b}$ during the iterative optimization, where $\lambda^{(k)}$ is the estimate of the reconstruction available at iteration k . A drawback of the first approach is that the noise on the data affects the weights, with a tendency to give higher weight when the noise contribution

happens to be negative. A complication of the second approach is that it makes σ_i a function of λ . In this case, the normalizing amplitude $1/(\sqrt{2\pi}\sigma_i)$ of the Gaussians cannot be dropped, implying that an additional term $-\sum_i \ln \sigma_i$ should be added to Eq. (13.44).

It is possible to use the Poisson distribution itself, instead of approximating it with Gaussians. The probability of the noise realization y_i then becomes:

$$p_{\text{ML}}(y_i | \sum_j A_{ij}\lambda_j + \bar{b}_i) = \frac{e^{-\sum_j A_{ij}\lambda_j + \bar{b}_i} (\sum_j A_{ij}\lambda_j + \bar{b}_i)^{y_i}}{y_i!} \quad (13.46)$$

Proceeding as before, the log-likelihood function is:

$$\ln \left(\prod_i p_{\text{LS}}(y_i | \sum_j A_{ij}\lambda_j + \bar{b}_i) \right) = \sum_i y_i \ln(\sum_j A_{ij}\lambda_j + \bar{b}_i) - (\sum_j A_{ij}\lambda_j + \bar{b}_i) - \ln y_i!$$

$$L_{\text{ML}} = \sum_i y_i \ln(\sum_j A_{ij}\lambda_j + \bar{b}_i) - (\sum_j A_{ij}\lambda_j + \bar{b}_i) \quad (13.47)$$

It should be noted that the term $\ln y_i!$ can be dropped, because it is not a function of λ . As L_{ML} is a non-linear function of λ , the solution cannot be written as a product of matrixes. However, it is sometimes helpful to know that the features of the Poisson-objective function are often very similar to those of the WLS function (Eq. (13.44)).

13.3.2. Optimization algorithms

Many iterative reconstruction algorithms have been proposed to optimize the objective functions L_{WLS} and L_{ML} . Here, only two approaches are briefly described: preconditioned conjugate gradient methods and optimization transfer, with expectation maximization (EM) as a special case of the latter.

13.3.2.1. Preconditioned gradient methods

The objective function will be optimized when its first derivatives are zero:

$$\hat{y}_i = \sum_j A_{ij}\lambda_j + \bar{b}_i \quad (13.48)$$

$$\frac{\partial L_{\text{WLS}}(\boldsymbol{\lambda})}{\partial \lambda_j} = \sum_i A_{ij} \frac{y_i - \hat{y}_i}{\sigma_i^2} \quad (13.49)$$

$$\frac{\partial L_{\text{ML}}(\boldsymbol{\lambda})}{\partial \lambda_j} = \sum_i A_{ij} \frac{y_i - \hat{y}_i}{\hat{y}_i} \quad (13.50)$$

The optimization can be carried out by a steepest ascent method, which can be formulated as follows:

$$\mathbf{d}^k = \nabla L(\boldsymbol{\lambda}^{k-1})$$

$$\boldsymbol{\lambda}^k = \boldsymbol{\lambda}^{k-1} + \alpha_k \mathbf{d}^k \quad (13.51)$$

$$\alpha_k = \arg \max_{\alpha} L(\boldsymbol{\lambda}^{k-1} + \alpha \mathbf{d}^k)$$

where the superscripts k and $k-1$ denote the iteration numbers and ∇L is the vector of the first derivatives of L with respect to λ_j .

Steepest gradient ascent is known to be suboptimal, requiring many iterations for reasonable convergence. To find a better update, it is required that after the update, the first derivatives of L are zero as intended. Approximating this with a first order Taylor expansion yields:

$$\nabla L(\boldsymbol{\lambda}^{k-1} + \mathbf{p}^k) = 0$$

$$\nabla L(\boldsymbol{\lambda}^{k-1}) + \mathbf{H} \mathbf{p}^k \approx 0 \quad (13.52)$$

$$\mathbf{p}^k \approx -\mathbf{H}^{-1} \nabla L(\boldsymbol{\lambda}^{k-1}) = -\mathbf{H}^{-1} \mathbf{d}^k$$

where the Hessian \mathbf{H} is the matrix of second derivatives of L . This is obviously a very large matrix, but its elements are relatively easy to compute:

$$\text{for WLS: } H_{jk} = -\sum_i \frac{A_{ij} A_{ik}}{\sigma_i^2} = -(\mathbf{A}' \mathbf{C}_y^{-1} \mathbf{A})[j, k] \quad (13.53)$$

$$\text{for ML: } H_{jk} = -\sum_i \frac{A_{ij} A_{ik} y_i}{\hat{y}_i^2} \approx \sum_i \frac{A_{ij} A_{ik}}{\hat{y}_i} \quad (13.54)$$

$$\approx -(\mathbf{A}' \mathbf{C}_y^{-1} \mathbf{A})[j, k] \quad \text{if } \hat{y} \approx \bar{y} \quad (13.55)$$

For a Gaussian likelihood, Eq. (13.52) is in fact exact, and a single iteration would suffice. As shown before, however, it is usually impossible to compute \mathbf{H}^{-1} . Instead, approximations to the Hessian (or other heuristics) can be used to obtain a good \mathbf{M} to derive a so-called preconditioned gradient ascent algorithm:

$$\mathbf{d}^k = \nabla L(\boldsymbol{\lambda}^{k-1}) \tag{13.56}$$

$$\boldsymbol{\lambda}^k = \boldsymbol{\lambda}^{k-1} + \alpha_k \mathbf{M} \mathbf{d}^k$$

To ensure that the convergence is preserved, the matrix \mathbf{M} must be symmetric positive definite (it should be noted that $-\mathbf{H}^{-1}$ is symmetric positive definite, since \mathbf{H} is symmetric negative definite, if \mathbf{A} has maximum rank).

A simple way to obtain a reasonable \mathbf{M} is to use only the diagonal elements of \mathbf{H} : $M_{ii} = -1/H_{ii}$ and $M_{ij} = 0$ if $i \neq j$. A more sophisticated approach is discussed in Ref. [13.17]: a circulant, i.e. shift-invariant approximation of the Hessian is proposed. Such an approximation is easily computed by fixing j at a particular location in the image in Eqs (13.53) or (13.54), which yields an image that can be considered as the point spread function of a convolution operator. This shift-invariant operator is then inverted via the Fourier transform, yielding a non-diagonal matrix \mathbf{M} . For cases where the true Hessian depends heavily on position, the approach could be repeated for a few well chosen positions j , applying linear interpolation for all other positions.

13.3.2.2. Conjugate gradient methods

Figure 13.11 shows the convergence of the steepest gradient ascent algorithm for a nearly quadratic function of two variables. In every iteration, the algorithm starts moving in the direction of the maximum gradient (i.e. perpendicular to the isocontour), and keeps moving along the same line until a maximum is reached (i.e. until the line is a tangent to the isocontour). This often leads to a zigzag line, requiring many iterations for good convergence.

The conjugate gradient algorithm is designed to avoid these oscillations [13.18]. The first iteration is identical to that of the steepest gradient ascent. However, in the following iterations, the algorithm attempts to move in a direction for which the gradient along the previous direction(s) remains the same (i.e. equal to zero). The idea is to eliminate the need for a new optimization along these previous directions. Let \mathbf{d}_{old} be the previous direction and \mathbf{H} the Hessian matrix (i.e. the second derivatives). It is now required that the new direction \mathbf{d}_{new} be such that the gradient along \mathbf{d}_{old} does not change. When moving in direction \mathbf{d}_{new} , the




FIG. 13.11. The dotted lines are isocontours of the objective function. The solid line shows the convergence of the steepest gradient ascent algorithm, the dashed line the convergence of conjugate gradient ascent. It should be noted that the starting points are equivalent because of the symmetry. The objective function equals $-(a|x-x_0|^p + b|y-y_0|^p)$, with $p = 2.15$.

gradient will change (using a quadratic approximation) as $\mathbf{H}\mathbf{d}_{\text{new}}$. Requiring that the resulting change along \mathbf{d}_{old} is zero yields the condition:

$$\mathbf{d}'_{\text{old}} \mathbf{H} \mathbf{d}_{\text{new}} = 0 \tag{13.57}$$

This behaviour is illustrated by the dashed line in Fig. 13.11: in the second iteration, the algorithm moves in a direction such that the trajectory cuts the isocontours at the same angle as in the starting point. For a quadratic function in n dimensions, convergence is obtained after no more than n iterations. As the function in Fig. 13.11 is not quadratic, more than two iterations are required for full convergence.

The new direction can be easily computed from the previous ones, without computation of the Hessian \mathbf{H} . The Polak–Ribiere algorithm is given by [13.18]:

$$\begin{aligned} \mathbf{g}_{\text{new}} &= \nabla L(\boldsymbol{\lambda}_{\text{old}}) \\ \gamma &= \frac{(\mathbf{g}_{\text{new}} - \mathbf{g}_{\text{old}})' \mathbf{g}_{\text{new}}}{\mathbf{g}'_{\text{old}} \mathbf{g}_{\text{old}}} \end{aligned} \tag{13.58}$$

$$\mathbf{d}_{\text{new}} = \mathbf{g}_{\text{new}} + \gamma \mathbf{d}_{\text{old}}$$

$$\alpha = \arg \max_{\alpha} L(\boldsymbol{\lambda}_{\text{old}} + \alpha \mathbf{d}_{\text{new}})$$

$$\boldsymbol{\lambda}_{\text{new}} = \boldsymbol{\lambda}_{\text{old}} + \alpha \mathbf{d}_{\text{new}}$$

This algorithm requires storage of the previous gradient \mathbf{g}_{old} and the previous search direction \mathbf{d}_{old} . In each iteration, it computes the new gradient and search direction, and applies a line search along the new direction.

13.3.2.3. Preconditioned conjugate gradient methods

Both techniques mentioned above can be combined to obtain a fast reconstruction algorithm, as described in Ref. [13.17]. The preconditioned conjugate gradient ascent algorithm (with preconditioning matrix \mathbf{M}) can be written as follows:

$$\begin{aligned} \mathbf{g}_{\text{new}} &= \nabla L(\boldsymbol{\lambda}_{\text{old}}) \\ \mathbf{p}_{\text{new}} &= \mathbf{M}\mathbf{g}_{\text{new}} \\ \gamma &= \frac{(\mathbf{g}_{\text{new}} - \mathbf{g}_{\text{old}})' \mathbf{p}_{\text{new}}}{\mathbf{g}'_{\text{old}} \mathbf{p}_{\text{old}}} \end{aligned} \tag{13.59}$$

$$\mathbf{d}_{\text{new}} = \mathbf{p}_{\text{new}} + \gamma \mathbf{d}_{\text{old}}$$

$$\alpha = \arg \max_{\alpha} L(\boldsymbol{\lambda}_{\text{old}} + \alpha \mathbf{d}_{\text{new}})$$

$$\boldsymbol{\lambda}_{\text{new}} = \boldsymbol{\lambda}_{\text{old}} + \alpha \mathbf{d}_{\text{new}}$$

13.3.2.4. Optimization transfer

The log-likelihood function (Eq. (13.47)) can be maximized by setting its gradients (Eq. (13.50)) to zero for all $j = 1 \dots J$. A problem is that each of these derivatives is a function of many voxels of $\boldsymbol{\lambda}$, which makes the set of equations very hard to solve. The idea of ‘optimization transfer’ is to replace the problematic log-likelihood function with another function $\Phi(\boldsymbol{\lambda})$ that leads to a simpler set of equations, usually one where the derivative with respect to λ_j is only a function of λ_j and not of the other voxels of $\boldsymbol{\lambda}$. That makes the problem separable into J 1-D optimizations, which are easily solved. Ideally, Φ and L should have the same optimum, but that is asking for too much. The key is to design $\Phi(\boldsymbol{\lambda})$ in such a way that maximization of $\Phi(\boldsymbol{\lambda})$ is guaranteed to increase $L(\boldsymbol{\lambda})$. This leads to an iterative algorithm, since new functions Φ will have to be designed and

maximized repeatedly to maximize L . At iteration k , the surrogate function $\Phi(\lambda)$ needs to satisfy the following conditions (illustrated in Fig. 13.12):

$$\Phi(\lambda^{(k)}) = L(\lambda^{(k)}) \quad (13.60)$$

$$\Phi(\lambda) \leq L(\lambda) \quad (13.61)$$




FIG. 13.12. Optimization transfer: a surrogate function is designed, which is equal to the likelihood in the current reconstruction, and less or equal everywhere else.

It follows that the new reconstruction image $\lambda^{(k+1)}$ which maximizes $\Phi(\lambda)$ has a higher likelihood than $\lambda^{(k)}$:

$$L(\lambda^{(k)}) = \Phi(\lambda^{(k)}) \leq \Phi(\lambda^{(k+1)}) \leq L(\lambda^{(k+1)}) \quad (13.62)$$

Several algorithms for maximum-likelihood and MAP reconstruction in emission and transmission tomography have been developed with this approach. De Pierro [13.19] has shown how the well known maximum-likelihood expectation-maximization (MLEM) algorithm can be derived using the optimization transfer principle. He also showed how this alternative derivation provides a natural way to extend it to an MAP algorithm.

13.3.3. Maximum-likelihood expectation-maximization

13.3.3.1. Reconstruction from sinogram data

There are many ways to derive the MLEM algorithm, including the original statistical derivation by Shepp and Vardi [13.20] (based on the work by Dempster et al. [13.21]) and the optimization transfer approach by De Pierro [13.19]. Only the EM recipe is given below.

Recall that we wish to find the image λ that maximizes the likelihood function L_{ML} of Eq. (13.47). The EM does this in a remarkable way. Instead of concentrating on L_{ML} , an alternative (different) likelihood function is derived by introducing a set of so-called ‘complete data’ x_{ij} , defined as the number of photons that were emitted at voxel j and detected in LOR i during the measurement. These unobserved data are ‘complete’ in the sense that they describe in more detail than the observed data y_i what happened during the measurement. These variables x_{ij} are Poisson distributed. Just as for the actual data y_i , one can write the log-likelihood function for observing the data x_{ij} while $\bar{x}_{ij} = A_{ij}\lambda_j$ were expected:

$$L_x(\lambda) = \sum_i \sum_j x_{ij} \ln(A_{ij}\lambda_j) - A_{ij}\lambda_j \quad (13.63)$$

However, this likelihood cannot be computed, because the data x_{ij} are not available. The emission measurement only produces sums of the complete data, since:

$$y_i = \sum_j A_{ij}x_{ij} + b_i \quad (13.64)$$

where b_i represents the actual (also unobserved) additive contribution b_i in LOR i .

The EM recipe prescribes computing the expectation of L_x , based on the available data and on the current reconstruction $\lambda^{(k)}$. Based on the reconstruction alone, one would write $E(x_{ij} | \lambda^{(k)}) = A_{ij}\lambda_j^{(k)}$. However, it is also known that x_{ij} should satisfy Eq. (13.64). It can be shown that this leads to the following estimate:

$$E(x_{ij} | \lambda^{(k)}, \mathbf{y}) = \frac{y_i}{\sum_j A_{ij}\lambda_j^{(k)} + \bar{b}_i} A_{ij}\lambda_j^{(k)} \quad (13.65)$$

where \bar{b}_i is the noise-free estimate of b_i , which is assumed to be available.

Inserting this in Eq. (13.63) produces the expectation of $L_x(\lambda)$ and completes the expectation (E) step. For the maximization (M) step, the first derivatives are simply set to zero:

$$\frac{\partial L_x(\lambda)}{\partial \lambda_j} = \sum_i \left(\frac{y_i}{\sum_j A_{ij} \lambda_j^{(k)} + \bar{b}_i} A_{ij} \lambda_j^{(k)} \frac{1}{\lambda_j} - A_{ij} \right) = 0 \quad (13.66)$$

This is easily solved for λ_j , yielding the new reconstruction $\lambda_j^{(k+1)}$:

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)}}{\sum_i A_{ij}} \sum_i A_{ij} \frac{y_i}{\sum_j A_{ij} \lambda_j^{(k)} + \bar{b}_i} \quad (13.67)$$

This is the well known MLEM algorithm for emission tomography.

It can be shown that this recipe has the wonderful feature that each new EM iteration increases the value of the likelihood L_{ML} . It should be noted that the complete data x_{ij} do not appear in Eq. (13.67); they are needed in the derivation but they do not need to be computed explicitly. This is very fortunate as there is a huge number of them.

An initial image $\lambda^{(1)}$ is required to start the iterations. As experience (and theoretical analysis) has shown that higher spatial frequencies have slower convergence, and because smooth images are preferred, the initial image is usually chosen to be uniform, by setting $\lambda_j^{(1)} = C$ and $j = 1 \dots J$, where C is a strictly positive constant.

The MLEM algorithm is multiplicative, implying that it cannot change the value of a reconstruction voxel, when the current value is zero. For this reason, the voxels in the initial image should only be set to zero if it is known a priori that they are indeed zero. The derivation of the MLEM algorithm uses the assumption that all y_i , all x_{ij} and all λ_j are non-negative. Assuming that $y_i \geq 0$ and $i = 1 \dots I$, and considering that the probabilities A_{ij} are also non-negative, it is clear that when the initial image $\lambda^{(1)}$ is non-negative, all subsequent images $\lambda^{(k)}$ will be non-negative as well. However, when, for some reason, a reconstruction value becomes negative (e.g. because one or a few sinogram values y_i are negative), then convergence is no longer guaranteed. In practice, divergence is almost guaranteed in that case. Consequently, if the sinogram is pre-processed with a procedure that may produce negatives (e.g. randoms subtraction in PET), MLEM reconstruction will only work if all negative values are set to a non-negative value.

13.3.3.2. Reconstruction from list-mode data

The measured data y_i considered in the derivations above (so-called ‘binned’ data) represent the number of counts acquired within an individual crystal pair i (LOR i), that is, y_i represents the sum of those acquired events (indexed by m) that were assigned (histogrammed) to the i -th LOR: $y_i = \sum_{m \in i} 1$. However, in modern PET systems, the number of possible LORs within the FOV typically exceeds (often by many times) the number of events acquired in a clinical PET study. Consequently, the binned data are very sparse and it is more efficient to store and process each acquired event (with all of its relevant information) separately, in the so-called ‘list-mode’ format.

Modification of the maximum-likelihood algorithms is straightforward (whether MLEM or accelerated algorithms based on ordered subsets discussed later), as shown in works by Parra and Barrett [13.22], and by Reader et al. [13.23]. It should be noted that the same is not true about other algorithms, for example, algorithms with additive updates. The MLEM algorithm for the list-mode data can be obtained by replacing y_i in the MLEM equation (Eq. (13.67)) by the above mentioned sum over events, skipping the LORs with zero counts (which do not contribute to the MLEM sum), and combining the sum over LORs i with the sum over events m :

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)}}{\sum_{i \in \text{LORs}} A_{ij}} \sum_{m \in \text{event-list}} A_{i_m j} \frac{1}{\sum_j A_{i_m j} \lambda_j^{(k)} + \bar{b}_{i_m}} \tag{13.68}$$

where i_m represents the LOR index in which the m -th event has been recorded.

The main difference is that the MLEM sum is now evaluated (including calculations of the relevant forward and back projections) only over the list of the available events (in any order). However, it is important to mention here that the normalizing term in front of the sum (sensitivity matrix $\sum_i A_{ij}$) still has to be calculated over all possible LORs, and not only those with non-zero counts. This represents a challenge for the attenuated data (attenuation considered as part of the system matrix A), since the sensitivity matrix has to be calculated specifically for each object and, therefore, it cannot be pre-computed. For modern systems with a large number of LORs, calculation of it often takes more time than the list-mode reconstruction itself. For this reason, alternative approaches (involving certain approximations) have been considered for the calculation of the sensitivity matrix, such as subsampling approaches [13.24] or Fourier based approaches [13.25].

13.3.3.3. Reconstruction of time of flight PET data

In the TOF case, the probability of a pair of photons arriving from a particular point along the LOR (as reported based on the difference of their detection times) is given by a Gaussian kernel having a width determined by the timing uncertainty of the detection system. In contrast, in the non-TOF case, the probability of detecting the event is approximately uniform along the LOR. Modification of iterative reconstruction algorithms (whether for binned or list-mode data) to account for the TOF is straightforward. Integrations along the LORs (the main component of the system matrix A) just need to be replaced with the TOF kernel weighted integrations along the LORs. The forward projection (or back projection) in a certain direction can now be viewed, and performed, as a convolution of the image with a proper TOF kernel in the LOR direction (see Fig. 13.13). The rest of the algorithm, i.e. formulas derived in the previous subsections, stays exactly the same (only the form of the system matrix A is changed). Additional information provided by the TOF measurements, leading to more localized data, results in faster, and more uniform, convergence, as well as in improved signal to noise ratios in reconstructed images, as widely reported in the literature.




FIG. 13.13. Comparison of the data formats for binned time of flight (TOF) data (left: histo-projection for a 45° view) and for the DIRECT (direct image reconstruction for TOF) approach (right: histo-image for a 45° view). Histo-projections can be viewed as an extension of individual non-TOF projections into TOF directions (time bins), and their sampling intervals relate to the projection geometry and timing resolution. Histo-images are defined by the geometry and desired sampling of the reconstructed image. Acquired events and correction factors are directly placed into the image resolution elements of individual histo-images (one histo-image per view) having a one to one correspondence with the reconstructed image voxels.

The TOF mode of operation has some practical consequences (and novel possibilities) for the ways the acquired data are stored and processed. The

list-mode format is very similar to the non-TOF case. The event structure is just slightly expanded by a few bits (5–8 bits/event) to include the TOF information, and the events are processed event by event as in the non-TOF case.

On the other hand, the binned data undergo considerable expansion when accommodating the TOF information. The projection (X ray transform) structures are expanded by one dimension, that is, each projection bin is expanded in the LOR direction into the set of time bins forming the so-called histo-projections (see Fig. 13.13 (left)). In practice, the effect of this expansion on the data size is not as bad as it appears, because the localized nature of TOF data allows decreased angular sampling (typically about 5–10 times) in both azimuthal and co-polar directions (views), while still satisfying angular sampling requirements. The resulting data size, thus, remains fairly comparable to the non-TOF case. During the reconstruction process, the histo-projection data are processed time-bin by time-bin (instead of projection line by line in the non-TOF case). It should be noted that hybrid approaches also exist between the two aforementioned approaches, in which the data are binned in the LOR space, but events are stored in list-mode for each LOR bin.

TOF also allows a conceptually different approach of data partitioning, leading to more efficient reconstruction implementations, by using the DIRECT (direct image reconstruction for TOF) approach utilizing so-called histo-images (see Fig. 13.13 (right)) [13.25]. In the DIRECT approach, the data are directly histogrammed (deposited), for each view, into image resolution elements (voxels) of desired size. Similarly, all correction arrays and data are estimated or calculated in the same histo-image format. The fact that all data and image structures are now in image arrays (of the same geometry and size) makes possible very efficient computer implementations of the data processing and reconstruction operations.

13.3.3.4. Reconstruction of dynamic data

Data acquired from an object dynamically changing with time in activity distribution, or in morphology (shape), or in both is referred to as dynamic data. An example of the first case would be a study looking at temporal changes in activity uptake in individual organs or tissues, so-called time–activity curves. An example of the second case would be a gated cardiac study providing information about changes of the heart morphology during the heart beat cycle (such as changes of the heart wall thickness or movements of the heart structures).

The dynamic data can be viewed as an expansion of static (3-D) data by the temporal information into 4-D (or 5-D) data. The dynamic data are usually subdivided (spread) into a set of temporal (time) frames. In the first application, each time frame represents data acquired within a certain

sequential time subinterval of the total acquisition time. The subintervals can be uniform, or non-uniform with their durations adjusted, for example, to the speed of the change of the activity curves. In the second application, each time frame represents the total counts acquired within a certain stage (gate) of the periodic organ movement (e.g. gated based on the electrocardiogram signal). In the following, issues of the reconstruction of dynamic data in general are addressed. Problems related to the motion and its corrections are discussed in Section 13.3.6.4.

Once the data are subdivided (during acquisition) or sorted (acquired list-mode data) into the set of time frames, seemingly the most natural way of reconstructing them is to do it for each time frame separately. It should be noted that this is the only available option for the analytical reconstruction approaches, while the iterative reconstruction techniques can also reconstruct the dynamic data directly in 4-D (or 5-D). A problem with frame by frame reconstruction is that data in the individual time frames are quite noisy, since each time frame only has a fraction of the total acquired counts, leading to noisy reconstructions. Consequently, the resulting reconstructions often have to be filtered in the spatial and/or temporal directions to obtain images of any practical value. Temporal filtering takes into account time correlations between the signal components in the neighbouring time frames, while the noise is considered to be independent. Filtering, however, leads to resolution versus noise trade-offs.

On the other hand, reconstructing the whole 4-D (or 5-D) dataset together, while using this correlation information in the (4-D) reconstruction process via proper temporal (resolution) kernels or basis functions, can considerably improve those trade-offs as reported in the literature (similarly to the case of spatial resolution modelling). The temporal kernels (basis functions) can be uniform in shape and distribution, or can have a non-uniform shape (e.g. taking into account the expected or actual shape of the time–activity curves) and can be distributed on a non-uniform grid (e.g. reflecting count levels at individual frames or image locations). The kernel shapes and distributions can be defined, or determined, beforehand and be fixed during the reconstruction. During the reconstruction process, just the amplitudes of the basis functions are reconstructed. The algorithms derived in the previous subsections basically stay the same, where the temporal kernels can be considered as part of the system matrix A (comparable to including the TOF kernel in TOF PET). Another approach, more accurate but mathematically and computationally much more involved, is to iteratively build up the shape (and distribution) of the temporal kernels during the reconstruction in conjunction with the reconstruction of the emission activity (that is, the amplitude of the basis functions).

While iterative methods lead to a clear quality improvement when reconstructing dynamic data, thanks to the more accurate models of the signal and

data noise components, for the quantitative dynamic studies their shortcoming is their non-linear behaviour, especially if they are not fully converged. For example, the local bias levels can vary across the time frames as the counts, local activity levels and object morphology change, which can lead to less accurate time–activity curves. On the other hand, analytical techniques which are linear and consequently do not depend on the count levels and local activity, might provide a more consistent (accurate) behaviour across the time frames in the mean (less bias of the mean), but much less consistent (less precise) behaviour in the variance due to the largely increased noise. It is still an open issue which of the two approaches provides more clinically useful results, and the discussions and research on this topic are still open and ongoing.

13.3.4. Acceleration

13.3.4.1. Ordered-subsets expectation-maximization

The MLEM algorithm requires a projection and a back projection in every iteration, which are operations involving a large number of computations. Typically, MLEM needs several tens to hundreds of iterations for good convergence. Consequently, MLEM reconstruction is slow and many researchers have studied methods to accelerate convergence.

The method most widely used is ordered-subsets expectation-maximization (OSEM) [13.26]. The MLEM algorithm (Eq. (13.67)) is rewritten here for convenience:

$$\hat{y}_i^{(k)} = \sum_j A_{ij} \lambda_j^{(k)} + \bar{b}_i \quad (13.69)$$

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)}}{\sum_i A_{ij}} \sum_i A_{ij} \frac{y_i}{\hat{y}_i^{(k)}} \quad (13.70)$$

where k is the iteration number and $\lambda^{(1)}$ is typically set to a uniform, strictly positive image.

In OSEM, the set of all projections $\{1 \dots I\}$ is divided into a series of subsets S_t , $t = 1 \dots T$. Usually, these subsets are exhaustive and non-overlapping, i.e. every projection element i belongs to exactly one subset S_t . In SPECT and PET, the data \mathbf{y} are usually organized as a set of (parallel- or fan-beam) projections, indexed by projection angle ϕ . Therefore, the easiest way to produce subsets of \mathbf{y} is by assigning all of the data for each projection angle to exactly one of the subsets.

However, if the data \mathbf{y} are stored in list-mode (see Section 13.3.2), the easiest way is to simply cut the list into blocks, assigning each block to a different subset.

The OSEM algorithm can then be written as:

$$\begin{aligned}
 &\text{initialize } \lambda_j^{\text{old}}, j = 1, \dots, J \\
 &\text{for } k = 1, \dots, K \\
 &\quad \text{for } t = 1, \dots, T \\
 &\quad\quad \hat{y}_i = \sum_j A_{ij} \lambda_j^{\text{old}} + \bar{b}_i, \quad i \in \mathcal{S}_t \\
 &\quad\quad \text{for } j = 1, \dots, J \\
 &\quad\quad\quad \lambda_j^{\text{new}} = \frac{\lambda_j^{\text{old}}}{\sum_{i \in \mathcal{S}_t} A_{ij}} \sum_{i \in \mathcal{S}_t} A_{ij} \frac{y_i}{\hat{y}_i}
 \end{aligned} \tag{13.71}$$

If all of the projections are combined into a single subset, the OSEM algorithm is identical to the MLEM algorithm. Otherwise, a single OSEM iteration k consists of T sub-iterations, where each sub-iteration is similar to an MLEM iteration, except that the projection and back projection are only done for the projections of the subset \mathcal{S}_t . If every sinogram pixel i is in exactly one subset, the computational burden of a single OSEM iteration is similar to that of an MLEM iteration. However, MLEM would update the image only once, while OSEM updates it T times. Experience shows that this improves convergence by a factor of about T , which is very significant.

Convergence is only guaranteed for consistent data and provided that there is subset balance, which requires:

$$\sum_{i \in \mathcal{S}_t} A_{ij} = \sum_{i \in \mathcal{S}_u} A_{ij} \tag{13.72}$$

where \mathcal{S}_t and \mathcal{S}_u are different subsets.

In practice, these conditions are never satisfied, and OSEM can be shown to converge to a limit cycle rather than to a unique solution, with the result that the OSEM reconstruction is noisier than the corresponding MLEM reconstruction. However, in many applications, the difference between the two is not clinically relevant.

The procedure is illustrated with a simple simulation in Fig. 13.14. As there was no noise and no attenuation, convergence of OSEM is guaranteed in this example. In more realistic cases, it may be recommended to have four or more

projections in a single subset, to prevent excessive noise amplification at higher iteration numbers.




FIG. 13.14. A simulation comparing a single ordered-subsets expectation-maximization (OSEM) iteration with 40 subsets, to 40 maximum-likelihood expectation-maximization (MLEM) iterations. The computation time of the MLEM reconstruction is about 40 times longer than that of OSEM. In this example, there were only two (parallel-beam) projection angles per subset, which is clearly visible in the first OSEM iteration.

13.3.4.2. Refinements of the ordered-subsets expectation-maximization algorithm

As mentioned above, OSEM converges to a limit cycle: after many iterations, it starts cycling through a series of solutions rather than converging to the maximum-likelihood solution. When compared to the initial image (usually a uniform image), these series of solutions are ‘relatively close’ to the maximum-likelihood solution. Consequently, the convergence of OSEM is initially much faster but otherwise similar to that of MLEM; the better performance of MLEM only becomes noticeable at high iteration numbers. Thus, a simple solution to avoid the limit cycle is to gradually decrease the number of subsets: this approach preserves the initial fast convergence of OSEM, avoiding the limit cycle by returning to MLEM at high iteration numbers. A drawback of this approach is that convergence becomes slower each time the number of subsets is reduced. In addition, there is no theory available that prescribes how many sub-iterations should be used for each OSEM iteration.

Many algorithms have been proposed that use some form of relaxation to obtain convergence under less restrictive conditions than those of OSEM. As an example, relaxation can be introduced by rewriting the OSEM Eq. (13.71) in an

additive way. Then, a relaxation factor α is inserted to scale the update term to obtain RAMLA (row-action maximum-likelihood algorithm [13.27]):

$$\lambda_j^{\text{new}} = \lambda_j^{\text{old}} + \alpha \lambda_j^{\text{old}} \sum_{i \in \mathcal{S}_t} A_{ij} \left(\frac{y_i}{\hat{y}_i} - 1 \right) \quad \text{with} \quad \alpha < \frac{1}{\max_t \left(\sum_{i \in \mathcal{S}_t} A_{ij} \right)} \quad (13.73)$$

The relaxation factor α decreases with increasing iteration number to ensure convergence. It should be noted that setting $\alpha = 1 / \sum_{i \in \mathcal{S}_t} A_{ij}$ for all (sub-)iterations yields OSEM. Several alternative convergent block iterative algorithms have been proposed. They are typically much faster than MLEM but slightly slower than the (non-convergent) OSEM algorithm.

13.3.5. Regularization

MLEM maximizes the likelihood, by making the computed projections (from the current reconstruction) as similar as possible to the measured projections, where the similarity is measured based on the Poisson distribution. An upper limit of the likelihood would be obtained when the measured and calculated projections are identical. However, this is never possible, because Poisson noise introduces inconsistencies. Nevertheless, a large part of the noise is consistent, which means that it can be obtained as the projection of a (noisy) activity distribution. This part of the noise propagates into the reconstructed image, and is responsible for the so-called ‘deterioration’ of the MLEM image at high iterations.

13.3.5.1. Stopping iterations early

An ‘accidental’ feature of the MLEM algorithm is its frequency dependent convergence: low spatial frequencies converge faster than higher frequencies. This is due to the low-pass effect of the back projection operation. This effect is easily verified for the reconstruction of the activity in a point source, if the MLEM reconstruction is started from a uniform image. The first iteration then yields the back projection of the point source measurement. As discussed in Section 13.2.1, this yields an image with intensity $\lambda(x, y) \propto 1/\sqrt{x^2 + y^2}$, if the point source was located at (0,0). Each iteration multiplies with a similar back projection, implying that after t iterations, the image intensity at (x, y) is proportional to $1/(x^2 + y^2)^{t/2}$, so that the peak at (0,0) becomes a bit sharper with every iteration. For more complicated objects, the evolution is more subtle.

IMAGE RECONSTRUCTION




FIG. 13.15. Simulation study illustrating position dependent convergence in PET with attenuation. After 8 iterations (*iter*), convergence in highly attenuated regions is poor. After 100 iterations, good convergence is obtained, but with strong noise propagation. Post-smoothing yields a fair compromise between noise and nearly position independent resolution. FBP: filtered back projection.

It follows that reducing the number of iterations has an effect which is similar to reducing the cut-off frequency of a low-pass filter. However, the effect on the resolution is position dependent, as illustrated in Fig. 13.15. Attenuated PET projections of a highly radioactive uniform ring inside a less active disc were simulated with and without Poisson noise. After eight MLEM iterations, the reconstructed ring has non-uniform activity. In the centre of the phantom, convergence is slower, resulting in poorer resolution and poorer recovery of the activity in the ring. After 100 iterations, convergence is much better everywhere in the phantom, but for noisy data, there is very disturbing noise propagation.

If the image was acquired for detection (e.g. to see if there is a radioactive ring inside the disc or not), then the image produced after eight iterations is excellent. However, if the aim is quantification (e.g. analysing the activity distribution along the ring), then quantification errors can be expected at low iteration numbers.

13.3.5.2. Post-smoothed maximum-likelihood

The noise in the higher MLEM iterations is high frequency noise, and there are strong negative correlations between neighbouring pixels. As a result, a

modest amount of smoothing strongly suppresses the noise at the cost of a mild loss of resolution. This is illustrated in the third row of Fig. 13.15.

If the MLEM implementation takes into account the (possibly position dependent) spatial resolution effects, then the resolution should improve with every MLEM iteration. After many iterations, the spatial resolution should be rather good, similar or even better than the sinogram resolution, but the noise will have propagated dramatically. It is assumed that the obtained spatial resolution corresponds to a position dependent point spread function which can be approximated as a Gaussian with a full width at half maximum (FWHM) of $F_{ML}(x, y)$. Assume further that this image is post-smoothed with a (position independent) Gaussian convolution kernel with an FWHM of F_p . The local point spread function in the smoothed image will then have an FWHM of $\sqrt{(F_{ML}(x, y))^2 + F_p^2}$. If enough iterations are applied and if the post-smoothing kernel is sufficiently wide, the following relation holds $F_p \gg F_{ML}(x, y)$ and, therefore, $\sqrt{(F_{ML}(x, y))^2 + F_p^2} \approx F_p$. Under these conditions, the post-smoothed MLEM image has a nearly position independent and predictable spatial resolution. Thus, if PET or SPECT images are acquired for quantification, it is recommended to use many iterations and post-smoothing, rather than a reduced number of iterations, for noise suppression.

13.3.5.3. Smoothing basis functions

An alternative approach to counter noise propagation is to use an image representation that does not accommodate noisy images. Instead of representing the image with a grid of non-overlapping pixels, a grid of smooth, overlapping basis functions can be used. The two mostly used approaches are the use of spherical basis functions or ‘blobs’ [13.28] and the use of Gaussian basis functions or sieves [13.29].

In the first approach, the projector and back projector operators are typically adapted to work directly with line integrals of the basis functions. In the sieves approach, the projection of a Gaussian blob is usually modelled as the combination of a Gaussian convolution and projection along lines. The former approach produces a better approximation of the mathematics, while the latter approach yields a faster implementation.

The blobs or sieves are probably most effective when their width is very similar to the spatial resolution of the tomographic system. In this setting, the basis function allows accurate representation of the data measured by the tomographic system, and prevents reconstruction of much of the (high frequency) noise. It has been shown that using the blob during reconstruction is more effective than using the same blob only as a post-smoothing filter. The reason is that the post-filter

always reduces the spatial resolution, while a sufficiently small blob does not smooth data if it is used during reconstruction.

If the blob or sieve is wider than the spatial resolution of the tomographic system, then its use during reconstruction produces Gibbs over- and undershoots, also known as ‘ringing’. This effect always arises when steep edges have to be represented with a limited frequency range, and is related to the ringing effects observed with very sharp low-pass filters. For some imaging tasks, these ringing artefacts are a disadvantage.

13.3.5.4. Maximum a posteriori or penalized likelihood

Smoothing the MLEM image is not a very elegant approach: first, the likelihood is maximized, and then it is decreased again by smoothing the image. It seems more elegant to modify the objective function, such that the image that maximizes it does not need further processing. This can be done with a Bayesian approach, which is equivalent to combining the likelihood with a penalty function.

It is assumed that a good reconstruction image λ will be obtained if that image maximizes the (logarithm of the) probability $p(\lambda|y)$ given by Eq. (13.39) and repeated here for convenience:

$$\hat{\lambda} = \arg \max_{\lambda} (\ln p(y|\lambda) + \ln p(\lambda)) \quad (13.74)$$

The second term represents the a priori knowledge about the tracer distribution, and it can be used to express our belief that the true tracer distribution is fairly smooth. This is usually done with a Markov prior. In a Markov prior, the a priori probability for a particular voxel, given the value of all other voxels, is only a function of the direct neighbours of that voxel:

$$p(\lambda_j | \lambda_k, \forall k \neq j) = p(\lambda_j | \lambda_k, k \in N_j) \quad (13.75)$$

where N_j denotes the set of neighbour voxels of j .

Such priors are usually written in the following form:

$$P(\lambda) = \ln p(\lambda) = \sum_j \ln p(\lambda_j | \lambda_k, k \in N_j) = -\beta \sum_j \sum_{k \in N_j} E(\lambda_j \lambda_k) \quad (13.76)$$

where

the ‘energy’ function E is designed to obtain the desired noise suppressing behaviour and the parameter β is the weight assigned to the prior.

A higher weight results in smoother images, at the cost of a decreased likelihood, i.e. poorer agreement with the acquired data. In most priors, the expression is further simplified by making E a function of a single variable, the absolute value of the difference $|\lambda_j - \lambda_k|$.

Some popular energy functions $E(|\lambda_j - \lambda_k|)$ are shown in Fig. 13.16. A simple and effective one is the quadratic prior $E(x) = x^2$; an MAP reconstruction with this prior is shown in Fig. 13.17. Better preservation of strong edges is obtained with the Huber prior: it is quadratic for $|\lambda_j - \lambda_k| \leq \delta$ and linear for $|\lambda_j - \lambda_k| > \delta$, with a continuous first derivative at δ . Consequently, it applies less smoothing than the quadratic prior for differences larger than δ , as illustrated in Fig. 13.17. Even stronger edge tolerance is obtained with the Geman prior, which converges asymptotically to a constant for large differences, implying that it does not smooth at all over very large pixel differences.




FIG. 13.16. The energy function of the quadratic prior, the Huber prior and the Geman prior.




FIG. 13.17. Maximum-likelihood expectation-maximization (MLEM) and maximum a posteriori reconstructions of the Shepp–Logan phantom. Three different smoothing priors were used: quadratic, Huber and Geman. The latter smooth small differences quadratically, but are more tolerant for large edges.

It can be shown that the prior (Eq. (13.76)) is a concave function of λ if $E|\lambda_j - \lambda_k|$ is a convex function. Consequently, the quadratic and Huber energy

functions yield a concave prior: it has a single maximum. In contrast, the Geman prior is not concave (see Fig. 13.16) and has local maximums. Such concave priors require careful initialization, because the final reconstruction depends on the initial image and on the behaviour of the optimization algorithm.

Figure 13.18 shows that MAP reconstructions produce position dependent spatial resolution, similar to MLEM with a reduced number of iterations. The reason is that the prior is applied with a uniform weight, whereas the likelihood provides more information about some voxels than about others. As a result, the prior produces more smoothing in regions where the likelihood is ‘weaker’, e.g. regions that have contributed only a few photons to the measurement due to high attenuation.




FIG. 13.18. Maximum-likelihood expectation-maximization (MLEM), smoothed MLEM and maximum a posteriori (MAP) (quadratic prior) reconstructions of simulated PET data of a brain and a ring phantom. The ring phantom reveals position dependent smoothing for MAP.

The prior can be made position dependent as well, to ensure that the balance between the likelihood and the prior is about the same in the entire image. In that case, MAP with a quadratic prior produces images which are very similar to MLEM images with post-smoothing: if the prior and smoothing are tuned to produce the same spatial resolution, then both algorithms also produce nearly identical noise characteristics.

Many papers have been devoted to the development of algorithms for MAP reconstruction. A popular algorithm is the so-called ‘one step late’ algorithm. Inserting the derivative of the prior P in Eq. (13.66) yields:

$$\frac{\partial(L_x(\boldsymbol{\lambda})+P(\boldsymbol{\lambda}))}{\partial\lambda_j} = \sum_i \left(\frac{y_i}{\hat{y}_i^{(k)}} A_{ij} \lambda_j^{(k)} \frac{1}{\lambda_j} - A_{ij} \right) + \frac{\partial P(\boldsymbol{\lambda})}{\partial\lambda_j} = 0 \quad (13.77)$$

where $\hat{y}_i^{(k)}$ is the projection of the current reconstruction for detector i .

A problem with this equation is that $\partial P(\boldsymbol{\lambda})/\partial\lambda_j$ is itself a function of the unknown image $\boldsymbol{\lambda}$. To avoid this problem, the derivative of the prior is simply evaluated in the current reconstruction $\boldsymbol{\lambda}^{(k)}$. The equation can then be solved to produce the MAP update expression:

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)} \sum_i A_{ij} \frac{y_i}{\hat{y}_i^{(k)}}}{\sum_i A_{ij} - \left. \frac{\partial P(\boldsymbol{\lambda})}{\partial\lambda_j} \right|_{\boldsymbol{\lambda}^{(k)}}} \quad (13.78)$$

Owing to the approximation, convergence is not guaranteed. The algorithm usually works fine, except with very high values for the prior.

The MLEM algorithm can be considered as a gradient ascent algorithm (see also Eq. (13.50)):

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)} \sum_i A_{ij} \frac{y_i}{\hat{y}_i^{(k)}}}{\sum_i A_{ij}} \quad (13.79)$$

$$= \lambda_j^{(k)} + \frac{\lambda_j^{(k)} \left. \frac{\partial L_{ML}(\boldsymbol{\lambda})}{\partial\lambda_j} \right|_{\boldsymbol{\lambda}^{(k)}}}{\sum_i A_{ij}} \quad (13.80)$$

Extensions to an MAP gradient ascent algorithm typically have the form:

$$\lambda_j^{(k+1)} = \lambda_j^{(k)} + S(\boldsymbol{\lambda}^{(k)}) \left. \frac{\partial(L_{ML}(\boldsymbol{\lambda})+P(\boldsymbol{\lambda}))}{\partial\lambda_j} \right|_{\boldsymbol{\lambda}^{(k)}} \quad (13.81)$$

where the key is to determine a good preconditioner S .

Several methods with (almost) guaranteed convergence have been based on the previously described optimization transfer method, by designing useful surrogate functions for both the likelihood and the prior.

13.3.6. Corrections

In typical emission data, the true events (having a Poisson character) are distorted and contaminated by a number of physical effects. To make the best use of the acquired data and of our knowledge of the acquisition system, these effects should be included in the reconstruction model. The distortion effects include resolution effects (such as detector resolution, collimator effects, and in PET also non-collinearity and positron range) and motion effects. The contamination effects can be divided, by their character and the way they are treated, into multiplicative and additive terms. The multiplicative factors include: attenuation of the annihilation photons by the object, the probability of the detector elements detecting an event once they are hit by the photon (detector normalization factors), coefficients accounting for the decay time and the geometrical restriction of directions/LORs for which true events are detected (axial acceptance angle, detector gaps). The additive terms include scattered and random (in the PET case) coincidences. Details on calculation of the correction factors and terms are discussed in other chapters. This chapter is limited to the discussion of their utilization within the reconstruction process.

The most straightforward approach is to pre-correct the data before reconstruction for the contamination effects (multiplying by multiplicative correction coefficients and subtracting the scatter and random estimates), so as to approximate the X ray transform (or attenuated X ray transform in the SPECT case) of the reconstructed object. For analytical reconstruction approaches (derived for the ideal X ray transform data), the data always have to be pre-corrected.

For the statistical reconstruction methods, derived based on the statistical properties of the data, an attempt is made to preserve the Poisson character of the data as much as possible by including the correction effects inside the reconstruction model. Theoretically, the most appropriate way is to include the multiplicative and scatter effects directly into the system matrix. The system matrix would have to include not only an accurate model of the direct data (true events) but also of the physical processes of the generation of the contamination scatter data. In a sense, the contamination would then become valid data, bringing extra information to our model and, thus, adding valid (properly modelled) counts to the image. However, inclusion of the scatter model into the system matrix tremendously increases the number of non-zero elements of the system matrix, i.e. the matrix is not sparse anymore, and consequently the system is more

ill-posed (the contamination data are typically quite noisy) and computationally exceedingly expensive, and, thus, not feasible for routine clinical use.

The more practical, and commonly used, approach is to include correction effects as multiplicative factors and additive terms within the forward projection model of the iterative reconstruction approaches:

$$\mathbf{y} = \mathbf{A}\boldsymbol{\lambda} + \mathbf{b} \quad (13.82)$$

where the effects directly influencing the direct (true) data are included inside the system matrix \mathbf{A} and will be discussed in the following, while the additive terms \mathbf{b} (including scatter and randoms) will be discussed separately in Section 13.3.6.2 on additive terms.

13.3.6.1. Factors affecting direct events — multiplicative effects

In the PET case, the sequence of the physical effects (described in previous chapters) that occur as the true coincident events are generated and detected can be described by the following factorization of the system matrix \mathbf{A} as discussed in detail in Ref. [13.30]:

$$\mathbf{A} = \mathbf{A}_{\text{det.sens}} \mathbf{A}_{\text{det.blur}} \mathbf{A}_{\text{att}} \mathbf{A}_{\text{geom}} \mathbf{A}_{\text{tof}} \mathbf{A}_{\text{positron}} \quad (13.83)$$

where

- $\mathbf{A}_{\text{positron}}$ models the positron range;
- \mathbf{A}_{tof} models the timing accuracy for the TOF PET systems (TOF resolution effects, as discussed in Section 13.3.3.3);
- \mathbf{A}_{geom} is the geometric projection matrix, the core of the system matrix, which is a geometrical mapping between the source (voxel j) and data (projection bin i , defined by the LOR, or its time bin in the TOF case); the geometrical mapping is based on the probability (in the absence of attenuation) that photon pairs emitted from an individual image location (voxel) reach the front faces of a given crystal pair (LOR);
- \mathbf{A}_{att} is a diagonal matrix containing attenuation factors on individual LORs;
- $\mathbf{A}_{\text{det.blur}}$ models the accuracy of reporting the true LOR positions (detector resolution effects; discussed in Section 13.3.6.2);

and $\mathbf{A}_{\text{det.sens}}$ is a diagonal matrix modelling the probability that an event will be reported once the photon pair reaches the detector surface — a unique multiplicative factor for each detector crystal pair (LOR) modelled by

normalization coefficients, but can also include the detector axial extent and detector gaps.

In practice, the attenuation operation A_{att} is usually moved to the left (to be performed after the blurring operation). This is strictly correct only if the attenuation factors change slowly, i.e. they do not change within the range of detector resolution kernels. However, even if this is not the case, a good approximation can be obtained by using blurred (with the detector resolution kernels) attenuation coefficients. In this case, the multiplicative factors $A_{det.sens}$ and A_{att} can be removed from the system matrix A and applied only after the forward projection operation as a simple multiplication operation (for each projection bin). The rest of the system matrix (except $A_{positron}$, which is object dependent) can now be pre-computed, whether in a combined or a factorized form, since it is now independent of the reconstructed object. On the other hand, the attenuation factors A_{att} (and $A_{positron}$, if considered) have to be calculated for each given object.

In the SPECT case, the physical effects affecting the true events can be categorized and factorized into the following sequence:

$$A = A_{det.sens} A_{det.blur} A_{geom,att} \tag{13.84}$$

where

$A_{det.sens}$ includes multiplicative factors (such as detector efficiency and decay time);

$A_{det.blur}$ represents the resolution effects within the gamma camera (the intrinsic resolution of the system);

and $A_{geom,att}$ is the geometric projection matrix, also including the collimator effects (such as the depth dependent resolution) and the depth and view dependent attenuation factors.

For gamma cameras, the energy and linearity corrections are usually performed in real time, and the remaining (detector efficiency) normalization factors are usually very close to one and can be, for all practical purposes, ignored or pre-corrected. Similarly, the theory says that the decay correction should be performed during the reconstruction, because it is different for each projection angle. However, for most tracers, the decay during the scan is very modest, and in practice it is usually either ignored or done as a pre-correction. The attenuation component is object dependent and needs to be recalculated for each reconstructed object. Furthermore, its calculation is much more computationally

expensive than in the PET case, since it involves separate calculations of the attenuation factors for each voxel and for each view. This is one of the reasons why the attenuation factors have often been ignored in SPECT. More details on the inclusion of the resolution effects into the system matrix are discussed in Section 13.3.6.3.

13.3.6.2. Additive contributions

The main additive contaminations are scatter (SPECT and PET) and random events (PET). The simplest possibility of dealing with them is to subtract their estimates (\bar{s} and \bar{r}) from the acquired data. While this is a valid (and necessary) pre-correction step for the analytical reconstructions, it is not recommended for statistical approaches since it changes the statistical properties of the data, causing them to lose their Poisson character. As the maximum-likelihood algorithm is designed for Poisson distributed data, its performance is suboptimal if the data noise is different from Poisson. Furthermore, subtraction of the estimated additive terms from the noisy acquired data can introduce negative values into the pre-corrected data, especially for low count studies. The negative values have to be truncated before the maximum-likelihood reconstruction, since it is not able to correctly handle the negative data. This truncation, however, leads to a bias in the reconstruction.

On the other end of the spectrum of possibilities, would be considering the scatter and randoms directly in the (full) system model, that is, including a complete physical model of the scatter and random components into a Monte Carlo calculation of the forward projection. However, this approach is exceedingly computationally expensive and is not feasible for practical use. A practical and the most common approach for dealing with the additive contaminations is to add their estimate ($\bar{b} = \bar{s} + \bar{r}$) to the forward projection in the matrix model of the iterative reconstruction, i.e. the forward model is given by $A\lambda + \bar{b}$, as considered in the derivation of the MLEM reconstruction (Eq. (13.67)).

Special treatment has to be considered for clinical scanners in which the random events (r , estimated by delayed coincidences) are on-line subtracted from the acquired data (y , events in the coincidence window — prompts). The most important characteristic of the Poisson data is that their mean equals their variance: $\text{mean}(y_i) = \text{var}(y_i)$. However, after the subtraction of the delays from the prompts (both being Poisson variables), the resulting data (γ) are not Poisson anymore, since $\text{mean}(\gamma_i) = \text{mean}(y_i - r_i) = \text{mean}(y_i) - \text{mean}(r_i)$, while $\text{var}(\gamma_i) = \text{var}(y_i - r_i) = \text{var}(y_i) + \text{var}(r_i)$. To regain the main characteristic of the Poisson data (at least of the first two moments), the shifted Poisson approach can be used, utilizing the fact that adding a (noiseless) constant value to the Poisson variable changes the mean but preserves the variance

of the result. To modify the mean of the subtracted data γ to be equal to their variance (i.e. $\text{var}(y_i) + \text{var}(r_i)$), we need to add to the subtracted data an estimate (of the mean) of the randoms (\bar{r}) multiplied by two. This gives $\text{mean}(\gamma_i + 2\bar{r}_i) = \text{mean}(y_i - r_i + 2\bar{r}_i) = \text{mean}(y_i) + \text{mean}(r_i)$, which is equal to $\text{var}(\gamma_i + 2\bar{r}_i) = \text{var}(y_i) + \text{var}(y_i)$. The MLEM algorithm using the shifted Poisson model can then be written as:

$$\lambda_j^{(k+1)} = \frac{\lambda_j^{(k)} \sum_i A_{ij} \frac{\gamma_i + 2\bar{r}_i}{\sum_j A_{ij} \lambda_j^{(k)} + \bar{s}_i + 2\bar{r}_i}}{\sum_i A_{ij}} \quad (13.85)$$

It is worthwhile mentioning here that even in the shifted Poisson case, the negative values in the subtracted data and consequent truncation leading to the bias and artefacts cannot be completely avoided. However, the chance of the negative values decreases since the truncation of the negative values is being performed on the ‘value-shifted’ data $(\gamma_i + 2\bar{r}_i)$. Examples of reconstructions from data with a subtracted additive term, using the regular MLEM algorithm and using MLEM with the shifted Poisson model, are shown in Fig. 13.19. As the counts were relatively high in this simulation, the subtraction did not produce




FIG. 13.19. Illustration of (exaggerated case of) reconstructions from contaminated data y from which the additive contamination term r was subtracted (both data and contamination term are Poisson). The top row shows the sinograms. The increased noise level in the contaminated area in the sinogram $(y - r)$ should be noted. The bottom row shows the true image without and with the contaminator; the maximum-likelihood expectation-maximization (MLEM) reconstruction from the subtracted data $(y - r)$ and the shifted Poisson MLEM reconstruction, in which the estimated (noiseless) additive term $2\bar{r}$ is added to the subtracted data and forward projection as given by Eq. (13.85).

negatives. MLEM of $(\mathbf{y} - \mathbf{r})$ creates streaks because the reliability of the subtracted data is overestimated.

It should be noted that in the reconstruction model (as well as in the pre-correction approaches) the estimates of the scatter and randoms have to be treated in the same way as the estimates of the true events in the forward projection, including consideration of the normalized or un-normalized events, attenuation corrected or uncorrected data, gaps in the data, etc. Various challenges exist for the scatter and randoms estimations in general, such as modelling of the out of FOV scatter. This is addressed in Chapter 11.

13.3.6.3. Finite spatial resolution

There are a number of physical and geometrical effects and limitations (such as positron range, non-collinearity, depth of interaction, size of detector crystal elements, inter-crystal scatter, collimator geometry, etc.) affecting PET and SPECT resolution as described in more detail in Chapter 11. To get the most out of the acquired data and to correct for the resolution degradation, these effects have to be properly modelled in the system matrix of statistical reconstruction, as considered in the components ($A_{\text{det.blur}}$, A_{geom} , A_{positron}) of the factorized system matrix outlined in Section 13.3.6.1. This step does not influence the mathematical definition of the reconstruction algorithm (such as MLEM, as given by Eq. (13.67)); only the form of its system matrix is changed.

However, this step has very practical consequences for the complexity of the algorithm implementation, for computational demands and most importantly for the quality of the reconstructed images. By including the resolution effects into the reconstruction model, a larger fraction of the data is being used for the reconstruction within each point of the space, with the true signal component becoming more consistent, while the noise components becoming less consistent with the model. Thus, the resolution modelling helps twice, by improving the image resolution while at the same time reducing the image noise, as illustrated in Fig. 13.20 for simulated SPECT data. This is quite different from the filtering case, where the noise suppression is always accompanied by resolution deterioration. On the other hand, the resolution modelling has a price in terms of a considerable increase in the computational load (both in space/memory and time) since the system matrix is much less sparse, that is, it contains a larger proportion of non-zero elements. This not only leads to more computational load per iteration, but also to a slower convergence of the iterative reconstruction and, consequently, to the need for more iterations.

Resolution effects can be subdivided into the effects dependent on the particular object, such as the positron range, and the effects influenced by the scanner geometry, design and materials (which can be determined beforehand

for the given scanner). The positron range depends on the particular attenuation structures in which the the positrons annihilate, and also varies from isotope to isotope. Furthermore, the shape of the probability function (kernel) of the positron annihilation abruptly changes at the boundaries of two tissues, such as at the boundary of the lungs and surrounding soft tissues, and, thus, it strongly depends on the particular object's morphology and is quite challenging to model accurately. In general, the positron range has a small effect (compared to the other effects) for clinical scanners, particularly for studies using ^{18}F -labelled tracers, and can often be ignored. However, for small animal imaging and for other tracers (such as ^{82}Rb), the positron range becomes an important effect to be considered.




FIG. 13.20. Examples of the effects of resolution modelling within statistical iterative reconstruction. Data were simulated for a SPECT system with depth dependent resolution. It is clearly seen that using the proper resolution model within statistical reconstruction (lower two images on the right) not only improves the resolution of the images, but also helps to efficiently suppress the noise component.

There is a whole spectrum of approaches to determine and implement the scanner dependent resolution models. Only the main ones are addressed. The simplest, but least accurate, approach is to approximate the system resolution model by a spatially invariant resolution kernel, usually a spherically symmetric Gaussian, with the shape (FWHM) estimated from point source measurements

at one or more representative locations within the given scanner. This approach typically provides satisfactory results within the central FOV of large, whole body PET scanners. However, for PET systems with smaller ring diameters (relative to the reconstruction FOV), such as animal systems, and for SPECT systems with depth dependent resolution (and in particular with non-circular orbits), it is desirable to use more accurate spatially variant resolution models.

The second category is using analytically calculated resolution functions (usually spatially variant anisotropic kernels) for each location (LOR) as determined based on analytical models of physical effects affecting the resolution. This approach is usually limited to simple analytical models representing (or approximating) only basic physical characteristics of the system. The resolution kernels are usually calculated in real time during the reconstruction process when they are needed within the forward and back projection calculations. In SPECT, distance dependent collimator blurring requires convolution kernels that become wider and, therefore, need more computation, with increasing distance to the collimator. The computation time can be reduced considerably by integrating an incremental blurring step into the projector (and back projector), based on Gaussian diffusion. This method, developed by McCarthy and Miller in 1991, is described in more detail in chapter 22 of Ref. [13.5].

A more accurate but computationally very demanding approach is using Monte Carlo simulations of the resolution functions based on a set of point sources at various (ideally all) image locations. Setting up an accurate mathematical model (transport equations tracing the photon paths through the detector system/crystals) is relatively easy within the Monte Carlo simulations, compared to the analytical approach of determining the resolution function. However, to obtain sufficient statistics to get the desired accuracy of the shape of the resolution functions is extremely time consuming. Consequently, simplifications often have to be made in practice, such as determining the resolution kernels only at a set of representative locations and interpolating/extrapolating from them the resolution kernels at other locations.

The most accurate but also most involved approach is based on experimental measurements of the system response by measuring physical point sources at a set of image locations within the scanner. This is a tedious and very time consuming process, involving point sources with long half-life isotopes and usually requiring the use of accurate robotic stages to move the point source. Among the biggest challenges is to accumulate a sufficient number of counts to obtain an accurate point spread function, even at a limited number of locations. Consequently, the actual resolution kernels used in the reconstruction model are often estimated by fitting analytical functions (kernels) to the measured data, rather than directly using the measured point spread functions.

At the conclusion of this subsection, it is worth making the following general comment. In the light of the resolution modelling possibilities discussed above, one might wonder whether it is worth spending energy and resources on building new PET and SPECT systems with improved resolution properties. However, although it has been shown in the literature that proper system models lead to improved reconstructed image quality, they can never fully recover information that has been lost through resolution effects and other instrumentation limitations. Furthermore, due to the increased level of modelling, the system matrix becomes more dense, and consequently the inverse problem (reconstruction) becomes more ill-posed, thus making it impossible to attain perfect recovery for the realistic data. There is no doubt that improved instrumentation as well as novel and more accurate reconstruction models play an important role in improving image quality and quantitative accuracy, and eventually increasing the general clinical utility of emission tomography systems.

13.3.6.4. Motion corrections

Owing to the relatively long acquisition times, motion effects, caused by patient movement and organ motion and deformation, cannot be avoided in emission tomography. In the following, all of these effects are covered under the simple term 'motion'. With the continuous improvements of PET and SPECT technology, leading to improved spatial resolution, signal to noise ratio, image quality and accuracy of quantitative studies, corrections for motion effects become more important. In fact, artefacts caused by motion are becoming the single most important factor for image degradation, especially in PET or PET/computed tomography (CT) imaging of the upper torso region. For example, motion effects can lead to the loss of small lesions by blurring them out completely in regions with strong motion (such as near the lower lung wall), or to their misplacement into the wrong anatomical region (e.g. into the liver from the lungs, or vice versa). Motion correction has become an important research topic; however, a thorough discussion of this topic is out of the scope of this chapter and interested readers are referred to the literature on this topic. In the following, the main concepts of motion correction as dealt with within the reconstruction process are outlined.

The two main sources of motion related artefacts in emission studies are the motion during the emission scan and the discrepancy (caused by the motion) between the attenuation and emission data. The motion during the emission scan means that the emission paths (LORs) through the object (as considered in the system matrix) change during the scan time. If this time dependent change is not accounted for, the system model becomes inconsistent with the data, which results in artefacts and motion blurring in the reconstructed images. On the other

hand, the transmission scan (CT) is relatively short and can usually be done in a breath-hold mode. Consequently, the attenuation image is usually motion-free and captures only one particular patient position and organ configuration (time frame). If the attenuation factors obtained from this fixed-time position attenuation image are applied to the emission data acquired at different time frames (or averaged over many time frames), this leads to artefacts in the reconstructed images, which tend to be far more severe in PET than in SPECT. This is, for example, most extremely pronounced at the bottom of the lungs which can typically move several centimetres during the breathing cycle, causing motion between two regions with very different attenuation coefficients.

Emission data motion: Correction approaches for motion during the emission scan are discussed first. The first step is subdividing the data (in PET, typically list-mode data) into a sufficient number of time frames to ensure that the motion within each frame is small. For the organ movement, the frames can be distributed over a period of the organ motion (e.g. breathing cycle). For the patient motion, the frames would be typically longer and distributed throughout the scan time. Knowledge about the motion can be obtained using external devices, such as cameras with fiducial markers, expansion belts or breathing sensors for respiratory motion, the electrocardiogram signal for cardiac motion, etc. There are also a limited number of approaches for estimating the motion directly from the data.

Once the data are subdivided into the set of the frames, the most straightforward approach is to reconstruct data independently in each frame. The problem with this approach is that the resulting images have a poor signal to noise ratio because the acquired counts have been distributed into a number of individual (now low count) frames. To improve the signal to noise ratio, the reconstructed images for individual frames can be combined (averaged) after they are registered (and properly deformed) to the reference time frame image. However, for statistical non-linear iterative reconstruction algorithms, this is not equivalent to (and typically of a lower quality than) the more elaborate motion correction approaches, taking into account all of the acquired counts in a single reconstruction, as discussed below.

For rigid motion (e.g. in brain imaging), the events on LORs (LOR_i) from each time frame, or time position, can be corrected for motion by translation (using affine transformations) into the new LORs (LOR_j) in the reference frame (see Fig. 13.21 (top right, solid line)), in which the events would be detected if there were no motion. Reconstruction is then done in a single reference frame using all acquired counts, leading to a better signal to noise ratio in the reconstructed images. Care has to be taken with the detector normalization factors so that the events are normalized using the proper factors (N_i) for the LORs on which they were actually detected (and not into which they were translated).

IMAGE RECONSTRUCTION

Attenuation factors are obtained on the transformed lines (att_i) through the attenuation image in the reference frame. Care also has to be given to the proper treatment of data LORs with events being translated into, or out of, the detector gaps or detector ends. This is important, in particular for the calculation of the sensitivity matrix, which then becomes a very time consuming process.




FIG. 13.21. Illustration of motion corrections for events acquired within line of response LOR_i with corresponding normalization N_i and attenuation att_i factors. Left top: positions and shapes of the object in the reference time frame 0 and frame k . Left bottom: illustration of blurring in the reconstruction combining events from all frames without motion correction (attenuation factors are also averaged over the whole range of the frames att_i^{0-k}). Middle column: processing within the reference time frame. Right top: LOR based motion correction for frame k — the LOR_i (dashed line) has to be transformed to the LOR_i (solid line for rigid motion, dotted line for non-rigid motion) which represents the paths that the photons would travel through the reference object if there were no motion. It should be noted that although the LORs are transformed, the normalization factors are used for the crystal pairs (LORs) in which the events were detected (N_i), while the used attenuation factors are for the transformed paths (att_i). Right bottom: image based motion correction, including image morphing of the estimated image from the reference frame (dashed lines) into the given frame (solid line).

For non-rigid (elastic) motion, which is the case for most of the practical applications, the motion correction procedures become quite involved. There are two basic possibilities. The first approach is to derive the transformations of individual paths of events (LORs) from each frame into the reference frame (see Fig. 13.21 (top right, dotted line)). For the non-rigid motion, the transformed paths through the reference object frame are not straight lines anymore, thus

leading to very large computational demands for the calculations of the forward and back projection operations. The same care for normalization, gaps and detector ends has to be taken as above.

The second, more efficient, approach involves morphing the image estimate (of the reference image) into the frame for which current events (LORs) are being processed (see Fig. 13.21 (bottom right, solid line)). It should be noted that some pre-sorting of the data is considered, so that events from each frame are processed together (using a common image morphing operation). Here, the acquired LORs (LOR_i) and their normalization coefficients (N_i) are directly used without modification. However, the sensitivity matrix still needs to be carefully calculated, taking into consideration update and subset strategy, e.g. including the morphing operation if subset data involve several frames. This is, however, a simpler operation than in the LOR based case since the morphing is done in the image domain. This image based approach is not only more efficient, but also better reflects/models the actual data acquisition process during which the acquired object is being changed (morphed).

Attenuation effects: In the following, it is considered that either attenuation information for each time frame is available, for example, having a sequence of CT scans for different time positions, or there is knowledge of the motion and tools to morph a fixed-time position CT image to represent attenuation images at individual time frames. It is further considered that tools are available to obtain the motion transformation of data and/or images between the individual time frames. If the emission data are stored or binned without any motion gating, they represent motion-blurred emission information over the duration of the scan. Using attenuation information for them for a fixed time position is not correct. It would be better to pre-correct those data using proper attenuation factors for each frame, but then the statistical properties (Poisson character) are lost due to the pre-correction. A good compromise (although not theoretically exact) is to use motion-blurred attenuation factors during the pre-correction or the reconstruction process.

For data stored in multiple time frames, separate attenuation factors (or their estimates) are used for each frame, such that they reflect attenuation factors (for each LOR) at that particular time frame. For the case when there are multiple CT images, this is simply obtained by calculation (forward projection) of the attenuation coefficients for each frame from the representative CT image for that frame. For the case when there is only one CT image, attenuation factors have to be calculated on the modified LORs (for each time frame) in the LOR based corrections, or to morph the attenuation image for each frame and then calculate the attenuation factors from the morphed images in the image based corrections.

13.4. NOISE ESTIMATION

13.4.1. Noise propagation in filtered back projection

The pixel variance in an image reconstructed with FBP can be estimated analytically, by propagating the uncorrelated Poisson noise in the data through the reconstruction operation. The FBP algorithm can be written as:

$$\Lambda(x, y) = \int_0^x d\phi \int_{-\infty}^{\infty} Y(x \cos \phi + y \sin \phi - s) h(s) ds \quad (13.86)$$

where $h(s)$ is the convolution kernel, combining the inverse Fourier transform of the ramp filter and a possible low-pass filter to suppress the noise.

The variance on the measured sinogram $Y(s, \phi)$ data equals its expectation $\bar{Y}(s, \phi)$; the covariance between two different sinogram values $Y(s, \phi)$ and $Y(s', \phi')$ is zero. Consequently, the covariance between two reconstructed pixel values $\Lambda(x, y)$ and $\Lambda(x', y')$ equals:

$$\begin{aligned} \text{covar}(\Lambda(x, y), \Lambda(x', y')) &= \int_0^x d\phi \int_{-\infty}^{\infty} \bar{Y}(x \cos \phi + y \sin \phi - s) ds \\ &\quad h(s) h(s + (x' - x) \cos \phi + (y' - y) \sin \phi) \end{aligned} \quad (13.87)$$

This integral is non-zero for almost all pairs of pixels. As $h(s)$ is a high-pass filter, neighbouring reconstruction pixels tend to have fairly strong negative correlations. The correlation decreases with increasing distance between (x, y) and (x', y') . The variance is obtained by setting $x = x'$ and $y = y'$, which produces:

$$\text{var}(\Lambda(x, y)) = \int_0^x d\phi \int_{-\infty}^{\infty} \bar{Y}(x \cos \phi + y \sin \phi - s) |h(s)|^2 ds \quad (13.88)$$

Figure 13.22 shows the variance image of the FBP reconstruction of a simulated PET sinogram of a heart phantom. The image was obtained by reconstructing 400 sets of noisy PET data. The figure also shows a noise-free and one of the noisy FBP images. The noise creates streaks that extend to the edge of the image. As a result, the variance is non-zero in the entire image.




FIG. 13.22. Simulated PET reconstructions of a heart phantom. Reconstructions were done with filtered back projection (FBP), maximum-likelihood expectation-maximization (MLEM) with Gaussian post-smoothing and with maximum a posteriori (MAP) using a quadratic prior. For each algorithm, a noise-free and a noisy reconstruction are shown, and also the pixel variance obtained from 400 independent Poisson noise realizations on the simulated PET data. All reconstructions (first two rows) are shown on the same grey value scale. A second scale was used to display the three variance images. The noisy FBP image contains negative pixels (displayed in white with this scale).

13.4.2. Noise propagation in maximum-likelihood expectation-maximization

The noise analysis of MLEM (and MAP) reconstruction is more complicated than that for FBP because these algorithms are non-linear. However, the MLEM algorithm has some similarity with the WLS algorithm, which can be described with matrix operations. The WLS reconstruction was described previously; Eq. (13.45) is repeated here for convenience (the additive term was assumed to be zero for simplicity):

$$\lambda = (A' C_y^{-1} A)^{-1} A' C_y^{-1} y \quad (13.89)$$

C_y is the covariance of the data, which is defined as $C_y = E(y - \bar{y})(y - \bar{y})'$, where E denotes the expectation and \bar{y} is the expectation of y .

The covariance of the reconstruction is then:

$$\begin{aligned}
 C_{\lambda} &= E(\lambda - \bar{\lambda})(\lambda - \bar{\lambda})' \\
 &= (A' C_y^{-1} A)^{-1} A' C_y^{-1} E(y - \bar{y})(y - \bar{y})' C_y^{-1} A (A' C_y^{-1} A)^{-1} \\
 &= (A' C_y^{-1} A)^{-1}
 \end{aligned}
 \tag{13.90}$$

This matrix gives the covariances between all possible pixel pairs in the image produced by WLS reconstruction. The projection A and back projection A' have a low pass characteristic. Consequently, the inverse $(A' C_y^{-1} A)^{-1}$ acts as a high-pass filter. It follows that neighbouring pixels of WLS reconstructions tend to have strong negative correlations, as is the case with FBP. Owing to this, the MLEM variance decreases rapidly with smoothing.

Figure 13.22 shows mean and noisy reconstructions and variance images of MLEM with Gaussian post-smoothing and MAP with a quadratic prior. For these reconstructions, 16 iterations with 8 subsets were applied. MAP with a quadratic prior produces fairly uniform variance, but with a position dependent resolution. In contrast, post-smoothed MLEM produces fairly uniform spatial resolution, in combination with a non-uniform variance.

REFERENCES

- [13.1] LEWITT, R.M., MATEJ, S., Overview of methods for image reconstruction from projections in emission computed tomography, Proc. IEEE Inst. Electr. Electron. Eng. **91** (2003) 1588–1611.
- [13.2] NATTERER, F., The Mathematics of Computerized Tomography, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA (1986).
- [13.3] KAK, A.C., SLANEY, M., Principles of Computerized Tomographic Imaging, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA (1988).
- [13.4] BARRETT, H.H., MYERS, K.J., Foundations of Image Science, John Wiley and Sons, Hoboken, NJ (2004).
- [13.5] WERNICK, M.N., AARSVOLD, J.N. (Eds), Emission Tomography, The Fundamentals of PET and SPECT, Elsevier Academic Press (2004).
- [13.6] NATTERER, F., Inversion of the attenuated radon transform, Inverse Probl. **17** (2001) 113–119.
- [13.7] XIA, W., LEWITT, R.M., EDHOLM, P.R., Fourier correction for spatially variant collimator blurring in SPECT, IEEE Trans. Med. Imaging **14** (1995) 100–115.
- [13.8] DEFRISE, M., CLACK, R., TOWNSEND, D.W., Image reconstruction from truncated, two-dimensional, parallel projections, Inverse Probl. **11** (1996) 287–313.

- [13.9] DEFRISE, M., KUIJK, S., DECONINCK, F., A new three-dimensional reconstruction method for positron cameras using plane detectors, *Phys. Med. Biol.* **33** (1988) 43–51.
- [13.10] KINAHAN, P.E., ROGERS, J.G., Analytic three-dimensional image reconstruction using all detected events, *IEEE Trans. Nucl. Sci.* **NS-36** (1990) 964–968.
- [13.11] DAUBE-WITHERSPOON, M.E., MUEHLLEHNER, G., Treatment of axial data in three-dimensional PET, *J. Nucl. Med.* **28** (1987) 1717–1724.
- [13.12] LEWITT, R.M., MUEHLLEHNER, G., KARP, J.S., Three-dimensional image reconstruction for PET by multi-slice rebinning and axial image filtering, *Phys. Med. Biol.* **39** (1994) 321–339.
- [13.13] DEFRISE, M., et al., Exact and approximate rebinning algorithms for 3D PET data, *IEEE Trans. Med. Imaging* **16** (1997) 145–158.
- [13.14] DEFRISE, M., A factorization method for the 3D X-ray transform, *Inverse Probl.* **11** (1995) 983–994.
- [13.15] TOMITANI, T., Image reconstruction and noise evaluation in photon time-of-flight assisted positron emission tomography, *IEEE Trans. Nucl. Sci.* **NS-28** (1981) 4582–4589.
- [13.16] QI, J., LEAHY, R.M., Iterative reconstruction techniques in emission computed tomography, *Phys. Med. Biol.* **51** (2006) R541–578.
- [13.17] FESSLER, J.A., BOOTH, S.D., Conjugate-gradient preconditioning methods for shift variant PET image reconstruction, *IEEE Trans. Image Process.* **8** (1999) 688–699.
- [13.18] PRESS, W.H., FLANNERY, B.P., TEUKOLSKY, S.A., VETTERLING, W.T., *Numerical Recipes, The Art of Scientific Computing*, Cambridge University Press (1986).
- [13.19] DE PIERRO, A.R., A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography, *IEEE Trans. Med. Imaging* **14** (1995) 132–137.
- [13.20] SHEPP, L.S., VARDI, Y., Maximum likelihood reconstruction for emission tomography, *IEEE Trans. Med. Imaging* **MI-1** (1982) 113–122.
- [13.21] DEMPSTER, A.P., LAIRD, N.M., RUBIN, D.B., Maximum likelihood from incomplete data via the EM algorithm, *J. R. Stat. Soc. Series B Stat. Methodol.* **39** (1977) 1–38.
- [13.22] PARRA, L., BARRETT, H.H., List-mode likelihood: EM algorithm and image quality estimation demonstrated on 2-D PET, *IEEE Trans. Med. Imaging* **17** 2 (1998) 228–235.
- [13.23] READER, A.J., ERLANDSSON, K., FLOWER, M.A., OTT, R.J., Fast accurate iterative reconstruction for low-statistics positron volume imaging, *Phys. Med. Biol.* **43** 4 (1998) 835–846.
- [13.24] QI, J., Calculation of the sensitivity image in list-mode reconstruction, *IEEE Trans. Nucl. Sci.* **53** (2006) 2746–2751.

IMAGE RECONSTRUCTION

- [13.25] MATEJ, S., et al., Efficient 3-D TOF PET reconstruction using view-grouped histogram images: DIRECT — Direct Image Reconstruction for TOF, *IEEE Trans. Med. Imaging* **28** (2009) 739–751.
- [13.26] HUDSON, M.H., LARKIN, R.S., Accelerated image reconstruction using ordered subsets of projection data, *IEEE Trans. Med. Imaging* **13** (1994) 601–609.
- [13.27] BROWNE, J., DE PIERRO, A.R., A row-action alternative to the EM algorithm for maximizing likelihoods in emission tomography, *IEEE Trans. Med. Imaging* **15** (1996) 687–699.
- [13.28] DAUBE-WITHERSPOON, M.E., MATEJ, S., KARP, J.S., LEWITT, R.M., Application of the row action maximum likelihood algorithm with spherical basis functions to clinical PET imaging, *IEEE Trans. Nucl. Sci.* **48** (2001) 24–30.
- [13.29] SNYDER, D.L., MILLER, M.I., THOMAS, L.J., Jr., POLITTE, D.G., Noise and edge artefacts in maximum-likelihood reconstructions for emission tomography, *IEEE Trans. Med. Imaging* **MI-6** (1987) 228–238.
- [13.30] LEAHY, R.M., QI, J., Statistical approaches in quantitative positron emission tomography, *Stat. Comput.* **10** (2000) 147–165.