ELEC9370:
Digital Image Processing
Chapter 9: Computed Tomography

Dr. D. S. Taubman

April 12, 2003

1 Introduction

Figure 1 illustrates a somewhat idealized tomographic imaging system. In this case a subject is illuminated with a uniform X-ray source. The X-rays are oriented at angle θ to the horizontal and the attenuated signal is picked up by a linear array of sensors. Ideally, there are infinitely many sensors and the image is acquired at an infinite number of angles, so that we are able to recover a two dimensional, spatially continuous intensity data set, $I(\theta, s)$, where θ ranges from 0 to π and s ranges from 0 to S. Here, S depends on the physical size of the object which is being imaged.
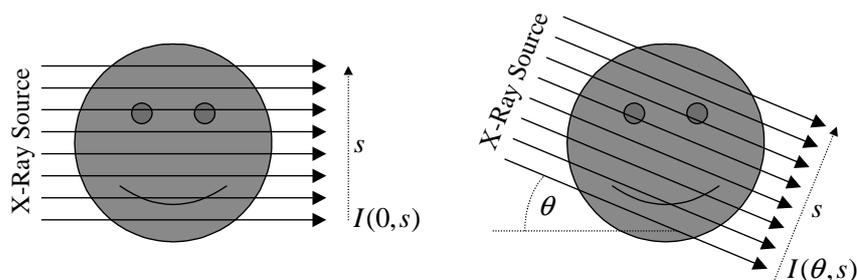
Figure 1: Idealized tomographic imaging.

1.1　The Radon Transform

X-rays are attenuated as they travel through the body, in accordance with the absorbtion characteristics (attenuation) of the different tissues and bone matter which are encountered. Of course, this is also an idealized model, since we ignore the effects of scattering. The idea of computed tomography is to compute the attenuation map, $a(x, y)$, of the subject's interior. In the idealized model we have

$$I(\theta, s) = I_0 \exp \left(\oint_{L_{\theta,s}} \ln a(x, y) \right)$$

where I_0 is the strength of the received signal in the absence of any attenuation, and $\oint_{L_{\theta,s}}$ denotes the line integral along the line at angle θ and position s. Taking the logarithm of both sides, we obtain

$$\ln I(\theta, s) = \ln I_0 - \oint_{L_{\theta,s}} d(x, y)$$

where $d(x, y) = -\log a(x, y)$ denotes the density at location (x, y) inside the subject.

Finally, writing

$$r(\theta, s) = \ln I_0 - \ln I(\theta, s) = \ln \frac{I_0}{I(\theta, s)}$$

for the reduction in log-intensity due to X-rate attenuation, we obtain

$$
\begin{aligned}
r(\theta, s) &= \oint_{L_{\theta,s}} d(x, y) \\
&= \int_{-\infty}^{\infty} d(t \cos \theta + s \sin \theta, -t \sin \theta + s \cos \theta) \, dt \quad\quad (1)
\end{aligned}
$$

This expansion of the line integral may be understood with the aid of Figure 2.Equation (1) is known as the Radon transform of $d(x, y)$ and inversion of the Radon transform is the objective of computed tomography. The coordinate transform expressed in equation (1) may be identified with the rotation matrix,

$$R_\theta = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \quad\quad (2)$$

where

$$\begin{pmatrix} x \\ y \end{pmatrix} = R_\theta \begin{pmatrix} t \\ s \end{pmatrix}$$
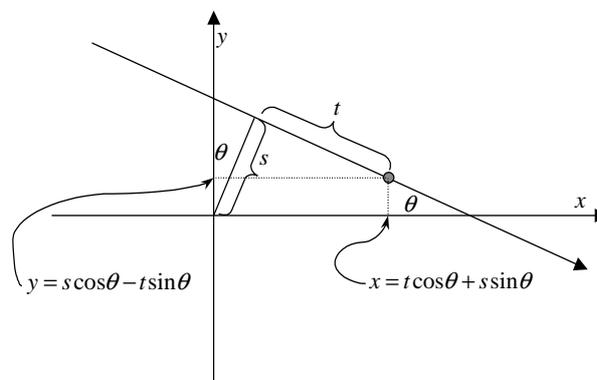
Figure 2: Coordinate system for the Radon transform

1.2   Other Applications

In the development above, we have considered only an idealized X-ray CT imaging system. However, there are a number of other important imaging applications whose outputs are related (directly or indirectly) to line integrals of some parameter field at various angles and positions. Examples include Positron Emission Tomography (PET) and Magnetic Resonance Imaging (MRI).

In MRI, the subject is exposed to a high intensity magnetic field, with a superimposed gradient field. In this way, the magnetic field is constant along lines running perpendicular to the field gradient. When the subject is exposed to RF radiation, protons absorb some of the RF energy in the form of an altered (flipped) magnetic moment. When the RF source is switched off, the protons relax (spin-lattice relaxation), emitting RF energy whose frequency is directly related to the magnetic field to which they are exposed through

$$\omega\left(s\right) = \gamma B\left(s\right)$$

where $B\left(s\right)$ is the magnetic field strength at position s along the magnetic field gradient and γ is the gyro-magnetic ratio for the material being imaged (usually hydrogen bound in water molecules). $\omega\left(s\right)$ is called the Lamour frequency and it is a function of the position s along the magnetic field gradient.

In this way, the energy detected at frequency $\omega\left(s\right)$ corresponds to the line integral of the density of water molecules along the line running perpendic-

ular to the magnetic field gradient, at position s. Equivalently, the received signal may be understood as the Fourier transform, along the s-direction, of the Radon transform, $r(\theta, s)$, of the water molecule density, $d(x, y)$, where θ is the orientation of the magnetic gradient field. The gradient orientation can be rotated electronically by mixing contributions from two separate gradient fields, one in the X-direction and one in the Y-direction.

2   Projection Slice Theorem

It is helpful to express the Radon transform relationship in the frequency domain. In particular, let $\hat{r}(\theta, \omega)$ denote the one dimensional Fourier transform in the s-direction of $r(\theta, s)$ and let $\hat{d}(\omega_x, \omega_y)$ denote the two dimensional Fourier transform of $d(x, y)$. Then

$$
\begin{aligned}
\hat{r}(\theta, \omega) &= \int_{-\infty}^{\infty} e^{-j\omega s} r(\theta, s)\, ds \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-j\omega s} d(t\cos\theta + s\sin\theta, -t\sin\theta + s\cos\theta)\cdot ds \cdot dt
\end{aligned}
$$

Now making the coordinate substitution,

$$
\begin{pmatrix} x \\ y \end{pmatrix} = R_\theta \begin{pmatrix} s \\ t \end{pmatrix}
$$

with R_θ as in equation (2), we have

$$
dx \cdot dy = \det(R_\theta)\, dx \cdot dy = dx \cdot dy
$$

and

$$
\begin{pmatrix} s \\ t \end{pmatrix} = R_\theta^{-1} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}
$$

so that

$$
s = x\cos\theta - y\sin\theta
$$

and

$$
\begin{aligned}
\hat{r}(\theta, \omega) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-j\omega(x\cos\theta - y\sin\theta)} d(x, y)\cdot dx \cdot dy \\
&= \hat{d}(\omega_x, \omega_y)\Big|_{\omega_x = \omega\cos\theta,\, \omega_y = -\omega\sin\theta}
\end{aligned}
$$
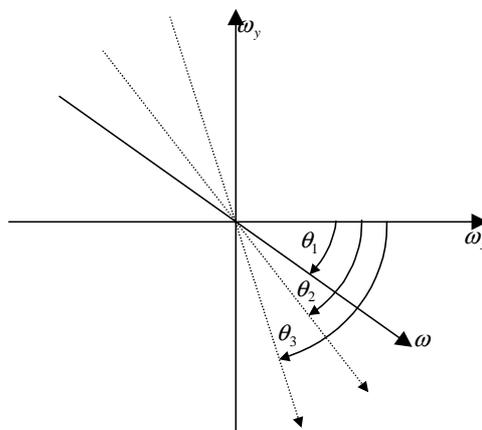
Figure 3: Illustration of the projection slice relationship between the Radon and Fourier transforms of a two dimensional signal, showing several different orientations, θ_i.

This result is known as the "projection slice theorem." It states that the one dimensional Fourier transform of the Radon transform, $\hat{r}(\theta, \omega)$, is nothing other than the two dimensional Fourier transform of the density function, $\hat{d}(\omega_x, \omega_y)$, transformed to polar coordinates. The relationship is illustated in Figure 3.

The projection slice theorem suggests an immediate approach to inverting the Radon transform. First form $\hat{r}(\theta, \omega)$. Then interpolate between the available orientations, θ, in the Fourier domain to estimate $\hat{d}(x, y)$ and take the inverse Fourier transform to recover $d(x, y)$. This approach is even more attractive when used with MRI, since the output of the MRI system is already in the Fourier domain, i.e., $\hat{r}(\theta, \omega)$. In fact, many MRI devices are able to implement the polar to rectangular coordinate transformation through the imaging electronics themselves (this is the so-called "Fourier geometry imaging mode"), so that the density field can be recovered simply by inverse Fourier transform of the collected data. Of course, there are many phenomena which have not been accommodated in our simple imaging model, including inhomogeneities in the magnetic field. A significant problem in MRI imaging is that of correctly recovering the Fourier phase characteristic.

3   Algebraic Reconstruction Techniques

We may view each measurement sample, $r(\theta, s)$, as a single constraint on the density field, $d(x, y)$, where the constraint is that $d(x, y)$ must satisfy

$$r(\theta, s) = \int_{-\infty}^{\infty} d(t\cos\theta + s\sin\theta, -t\sin\theta + s\cos\theta)\, dt$$

In this way, the constraint defines a linear sub-space of the space of all possible two dimensional images, $d(x, y)$. Let $\mathbb{D}_{\theta,s}$ denote this sub-space. The image, $d(x, y)$, must lie in the intersection of all of these sub-spaces, which we may write as

$$d(x, y) \in \bigcap_{\theta,s} \mathbb{D}_{\theta,s}$$

Now observe that a linear sub-space is also a convex set, since linear combinations of any two elements in a linear space also belongs to the space (by definition). Thus, we may consider finding $d(x, y)$ by the method of projections onto convex sets (POCS), as studied in the topic entitled "Inverse Problems," and also summarized for your convenience in Appendix A.

To run the POCS algorithm, we have only to find the projection of a given image, $d^{(k)}(x, y)$, onto the sub-space, $\mathbb{D}_{\theta,s}$. We then iteratively project any starting image onto each of the $\mathbb{D}_{\theta,s}$ in turn until the algorithm converges (guaranteed) at which point we have an image which satisfies all of the constraints. The algorithm can only be meaningfully implemented when the points, (x, y), lie on a discrete grid. Each observed sample may then be written as a finite linear combination,

$$r(\theta, s) = \sum_{\mathbf{n}} d[\mathbf{n}]\, w\left(|n_1 \Delta x \sin\theta + n_2 \Delta y \cos\theta - s|\right) \tag{3}$$

Here, $|n_1 \Delta x \sin\theta + n_2 \Delta y \cos\theta - s|$ is the Euclidean distance between the line at location s and orientation θ and the point $(x, y) = (n_1 \Delta_x, n_2 \Delta_y)$. The one dimensional function, $w()$, is a kind of blurring function (a Line Spread Function or LSF) identifying the sensitivity of the sensor at position s to the the density field, $d(x, y)$, at varying distances from the line.

Note that equation (3) may be rewritten as an inner product,

$$r(\theta, s) = \sum_{\mathbf{n}} d[\mathbf{n}]\, w_{\theta,s}[\mathbf{n}] = \langle \mathbf{d}, \mathbf{w}_{\theta,s} \rangle$$

where

$$w_{\theta,s}\left[\mathbf{n}\right] = w\left(\left|n_1 \Delta x \sin\theta + n_2 \Delta y \cos\theta - s\right|\right)$$

The projection of an initial image, $d^{(k)}\left[\mathbf{n}\right]$, onto $\mathbb{D}_{\theta,s}$, may now be written as

$$d^{(k+1)}\left[\mathbf{n}\right] = d^{(k)}\left[\mathbf{n}\right] + \alpha w_{\theta,s}\left[\mathbf{n}\right]$$

where

$$
\begin{aligned}
\alpha &= \frac{r\left(\theta,s\right) - \left\langle \mathbf{d}^{(k)}, \mathbf{w}_{\theta,s}\right\rangle}{\left\|\mathbf{w}_{\theta,s}\right\|^2} \\
&= \frac{r\left(\theta,s\right) - \sum_{\mathbf{k}} d\left[\mathbf{k}\right] w_{\theta,s}\left[\mathbf{k}\right]}{\sum_{\mathbf{k}} w_{\theta,s}^2\left[\mathbf{k}\right]}
\end{aligned}
$$

It is easy to verify that this selection yields $\left\langle \mathbf{d}^{(k+1)}, \mathbf{w}_{\theta,s}\right\rangle = r\left(\theta,s\right)$.

The POCS approach described here is also known as an Algebraic Reconstruction Technique (ART). It is essentially a general iterative algorithm for solving a large set of linear equations with sparse coefficients.

Comparing with equation (5), it is easy to see how the POCS approach may be augmented to handle quantization errors in the $r\left(\theta,s\right)$ field or to add additional constraints to the problem. For example, it is easy to constrain the density field $d\left(x,y\right)$ to be strictly non-negative,

$$d\left(x,y\right) \geq 0$$

or to have a known region of support (location of the subject in the imaging system), as part of the iterative POCS algorithm. For more on this, see Section **??**.

A   Review of POCS

In this Appendix we review a useful approach to solving inverse problems in which the observations form hard constraints on the image we are trying to recover. We assume that noise in our image observations can take on any value in side the range

$$|\nu[\mathbf{n}]| \leq \sigma_\nu, \quad \forall \mathbf{n} \tag{4}$$

and we look for an ideal image, $y[\mathbf{n}]$, which satisfies the observation model

$$x[\mathbf{n}] = \sum_{\mathbf{k}} y[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] + v[\mathbf{n}]$$

Here, $x[\mathbf{n}]$ denotes the individual observations (usually samples recovered from an imperfect image sensor), $v[\mathbf{n}]$ represents noise in these observations, and $h_{\mathbf{n}}[\mathbf{k}]$ represents the relationship (assumed linear) between the ideal image at location $\mathbf{k}$ and observation $x[\mathbf{n}]$. In many cases, $h_{\mathbf{n}}[\mathbf{k}]$ represents LSI filtering so that $h_{\mathbf{n}}[\mathbf{k}] = h[\mathbf{n} - \mathbf{k}]$.

In view of the simple noise model in equation (4), we look for images which satisfy the constraints

$$\left| x[\mathbf{n}] - \sum_{\mathbf{k}} y[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right| \leq \sigma_\nu, \quad \forall \mathbf{n} \tag{5}$$

Each observed sample, $x[\mathbf{n}]$, induces a separate constraint equation. If our model is correct, there should be at least one image, $y[\mathbf{n}]$, which satisifes all the constraints simultaneously. Of course, in many situations there will be many images which satisfy all the constraints and so we will often impose additional constraints to narrow the space of possible solutions.

A.1   Underlying Principle

In this section we present the underlying principle behind the method of iterated projections onto convex sets, usually simply called "POCS". We begin by reviewing the concept of a convex set. Let $\mathbb{Y}$ be any set; in our case, it will denote the set of all images which satisfy some constraint. Let $\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{Y}$ denote any two elements of this set; in our case, these are two different images which both satisfy the relevant constraint. The set $\mathbb{Y}$ is called convex if $\mathbf{y} = t\mathbf{y}_1 + (1 - t)\mathbf{y}_2 \in \mathbb{Y}$ for all $0 \leq t \leq 1$, whenever

$\mathbf{y}_1, \mathbf{y}_2 \in \mathbb{Y}$. That is, all images formed by mixing the two images, $\mathbf{y}_1$ and $\mathbf{y}_2$ in any proportion must also satisfy the relevant constraint.

Let us consider the observation constraints of equation (5). In this case, each observed sample, $x[\mathbf{n}]$, generates a constraint set, $\mathbb{Y}[\mathbf{n}]$. That is

$$\mathbf{y} \in \mathbb{Y}[\mathbf{n}] \iff \left| x[\mathbf{n}] - \sum_{\mathbf{k}} y[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right| \leq \sigma_\nu \tag{6}$$

Let $\mathbf{y}_1$ and $\mathbf{y}_2$ be any two images which belong to $\mathbb{Y}[\mathbf{n}]$ and let $\mathbf{y}$ be any image on the convex interpolation of $\mathbf{y}_1$ and $\mathbf{y}_2$; that is, $y[\mathbf{k}] = t y_1[\mathbf{k}] + (1 - t) y_2[\mathbf{k}]$, $\forall \mathbf{k}$, with $0 \leq t \leq 1$. Then it is easy to see that $\mathbf{y} \in \mathbb{Y}$ also. For

$$\left| x[\mathbf{n}] - \sum_{\mathbf{k}} \left[t y_1[\mathbf{k}] + (1 - t) y_2[\mathbf{k}] \right] h_{\mathbf{n}}[\mathbf{k}] \right|$$
$$= \left| t \left\{ x[\mathbf{n}] - \sum_{\mathbf{k}} y_1[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right\} + (1 - t) \left\{ x[\mathbf{n}] - \sum_{\mathbf{k}} y_2[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right\} \right|$$
$$\leq t \left| x[\mathbf{n}] - \sum_{\mathbf{k}} y_1[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right| + (1 - t) \left| x[\mathbf{n}] - \sum_{\mathbf{k}} y_2[\mathbf{k}] h_{\mathbf{n}}[\mathbf{k}] \right|$$
$$\leq t \sigma_\nu + (1 - t) \sigma_\nu = \sigma_\nu$$

where we have used the triangle inequality. Thus $\mathbb{Y}[\mathbf{n}]$ is a convex set. The image we seek, $\mathbf{y}$, lies in the intersection of all these convex constraint sets.

POCS is a general method for finding an element in the intersection of a collection of convex sets. Let $\mathbb{Y}_1, \mathbb{Y}_2, \ldots, \mathbb{Y}_K$ be an enumeration of all the convex sets. In our case K would be the number of observed image samples. Let $\mathbf{y}^{(0)}$ be any initial image, not necessarily satisfying any of the constraint sets. The algorithm iteratively refines this initial image, setting

$$\mathbf{y}^{(k)} = \mathcal{P}(\mathbb{Y}_k, \mathbf{y}^{(k-1)}) \tag{7}$$

That is, the next estimate, $\mathbf{y}^{(k)}$, is obtained by *projecting* the most recent estimate, $\mathbf{y}^{(k-1)}$, onto the convex set, $\mathbb{Y}_k$. Once we have finished projecting onto all the convex sets, $\mathbb{Y}_k$, $k = 1, 2, \ldots, K$, we start all over again, repeatedly projecting onto each of the convex sets in turn. This process is guaranteed to converge. That is, we will eventually reach a point at which the image lies in

$$\mathbb{Y} = \bigcap_k \mathbb{Y}_k$$

provided this intersection is non-empty.

The projection operation expressed in equation (7) means that $\mathbf{y}^{(k)}$ is the closest element of $\mathbb{Y}_k$ to $\mathbf{y}^{(k-1)}$. That is,

$$\mathbf{y}^{(k)} = \mathcal{P}(\mathbb{Y}_k, \mathbf{y}^{(k-1)}) = \operatorname*{argmin}_{\mathbf{y} \in \mathbb{Y}_k} \left\| \mathbf{y} - \mathbf{y}^{(k-1)} \right\|$$

Often, we will measure distance using the familiar Euclidean metric. In general, however, any valid distance metric may be used provided the same metric is used to perform the projection onto EVERY convex set. Of course, the projection operation has no effect on elements which already belong to the relevant set. Thus, each projection operator is *idempotent*.

A.2   Typical Constraint Sets

A.2.1   "Quantized" Observations

In this section we consider constraints which arise from the observation model in the manner identified in equation (6). As mentioned already, this type of constraint can be an appropriate model for a set of quantized observations. For notational simplicity, we rewrite this constraint in inner product form as

$$\mathbf{y} \in \mathbb{Y} \Longleftrightarrow |x\,[\mathbf{n}] - \langle \mathbf{y}, \mathbf{h_n} \rangle| \leq \sigma_\nu$$

We have already seen that this set is convex. Now, write

$$\mathbf{y}^{(1)} = \mathcal{P}(\mathbb{Y}, \mathbf{y}^{(0)})$$

for the projection of an initial vector, $\mathbf{y}^{(0)}$, onto the set, $\mathbb{Y}$. Then

$$\mathbf{y}^{(1)} = \mathbf{y}^{(0)} + \boldsymbol{\delta}$$

where $\boldsymbol{\delta}$ is the smallest displacement such that $\mathbf{y}^{(1)} \in \mathbb{Y}$. We will work with the conventional Euclidean norm,

$$\|\boldsymbol{\delta}\| = \sqrt{\sum_{\mathbf{k}} \delta^2\,[\mathbf{k}]}$$

In this case, we must have

$$\boldsymbol{\delta} = \alpha \mathbf{h_n}$$

for some real number, α. This is because vectors of this form produce the largest change in $|x\,[\mathbf{n}] - \langle \mathbf{y}, \mathbf{h_n} \rangle|$, per unit length. So we want to find the smallest value of α such that

$$
\begin{aligned}
\sigma_\nu &\geq \left| x\,[\mathbf{n}] - \left(\left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle + \alpha \left\langle \mathbf{h_n}, \mathbf{h_n} \right\rangle \right) \right| \\
&= \left| \left(x\,[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle \right) - \alpha \left\| \mathbf{h_n} \right\|^2 \right|
\end{aligned}
$$

The solution to this depends upon whether $x\,[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle$ is less than $-\sigma_\nu$, greater than σ_ν, or in between these two bounds. In particular, we have

$$
\alpha = \begin{cases} \dfrac{\left(x[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle \right) - \sigma_\nu}{\left\| \mathbf{h_n} \right\|^2} & \text{if } x\,[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle > \sigma_\nu, \\[2ex] \dfrac{\left(x[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle \right) + \sigma_\nu}{\left\| \mathbf{h_n} \right\|^2} & \text{if } \left(x\,[\mathbf{n}] - \left\langle \mathbf{y}^{(0)}, \mathbf{h_n} \right\rangle \right) < -\sigma_\nu, \\[2ex] 0 & \text{otherwise.} \end{cases} \tag{8}
$$

Thus, the procedure for projecting an initial image $\mathbf{y}^{(0)}$ onto the constraint set $\mathbb{Y}[\mathbf{n}]$ is as follows:

Step 1: Compute the error,

$$
\epsilon = x[\mathbf{n}] - \sum_{\mathbf{k}} y^{(0)}[\mathbf{k}] h_{\mathbf{n}}\,[\mathbf{k}]
$$

Step 2: Compute the value of $\left\| \mathbf{h_n} \right\|^2 = \sum_{\mathbf{k}} h_{\mathbf{n}}^2\,[\mathbf{k}]$, noting that this might be different for each value of $\mathbf{n}$.

Step 3: Compute the value of α from

$$
\alpha = \begin{cases} \dfrac{\epsilon - \sigma_\nu}{\left\| \mathbf{h_n} \right\|^2} & \text{if } \epsilon > \sigma_\nu, \\[2ex] \dfrac{\epsilon + \sigma_\nu}{\left\| \mathbf{h_n} \right\|^2} & \text{if } \epsilon < -\sigma_\nu, \\[2ex] 0 & \text{otherwise.} \end{cases}
$$

Step 4: Set $y^{(1)}[\mathbf{k}] = y^{(0)}[\mathbf{k}] + \alpha h_{\mathbf{n}}\,[\mathbf{k}]$.

In a practical application, $\left\| \mathbf{h_n} \right\|^2$ usually takes on one of a small number of different values, which can be computed ahead of time so that most of the computational effort for any given projection is due to the computation of ϵ.

When the constraint sets are all as above and the "noise" amplitude, $\sigma_\nu = 0$, the method of iterated projections onto convex sets is identical to the so-called "algebraic reconstruction" method for solving large sets of linear equations.

A.2.2   Band-limited Constraint

As mentioned earlier, it is sometimes necessary to add additional constraints on the class of images which can be reconstructed, beyond the constraints offered by the observation model. This frequently happens when there are too few observations so that the intersection of the observation constraint sets is too large. After all, the POCS algorithm is only guaranteed to find a single member of the intersection of the constraint sets and this is not guaranteed to be a meaningful image if the constraints are very sloppy. One common additional constraint is that of band-limitedness. That is, we might assume that the ideal image, $\mathbf{y}[\mathbf{n}]$, has the property that its Fourier transform, $\hat{y}(\boldsymbol{\omega})$, is zero outside some region, $\mathcal{R}_{\hat{y}} \subset (-\pi, \pi)^2$. It is easy to see that the corresponding constraint set is convex. Moreover, the projection operator is easy to deduce. Specifically, let $y^{(0)}[\mathbf{n}]$ be the initial image and let $y^{(1)}[\mathbf{n}]$ be its projection onto the set of band-limited images. Then writing

$$y^{(1)}[\mathbf{n}] = y^{(0)}[\mathbf{n}] - \delta[\mathbf{n}]$$

we have

$$\hat{y}^{(1)}(\boldsymbol{\omega}) = \hat{y}^{(0)}(\boldsymbol{\omega}) - \hat{\delta}(\boldsymbol{\omega}) = 0, \quad \text{for } \boldsymbol{\omega} \notin \mathcal{R}_{\hat{y}}.$$

Thus,

$$\hat{\delta}(\boldsymbol{\omega}) = \hat{y}^{(0)}(\boldsymbol{\omega}), \quad \text{for } \boldsymbol{\omega} \notin \mathcal{R}_{\hat{y}}. \tag{9}$$

Now we want to find the smallest displacement image, $\delta[\mathbf{n}]$, which satisfies equation (9), where size is measured with respect to the Euclidean distance metric, since this is the one we used to perform the projections associated with the observation constraint sets in Section **??** and we must use the same distance metric for all constraint sets in the POCS method. So

$$\|\boldsymbol{\delta}\|^2 = \sum_{\mathbf{n}} \delta^2[\mathbf{n}] = \frac{1}{4\pi^2} \int_{-\pi}^{\pi} d\omega_1 \cdot \int_{-\pi}^{\pi} d\omega_2 \cdot \left| \hat{\delta}(\boldsymbol{\omega}) \right|^2$$

by Parseval's theorem, which means that the smallest displacement image satisfying equation (9) has $\hat{\delta}(\boldsymbol{\omega}) = 0$, $\forall \boldsymbol{\omega} \in \mathcal{R}_{\hat{y}}$. Thus,

$$\hat{y}^{(1)}(\boldsymbol{\omega}) = \begin{cases} \hat{y}^{(0)}(\boldsymbol{\omega}) & \text{for } \boldsymbol{\omega} \in \mathcal{R}_{\hat{y}}, \\ 0 & \text{otherwise.} \end{cases}$$

To summarize, to project an image onto the convex set of band-limited functions one has only to take the Fourier transform (DSFT), set all terms

outside the relevant region of support, $\mathcal{R}_{\hat{y}}$, to zero, and then take the inverse Fourier transform. In practice, one generally performs this with the DFT using the fast FFT algorithm.

It should be noted that unlike the observation constraints, imposing a constraint such as band-limitedness on the image we are reconstructing may represent an inaccurate assumption concerning the actual image which generated the observations. Consequently there is no longer any guarantee that all the constraints can simultaneously be satisfied, that is, the intersection of all the constraint sets might be empty. When this happens, the POCS iterations will not converge. It is possible to detect lack of convergence during the POCS procedure and optionally modify one or more of the constraints to enlarge the corresponding constraint set. For example, we might gradually enlarge a circularly symmetric or rectangular band-limited region of support, $\mathcal{R}_{\hat{y}}$, until convergence is achieved, thereby recovering the most band-limited image which is consistent with all of the observations.

A.2.3   Spatial Region of Support Constraint

Suppose we know that $y[\mathbf{n}] = 0$ outside some region of support, $\mathcal{R}_y$. This constraint is similar to the bandlimited constraint, except that it applies in the image domain, rather than the Fourier domain. Accordingly, the projection operator involves setting

$$y^{(1)}[\mathbf{n}] = \begin{cases} y^{(0)}[\mathbf{n}] & \text{if } \mathbf{n} \in \mathcal{R}_y \\ 0 & \text{otherwise} \end{cases}$$

A.2.4   Boundedness Constraint

Suppose we know that $y[\mathbf{n}]$ must be strictly bounded according to

$$L \leq y[\mathbf{n}] \leq U, \quad \forall \mathbf{n}$$

It is easy to see that the corresponding constraint set is convex. Again, the projection operator is easy to deduce. Letting $y^{(0)}[\mathbf{n}]$ be the initial image and $y^{(1)}[\mathbf{n}]$ its projection onto the set of bounded images, we can write the projection operator as

$$y^{(1)}[\mathbf{n}] = \begin{cases} U & \text{if } y^{(0)}[\mathbf{n}] \geq U \\ L & \text{if } y^{(0)}[\mathbf{n}] \leq L \\ y^{(0)}[\mathbf{n}] & \text{otherwise} \end{cases}$$