The Layman's Guide to Tomographic Image Reconstruction: A Gentle Introduction

Javier A. Salazar II

December 19, 2017

Contents

1 Introduction

Although this paper mainly focuses on X-ray computed tomography (CT) scans, there is additional information in the broken ray radon chapter (chapter 8) on current research topics involving other forms of tomography scans. The chapter thereafter (chapter 9) focuses on localized tomography which is also a hot topic in modern research.

Why X-ray? Why not infrared or visible light? The problem lies in the balance of contrast and safety. Due to physical nature of high frequency photons, photons pass through the material without difficulty. One downside of X-ray (high frequency) is the radiation patients receive when undergoing scanning even though the resolution is high. In comparison, ultrasound is a safe imaging technology but will lead to poorer pictures.

Before we can understand what goes on mathematically to reconstruct the image, we must understand the physical workings on one of the beams. When a beam passes through matter, such as the brain or a bone, some of the X-ray photons are absorbed by the material. How much is determined based on the electron density of the material itself; something like bone has high Calcium content so more electrons than Helium. The change in intensity of the photons will determine how well or how badly the material can absorb the photons. If the human brain was homogeneous in nature, then we would not be having this discussion on tomography reconstruction since the changes in intensity from the source to the detector is all we would need. Fortunately for us, the problem is much more difficult.

1.1 Beer's law and The Attenuation Function

To create a model that will allow us to reconstruct the image, we will need to make the following assumptions: 1. All of the X-ray beams are monochromatic. In other words, each photon has the same energy level E with constant frequency along with the same number of photons passing per second through the path of the beam denoted as N(x). 2. We assume the beam has zero width and will not be refracted while inside a patient. It should be noted that, in the real world, these assumptions will not hold for certain beams but will still provide us with an approximation model to reconstruct the image. Using this information, the intensity of the beam at x is

$$I(x) = E \cdot N(x) \tag{1}$$

Every material will absorb photons differently and the ratio of photons that do pass through is called the attenuation coefficient of the material. Something like bone has a high attenuation meaning it absorbs a lot of the photons. We will use this attenuation coefficient symbolized as A(x) to describe the relationship between the change of intensity and the attenuation of the material. This is expressed by the differential equation described as Beer's law shown below

$$\frac{dI}{dx} = -A(x) \cdot I(x) \tag{2}$$

When the X-ray is taken, only the initial and final values of I(x) are known. Using this and solving the separable differential equation yield the following

$$\int_{x_0}^{x_1} A(x)dx = ln(\frac{I_0}{I_1}) \tag{3}$$

At first sight, this might be a confusing result to understand. What this equation means is that the X-ray will only give us the integral of A(x) even though we need A(x) to actually reconstruct the image. The reader might still be left confused on how the beams and the attenuation function are related so the picture below will be explained.

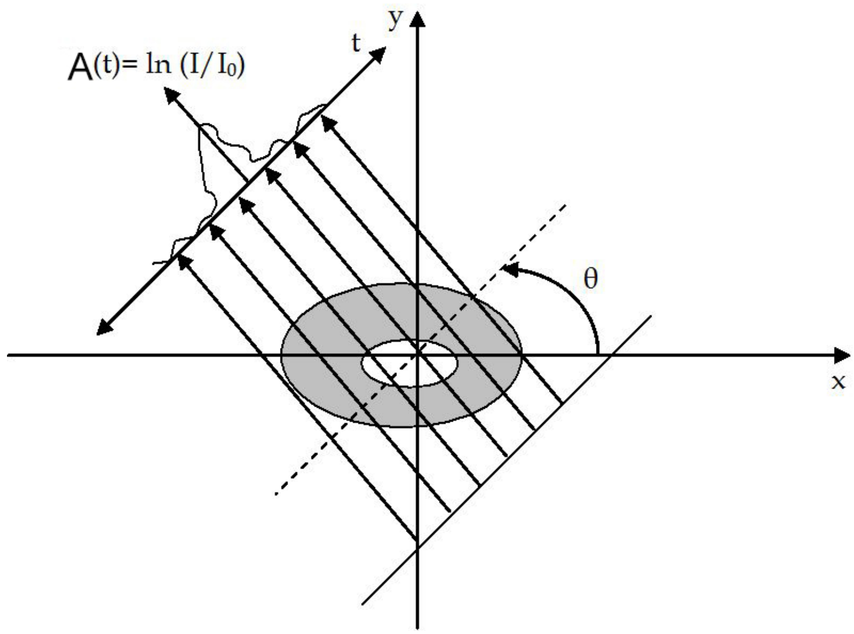

Figure 1: A “virtual patient” undergoing a CT scan with A(x) displayed.

The gray and white ellipse shapes shown represent a “virtual patient” where the different colors represent a different attenuation i.e. different materials in the “skull” of a patient. The rays are set by a specific angle illustrated as θ. The attenuation values across space are graphed as A(t). Note that the values themselves are integrals since photons are absorbed as they go through the different materials. We can design the X-ray machine that can measure I_0 and I_1 so that we can find the integral of A(x), but we wish to know the value of A(x) for every location. The question that this chapter poses becomes this, can we figure out A(x) if we only know the average value of A along the X-ray beam lines that pass through a specific region of the image? (The answer of course is yes otherwise we would be done here)

The basic premise of CT scanning is the following: If we have enough X-rays that pass through the patient with different angles and distances all evenly spaced, then we can reconstruct the image.

1.2 Lines, Lines, and More Lines

There are many ways to describe lines in the spatial domain, however, CT scans use polar geometry to describe the image. Later on, it will be understood why this decision was made to describe lines using polar instead of Cartesian form. Lines are normally represented as $y = m \cdot x + b$ but are unable to properly describe vertical lines. Instead, we will use the variable t and θ. T is the distance away from the origin and θ is the counter clockwise angle from the positive x axis. We know from algebra that $x = t \cdot cos(\theta)$ and $y = t \cdot sin(\theta)$. Taking this a step further, we will call a line perpendicular to this t line as $l_{t,\theta}$ as this is the actual X-ray that will pass through the body. A representation is shown below. Further, we discuss the characterization of the line into the t and s parameters.

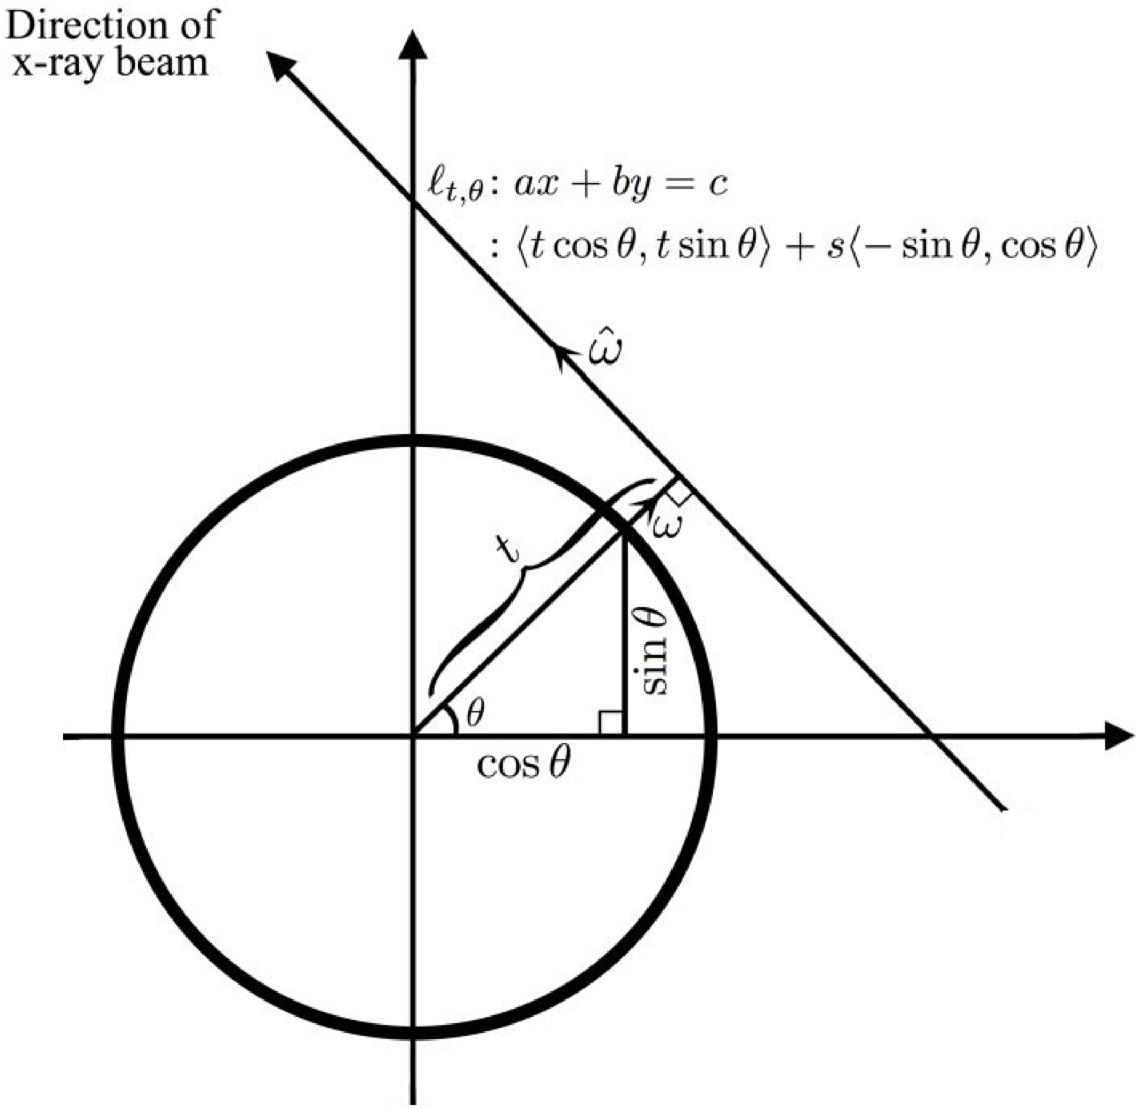

Figure 2: An X-Ray going through empty space with the t and θ values shown and the parameterization equation of the line.

The first term in the equation in Figure 2 tells a point how much to walk in the x and y directions from the origin to the line where the intersection occurs. The second term "s" describes how far along the line to walk. If s=1 then a point at the intersection would walk in the positive s direction for 1 unit. If we take all real values of s, then we have essentially drawn the entire line in the plane. Due to polar nature, we can tell that a line t at 30 degrees is the same as line -t at -150 degrees or line t at 390 degrees. In other words, $l_{t,\theta+2\pi} = l_{t,\theta}$ and $l_{t,\theta+\pi} = l_{-t,\theta}$. To restate the important discovery of this section that will prove useful in chapter 2:

$$l_{t,\theta} = t \cdot < cos(\theta), sin(\theta) > + s \cdot < -sin(\theta), cos(\theta) > \tag{4}$$

2 The Radon Transform

2.1 Phantom...of the opera?

Before we can simulate X-ray scanning, we must design a purely mathematical symbol known as the phantom. A phantom is a "virtual patient" that we will use to run through the image reconstruction process. The problem with using real patient data is we have no idea what it is "supposed" to look like. How do we know we have successfully created the algorithm? If the image doesn't look as expected, is that fault due to the algorithm or a patient slightly moving while collecting data? We use the phantom as a way to get around the issue to compare the reconstructed image with the actual image. There exists a famous phantom very commonly used to test the robustness of image reconstruction algorithms known as the Shepp-Logan phantom. An image of the famous phantom is shown below. Note: The phantom represents the head of a skull hence the white oval holding the other objects inside.

Figure 3: The Shepp-Logan phantom - The standard phantom in the world of image reconstruction.

2.2 A Very Bored Johann Radon

The Radon transform was introduced in 1917 by Johann Radon who also created the inverse Radon transform which we will discuss later on. As shown below in Definition 1, the lines projected by $l_{t,\theta}$ are integrated. Thus, the Radon transform takes the different values located along the line and integrates them; this process is repeated for all the other lines. The f function described can be any two-dimensional function such as the Shepp-Logan phantom shown in Figure 3. Of course, since the Shepp-Logan function is compactly supported (meaning it doesn't go on forever), then the integration variable s only goes from one end of the image to the other end depending on the line.

Definition 1. For an image defined in $\mathbb{R}^2$ for which f describes, let the Radon transform of f be defined in terms of real values t and θ as

$$Rf(t,\theta) := \int_{s=-\infty}^{s=\infty} f((t\cos(\theta) - s\sin(\theta)), t\sin(\theta) + s\cos(\theta))ds \tag{5}$$

The Radon transform defined above is the mathematical simulation of an X-Ray running through a phantom. We proceed by taking the line integral value and plotting it in a transformed coordinate system. Since the parameters are t and θ, we use the t value as the y axis and the θ values as the x axis. For example, a 256x256 pixel image will require t to be from -91 to 91 pixels from the origin to completely cover the image. As a result, the radon image has different dimensions than the original image. The Radon image of the Shepp-Logan phantom is illustrated below.

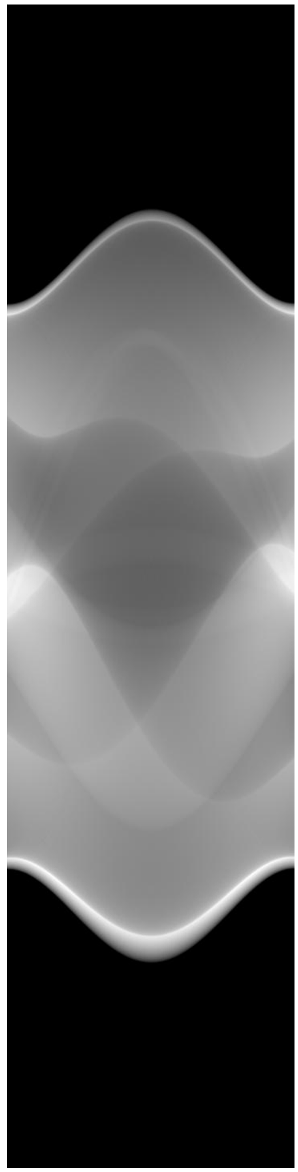

Figure 4: The Radon image of the Shepp-Logan phantom.

Figure 4 shows the data from $0°$ to $360°$ so that is why there is symmetry in the picture. We really only need the left or right side of the full image to reconstruct the phantom. Because the lines go from -t to t, we can transform the negative values of t into positive values with a degree shift of $\theta_2 = \theta_1 + 180°$ and thus all of the values are covered. Likewise, from another viewpoint, we only need the top half or bottom half of the image to completely reconstruct the image.

2.3 Properties of the Radon Transform

Given the Radon transform is a linear system, the following properties are as follows:

$$R(\alpha f + \alpha g) = \alpha Rf + \alpha Rg \tag{6}$$

Proposition 1. *The function f is defined in the plane, a and b are arbitrary real numbers, and c is a positive real number. The function g is a translated function of f in the form* $g(x, y) = f(x-a, y-b)$ *and the scaled function of f in the form of* $Rh(t, \theta) = (\frac{1}{c}) \cdot Rf(ct, \theta)$*. Then for a real t and* θ *:*

$$Rg(t, \theta) = Rf(t - a\cos(\theta) - b\sin(\theta), \theta) \tag{7}$$

$$Rh(t, \theta) = (\frac{1}{c}) \cdot Rf(ct, \theta) \tag{8}$$

For a function f defined in the plane a real angle ϕ, for a real x and y, function g is a rotated function f by $g(x, y) = f(x\cos(\phi) + y\sin(\phi), -x\sin(\phi) + y\cos(\phi)$ then the Radon is:

$$Rg(t, \theta) = Rf(t, \theta - \phi) \tag{9}$$

3 Unfiltered Back Projection

With only Rf known, we will need to recover the function f. Let's select a point in the original image f that we wish to recover and call it x_0 and y_0. This point has many radon lines that go through it, each with different t and θ values. We can define the value of t based on the θ value as:

$$t = x_0\cos(\theta) + y_0\sin(\theta) \tag{10}$$

The illustration below shows this relationship:

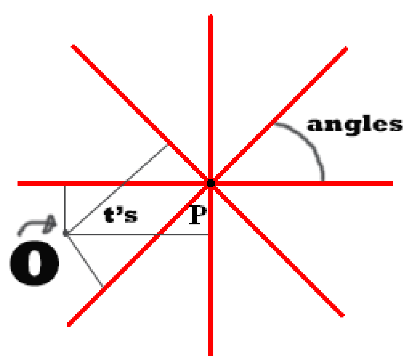

Figure 5: The radon lines of a point P in the function f with respect to origin O.

Definition 2. Using Eq. 9 and using a function $h = h(t, \theta)$ whose inputs are polar coordinates (i.e. the radon function of the image function f), let the back projection of h and the point (x,y) be defined by:

$$Bh(x, y) := \frac{1}{\pi}\int_{\theta=0}^{\theta=\pi} h(x\cos(\theta) + y\sin(\theta)), \theta)d\theta \tag{11}$$

The reader may be wondering why θ is evaluated from 0 to π and not the complete circle. Remember that Radon line at say 45 degrees is the same as the line at 225 degrees and thus is completly unnecessary to go around the whole circle. I'm sure the patient would also appreciate not being exposed to twice the radiation required as well. The $\frac{1}{\pi}$ term is there since we are taking an average value of all of the lines that go through the point at (x, y). Eg. 9 substitution was made for t within the h() function to perform what is visually described in Figure 5. Notice that the image function f is in Cartesian coordinates but the radon function h is in polar coordinates. The back projection of h goes back to Cartesian as expected since it represents the function f. Back projection, since it is simply an integration like Radon, also is a linear system. If we were to take Figure 4 (the Radon image of the Shepp-Logan phantom) and use the back projection formula we would get the following image:

Figure 6: The unfiltered back projection of the Radon Shepp-Logan phantom.

Note that Figure 6 is the unfiltered back projection. Unfiltered means the Radon image was plugged into the back projection formula without any modifications. In order to properly reconstruct the image, we must perform an operation to the Radon image before we use back projection. This operation will "fix" the image so that the blurry image will sharpen. If this operation is done, then we call it the filtered back projection because a filter was applied to the Radon image. To further demonstrate what back projection does, the image below graphically demonstrates the relationship between the spatial image, the radon functions, and the back projection image.

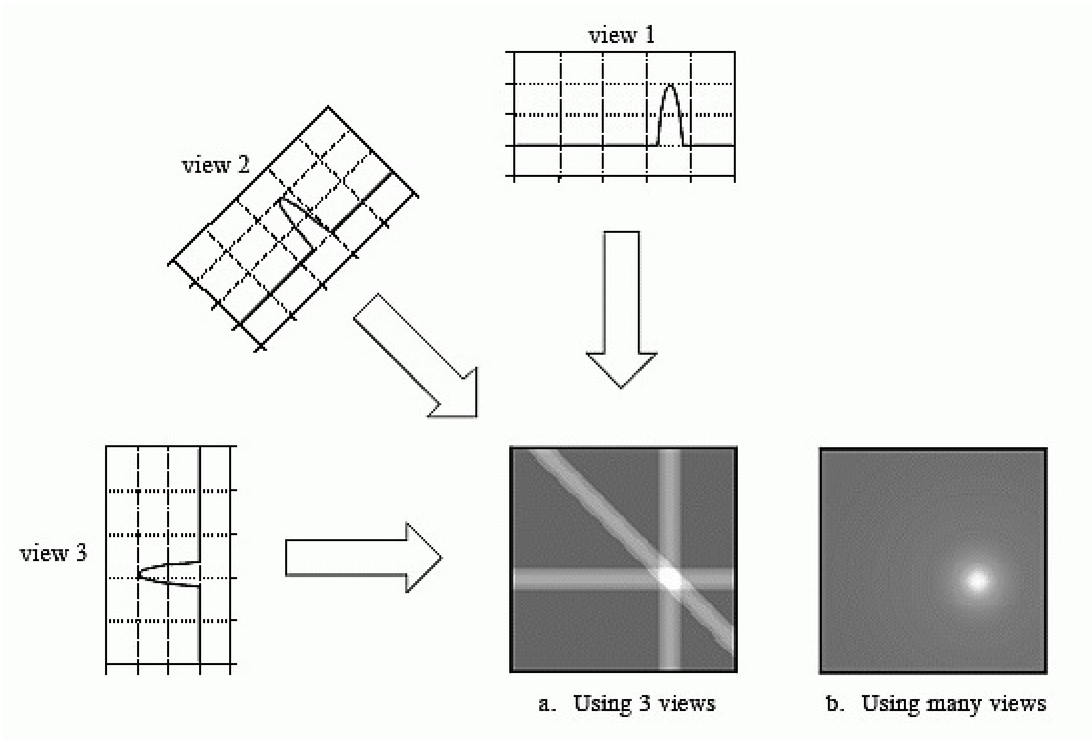

Figure 7: The visual representation of the back projection formula.

The three views represent different Radon slices of the original image. All of them appear the same since the image is only a circle, however the three views have different θ values. Note that the three views are smeared along the image as shown in section a. If enough views are used, then the image becomes section b.

3.1 Missing Data

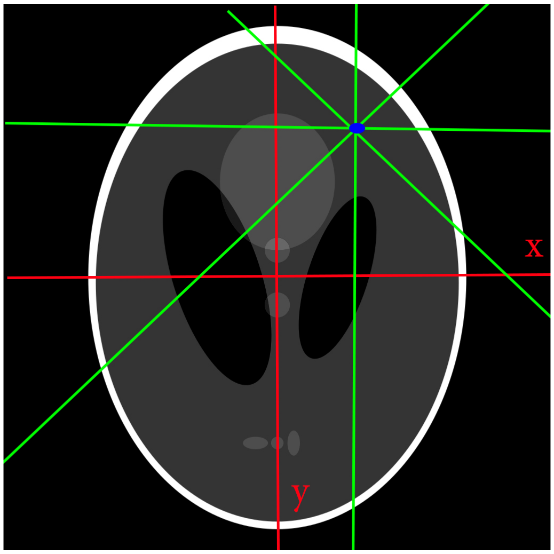

(a) Point undergoing reconstruction with relevant X-ray lines.

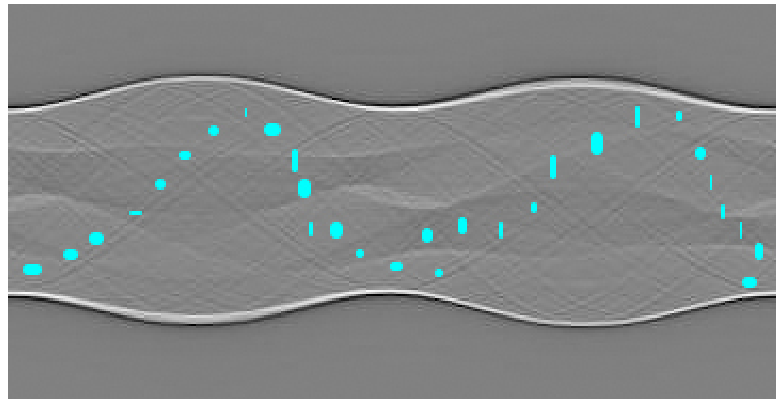

(b) Corresponding line parameters in Radon domain.

From the figure above, (a) shows a blue point that has not been reconstructed yet. We wish to get the value at this location and so the green lines correspond to X-rays taken with different t and θ values during the imaging process. We know that lines in the spatial domain are related to points in the Radon domain. From the sinogram at (b), the blue points are some of the lines that go through a point in the spatial domain. The points are arranged in such a way that it appears "like" a sine wave but not exactly. This interesting result is the reason why the X-ray image is called a sinogram despite the incorrect shape. Imagine that the phantom went through the Radon transform with an angle step size of 1. That means there are 180 lines going through each point in the image (keeping Radon resolution and backprojection resolution the same). All of those lines will correspond to different angle and t values. The angle values, due to same step size from Radon to backprojection, are all located perfectly in the sinogram. There will not be any values in between. There won't be a 90.5 degree line that goes through a point due to the step size chosen. Therefore, all of the data exists for the horizontal axis in the sinogram. However, because there is a step distance between the rays, there may not exist a value for certain locations. Some of the lines that go through a point in the phantom will not necessarily exist in the sinogram. When dealing with discrete data, it is always a possibility that the value needed is in between two other values. This is solved with interpolation.

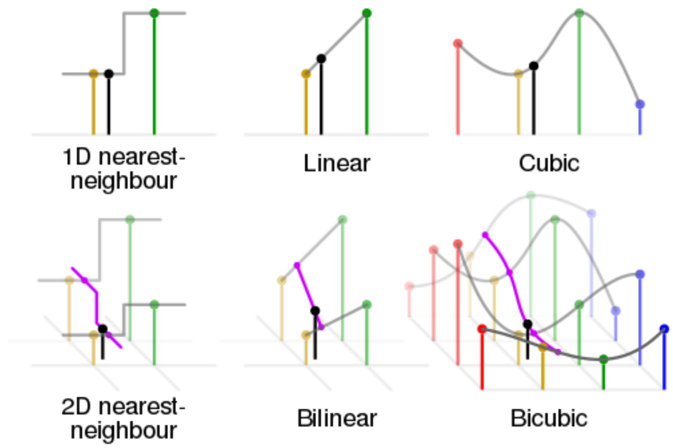

Figure 9: Different methods for interpolation.

Interpolation deals with discrete data by fitting a line or curve into the data and taking a value that lies between known data. From the picture shown above, we can see three important methods for interpolation: nearest neighbor, linear, and cubic. The yellow and green points are known data points and the black point in the middle is a value we wish to recover. The nearest neighbor method works by simply taking the value closest to the known data value. In the special case where the point lies in between, the higher data point is chosen. The cubic method works by fitting points along a cubic curve. Notice that there is also a blue and red point in the image. This is due to the fact that four points are required to completely characterize a cubic function. A line requires two points to fit, whereas a polynomial of degree two requires three points minimum. A cubic function will need the four points. The linear method fits a line between two points by finding the slope and then the y value is found from this. In other words, a weighted average value is gathered from the data and can be described by:

$$f(x) = (\frac{f(x_{x+1}) - f(x_k)}{x_{k+1} - x_k}) \cdot (x - x_k) + f(x_k) \tag{12}$$

where x is the value we wish to get that lies between x_k and x_{k+1}. It is important to note that our problem is one dimensional in nature due to missing t-values only. If the angle step size from the Radon transform and the backprojection algorithm do not match, then there will be angles for which there is no value in the sinogram. Thus, a 1D problem will turn into a 2D problem and 2D interpolation will be required. The images in the bottom row from the figure above show this process.

4 The Fourier Domain

4.1 A Brief Overview

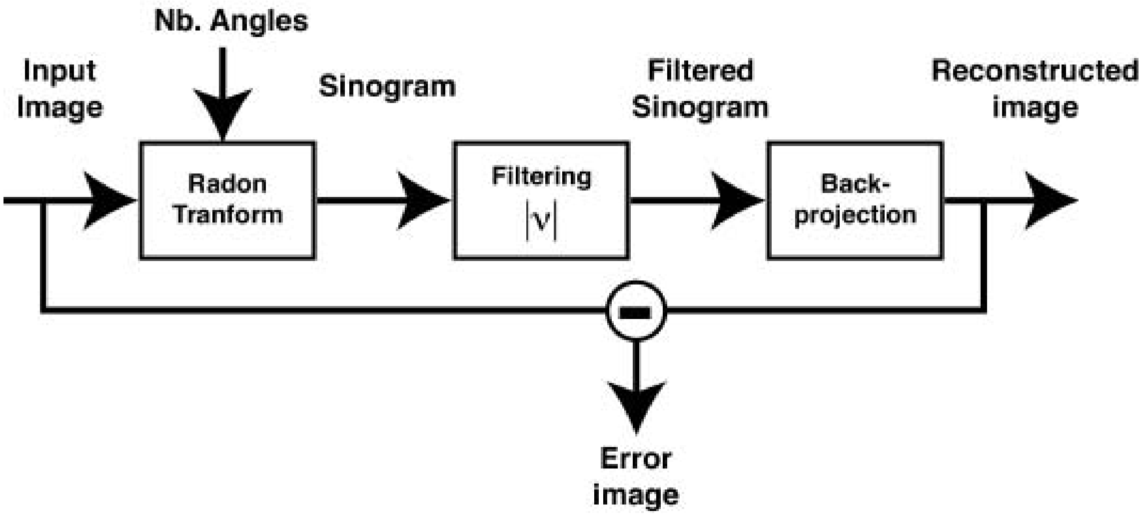

Figure 10: The complete tomography scanning progress.

As seen above, we have an input image that we wish to reconstruct. We feed the image and the number of angles to perform the line integrals. The more angles leads to a better resolution up to a certain extent. The radon image is also known as a sinogram due to its sine-looking graph. However, the radon image does not correlate to the sine function, the image just appears to resemble a sine function. The sinogram goes through a filtering process to give a filtered sinogram. The filtered image has the same structure as the unfiltered image: the amplitude values are changed. The final step is to perform back projection on the filtered sinogram to reconstruct the image. Afterwards, a simple error algorithm can compare the reconstructed image to the original image and display the areas that achieved good reconstruction as well as show areas that did poorly in the reconstruction. This chapter will focus on the filtering block. To filter the sinogram, we must take the image and transform it to the frequency domain (known as Fourier domain) and perform a multiplication. Once that is done, the image is converted back to the spatial domain using inverse Fourier transform. This next section will go through this process.

4.2 Joseph Fourier and His Radical Idea

Definition 3. For an absolutely integrable function f, that is $\int_{\infty}^{\infty} |f(x)|dx < \infty$, the Fourier transform of f is defined by ω such that:

$$Ff(\omega) := \int_{-\infty}^{\infty} f(x)e^{-i\omega x}dx \tag{13}$$

What does this mean? This means that a function with respect to time is converted to a function with respect to frequency. Imagine an audio signal of a bird chirping mixed with the sound of a horn. These sounds have very distinct sounds, but they can't be seperated when desling with the time domain. The two seperate signals superimpose each other and add up so it is impossible to see what signal belongs to the bird and which one belongs to the honk. However, all signals and functions (with the condition specified in Definition 3) can be decomposed to sine and cosine components. As long as the signal is compactly supported, meaning it doesn't go off to infinity, then we can use Fourier transform. Let the picture below demonstrate this concept.

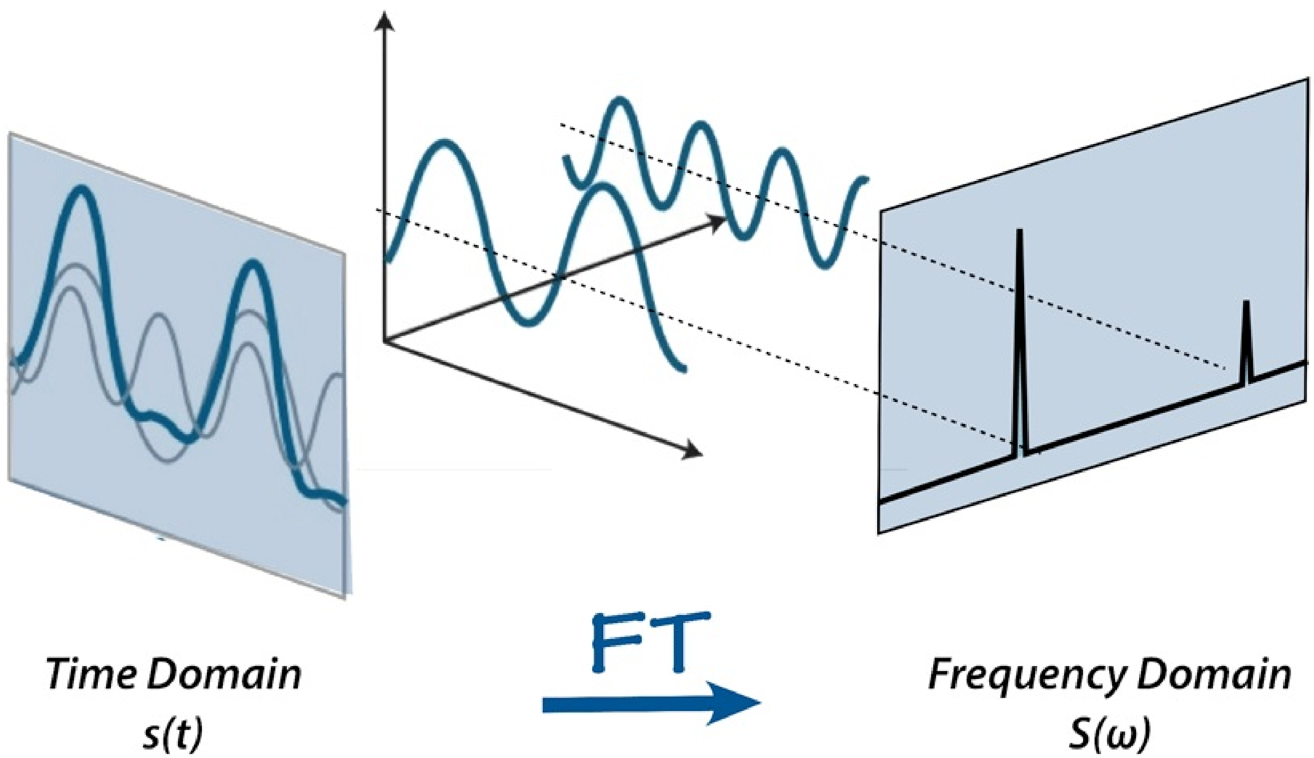

Figure 11: A time-domain signal and its respective frequency-domain equivalent.

We can see that the time-domain picture consists of the two gray sine waves that superimpose into the blue sine wave. The middle section shows the separation of the two sine waves. Looking at the graph from the front we see the normal signal, but as we turn and look at it from the side, we see the amplitudes of the signal and where they are located on the frequency axis. This specific signal has a lot of low frequency while only having a little high frequency components. If this was the car horn and the bird chirping, we can destroy that low frequency value and keep only the high frequency signal which is the bird chirp. After that, we can take the signal back to the time-domain effectively removing the car horn from the audio file. This process of messing with the frequency components is known as filtering. There is a similar process with medical reconstruction. Before going to the filtering process, I would like to mention the inverse Fourier transform. This transformation takes a frequency-domain signal and converts it back to the time-domain. It is defined as follows:

Definition 4. For a frequency-domain function g that is absolutely integrable, the inverse Fourier transform is defined for real x as follows:

$$F^{-1}g(x) := \frac{1}{2\pi}\int_{\omega=-\infty}^{\infty} g(\omega)e^{i\omega x}d\omega \tag{14}$$

Now that we have defined the transform definitions, it's time to show the filter process. This is the highlight of the chapter; the essential piece to properly reconstruct the signal.

4.3 Filters

For reasons that are difficult to explain, the correct filter to reconstruct the image is the absolute value of the frequency. This filter is shown in the picture below.

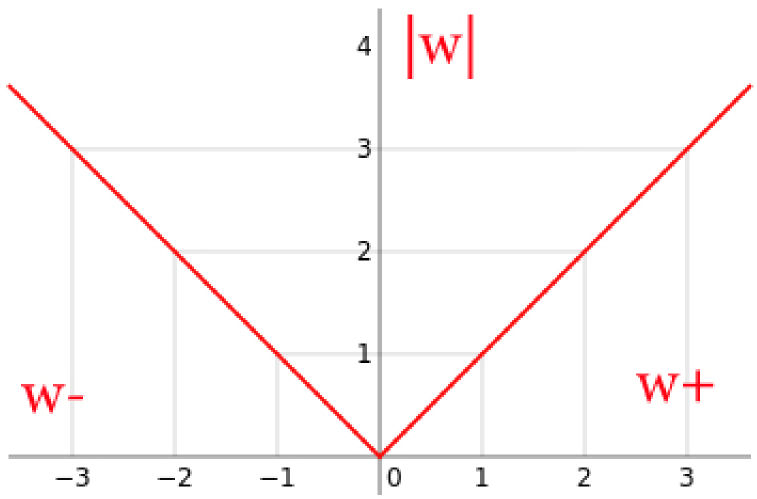

Figure 12: A visual representation of the $|w|$ filter .

This specific filter is what we call a high-pass filter. The high frequencies pass through with little trouble but the frequencies close to zero are de-emphasized since the filter value at that range is near zero. In image processing, this leads to a sharpening effect in images. The dual is the low-pass filter that allows low frequencies to pass and high frequencies to be blocked. This leads to a blurring effect as opposed to a sharpening effect. There are some filters that mix both concepts together and this is called a band-pass filter. Once the Fourier transform with respect to t is taken, we multiply the values by absolute omega, and transform it back with respect to t. The reader may seem confused by what is meant by what is meant by multiplying by the current frequency so that will be explained here. Once the image values are transformed to frequency values, what determines the x-domain frequency location of those values? The answer is Nyquist theorem. Sampling is the reduction of a continuous signal to a discrete signal. Sampling rate is the frequency of how many samples we take for second. Something like an audio wave is an analog signal but to convert it to discrete data, for say a digital speaker, requires that we take enough points of the signal to completely recreate the signal and thus no loss in information. An MP3 file has a sampling rate of 44,100 Hz so that's how many points it takes per second. If we have a 100 Hz sine wave then, per Nyquist theorem, we will need to sample the sine wave at least twice the frequency (200 Hz) to convert the continuous function to a discrete function. Likewise, the sinogram also carries a "sampling rate" but unlike an MP3, it is determined by distance instead of time. Let the sampling rate be defined as:

$$F_s = \left|\frac{2}{t_{n+1} - t_n}\right| \tag{15}$$

Where $t_{n+1} - t_n$ is the step size (distance) between two parallel rays in sequential order. Per Nyquist definition explained above, the domain of our signal becomes $[\frac{-F_s}{2}, \frac{F_s}{2}]$ with the number of points in between these two values matching the number of points of t values. With this said, the absolute values of these frequency points are multiplied by the amplitude value at that location to get the filtered sinogram after using inverse Fourier transform. The picture below shows what this filter process does to the sinogram originally shown in Figure 4.

Figure 13: Filtered sinogram of the Shepp-Logan phantom.

Applying back projection to the filtered sinogram above yields the following image below. Notice the change from Figure 6 which shows the unfiltered back projection. Also, compare the image with Figure 3 which is the original image that we are supposed to get.

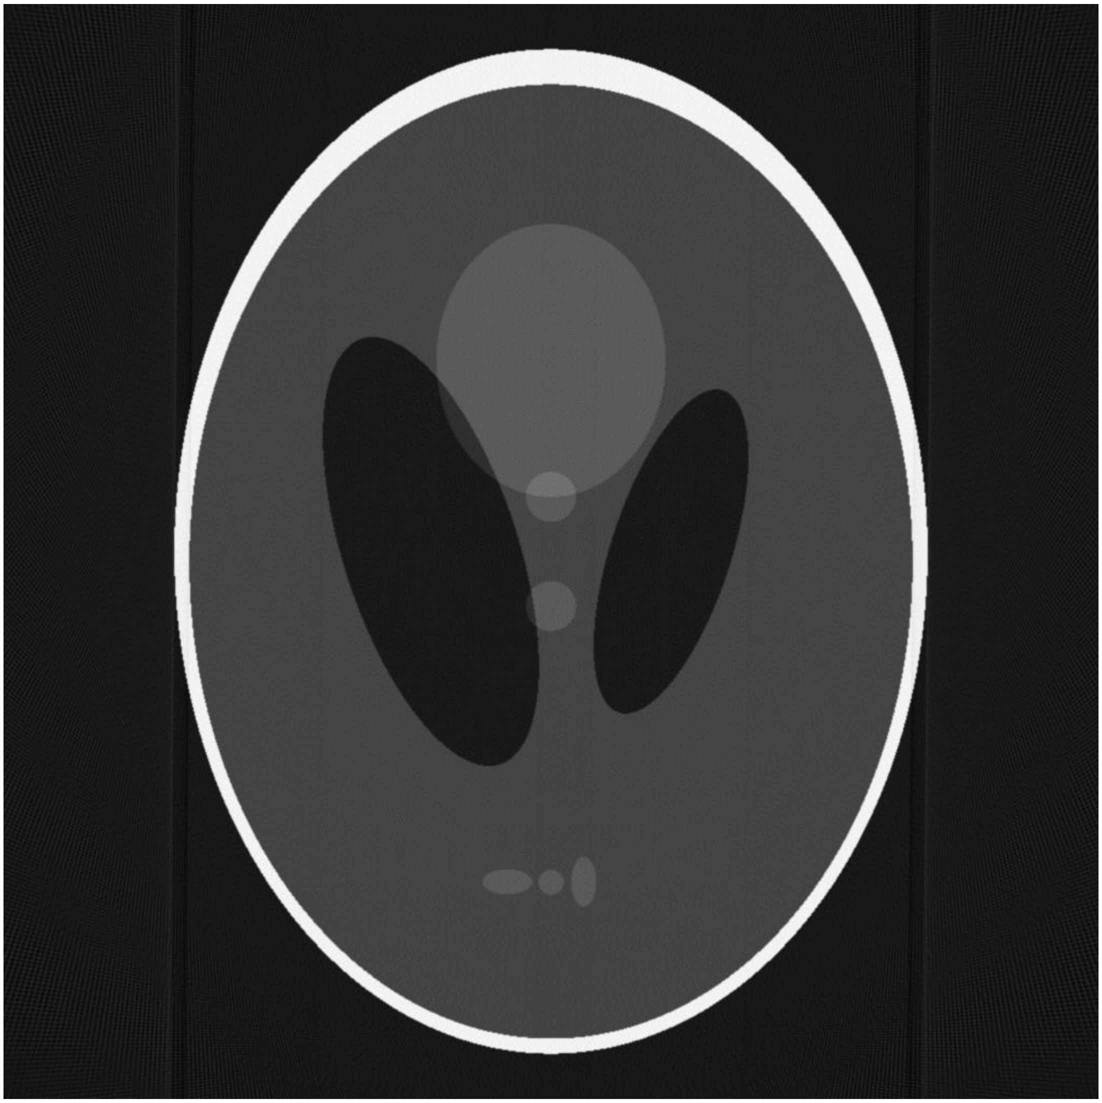

Figure 14: Reconstructed image of the Shepp-Logan phantom using $|w|$ filter.

4.4 Noise

After careful inspection of Figure 6 and 12, the reader might notice some lines or dots showing up where the original image does not. These artifacts are what we call noise. It's a mathematical consequence of the tools we are using to generate the image. It is not due to incomplete or partial data that causes these errors. This is a consequence of the filtering process. Some people spend their entire careers dealing with post-image processing where they attempt to remove noise from the images as much as possible. In the next chapter, we will briefly talk about different filter options to reduce some of that noise.

5 Filter Design

While the ideal filter option is the absolute omega option from a mathematical point of view, this specific filter can lead to more noise than others. In the discrete world, this filter is known as the Ram-Lak filter. When that filter is multiplied by the cosine function then it is called the hanning window. We can multiply absolute omega by different functions to get different frequency filtering effects. Some can remove more higher frequencies and some can remove lower frequencies. Sometimes sharpness is sacrificed for noise reduction and other times noise reduction is sacrificed for image sharpness. The pictures below show the filter design on the left and the right images are the reconstructed images.

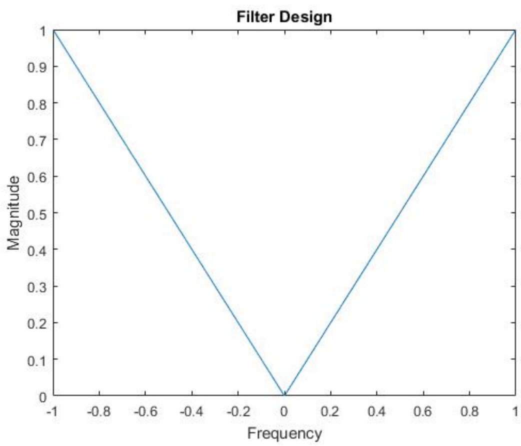

(a) Ram-Lak Filter

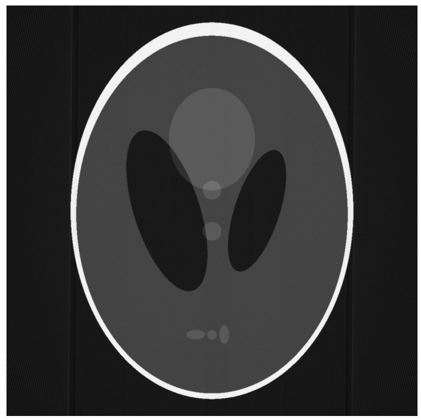

(b) Reconstructed Image

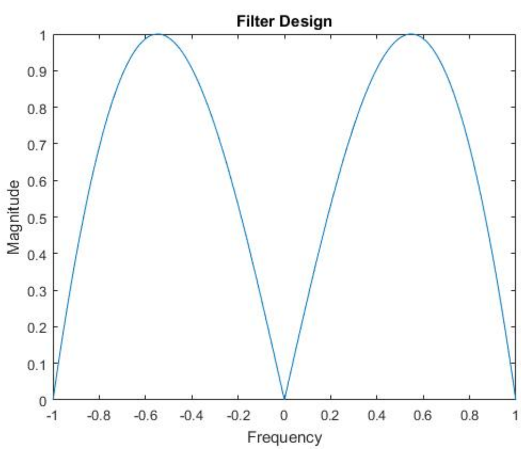

(a) Hanning Filter

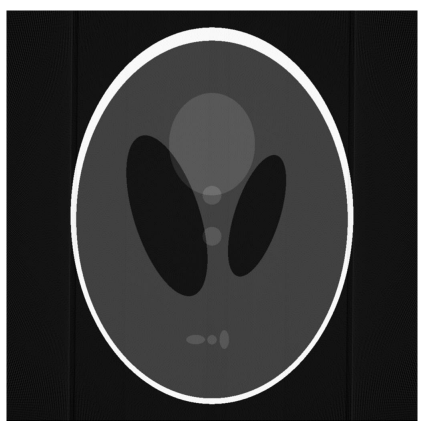

(b) Reconstructed Image

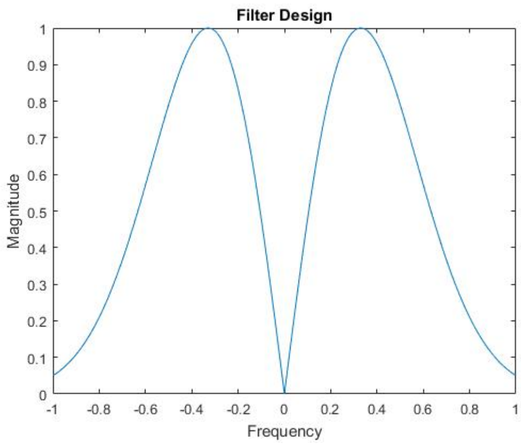

(a) Gaussian Filter

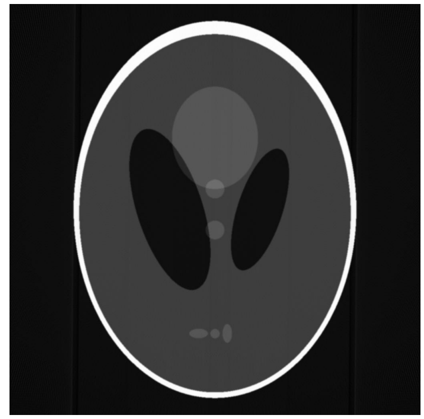

(b) Reconstructed Image

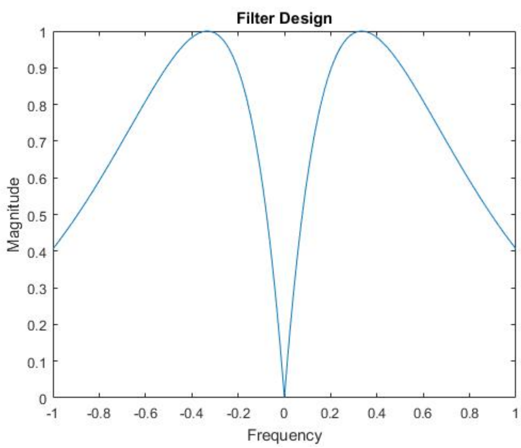

(a) Absolute Exponential Filter

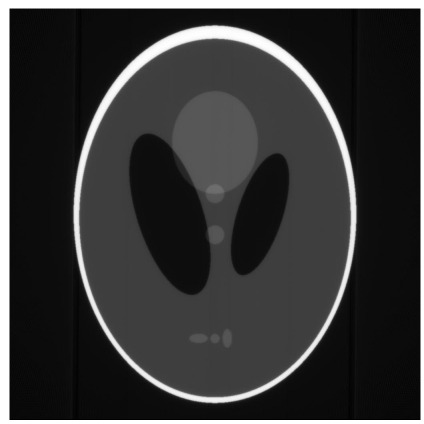

(b) Reconstructed Image

5.1 Error Images

A closer examination is needed since the images look nearly identical to the untrained eye. Zooming in closer to the images, it becomes easier to see where the differences exist. To further illustrate this, we will produce what is known as an error image. If we use the original phantom (meaning not a reconstructed one) as a basis, then we can subtract pixel values between the reconstructed phantoms from the filters and the pixels of the original phantom. Areas that are black means there was little difference. In other words, black areas have excellent reconstruction whereas bright (white) areas means there was a big difference between the reconstructed value and the correct value. An error image is shown below for filter 1 and filter 5 respectively.

Figure 19: Error image for Ram-Lak filter.

Figure 20: Error image for absolute exponential filter.

From the figures, we can see the absolute exponential filter produces an outline that is more white all around than the Ram-Lak filter. Therefore, it is blurrier or more incorrect in the skull outline area of the phantom. However, we notice the background black around the the skull is darker than the black areas of the Ram-Lak filter. This means the noise artifacts that occur around the skull from the reconstruction are minimized more than in the Ram-Lak filter; the price to pay in this case is a slightly blurrier image. There are other techniques used besides filter manipulation to remove noise from images, but this topic remains outside the scope of this paper.

6 The Complete Theorem

Theorem 1. *Let a phantom, or any function for that matter, be defined as f on the plane with real numbers x and y. It follows that the reconstructed function is described as*

$$f(x,y) = \frac{1}{2} B\left\{F^{-1}[|\omega| F(Rf)(\omega,\theta)]\right\}(x,y) \tag{16}$$

where B is the back projection integral operator and R is the Radon integral operator. The F and F^{-1} operators consist of the Fourier and inverse Fourier transform pair.

The function f is feed to the Radon integral where the output variables are changed from x and y to t and θ. We know that t is defined as a distance perpendicular to the line of integration. Also, the t line and the positive x axis form an angle θ and this relationship is expressed as $t = x\cos(\theta) + y\sin(\theta)$. Once this transformation is complete, the Radon function is fed to the 1D Fourier transform integral such that the variable of integration is with respect to t. The new frequency function is multiplied by absolute ω where ω values are defined by the Nyquist theorem. The filter is chosen so that high frequencies are emphasized and low frequencies are de-emphasized to create a sharpening effect. Once that is achieved, the inverse 1D Fourier transform is taken with the variable of integration as the frequency variable ω. At this point, the filtered sinogram is still in terms of t and θ and interpolation will be required to map polar values to Cartesian values. Finally, the back projection integral is applied to the filtered sinogram to give the reconstructed image. If we fully expand the theorem above, we get the following new theorem.

Theorem 2. *Let $f \in R^2$ such that the reconstructed function is*

$$f(x,y) = \frac{1}{(2\pi)^2}\int_{\theta=\pi}^{\pi}\left(\int_{\omega=-\infty}^{\infty}|\omega|\left(\int_{t=-\infty}^{\infty}\left(\int_{s=-\infty}^{\infty} f(t\cos(\theta) - s\sin(\theta), t\sin(\theta) + s\cos(\theta))ds\right)e^{-i\omega t}dt\right)e^{i\omega t}d\omega\right)d\theta$$

The theorems shown above are under the assumption that the data is continuous. However, when undergoing a CT scan in the doctor's office, it is safe to say that there will be discrete packets of information meaning only finite data exists. The next section discusses this concept to a greater extent.

6.1 The Discrete World

From the text so far, we have discussed the continuous world where Radon, Fourier, and back projection are all integrals. However, when dealing with real-world data, the integrals turn to summations over finite intervals with discrete increments. Thus, the following definitions are the discrete counter parts to the equations discussed up to this point.

An X-ray machine cannot take measurements along every line labeled $l_{t,\theta}$. Therefore, there is a finite number of angles labeled θ between 0 and π for finite values of t since there is spacing between the rays sent from the machine. If scans are produced at N different angles with an increment of $d\theta = \pi/N$ then the values of theta are defined as $\{k\pi/N : 0 \leq k \leq N-1\}$. The spacing between the rays is defined as τ so that the sample spacing contains $2 \cdot M + 1$ beams at each angle. Then the value of t are $\{j\tau : -M \leq j \leq M\}$. The values of M and τ will then depend on the specific construction of the machine itself.

Definition 5. The discrete Radon transform of a discrete function f is defined using the line parameters described above such that

$$R_D f_{j,k} = Rf(j\tau, k\pi/N). \tag{17}$$

Definition 6. The discrete Fourier transform is defined as follows where j is the Fourier coefficient number, k is the current data point in the summation, N is the number of points that defines an N-periodic discrete function f.

$$(F_D f)_j = \sum_{k=0}^{N-1} f_k e^{-i2\pi kj/N} \text{ for j = 0, 1, ..., (N - 1).} \tag{18}$$

Definition 7. The discrete inverse Fourier transform is defined as follows where n is the data point number, k is the current Fourier coefficient in the summation, N is the number of points that defines an N-periodic discrete function g.

$$(F_D^{-1} g)_n = \frac{1}{N} \sum_{k=0}^{N-1} g_k e^{i2\pi kn/N} \text{ for n = 0, 1, ..., (N - 1).} \tag{19}$$

Definition 8. The discrete back projection contains discrete theta steps such that θ is replaced by $k\pi/N$ defined from $0 \leq k \leq N-1$ and so a discrete function h has the back projection defined as

$$B_D h(x,y) = \frac{1}{N} \sum_{k=0}^{N-1} h(x\cos(k\pi/N) + y\sin(k\pi/N), k\pi/N). \tag{20}$$

Of course, the discrete linear interpolation function described in chapter 3 will be useful when performing the back projection on missing data that lies between two points. Moreover, the Nyquist theorem sampling rate described in chapter 4 helps us define the domain of the discrete frequency data. With all of these tools combined, it becomes straightforward to apply these reconstruction formulas in the discrete world.

7 A Final Note