

Improving Reconstructed Image Quality in a Limited-Angle Positron Emission Tomography System

David Fan-Chung Hsu

Department of Electrical Engineering, Stanford University
350 Serra Mall, Stanford CA 94305
fcdh@stanford.edu

Abstract

Limited-angle positron emission tomography systems can be placed next to the anatomical area of interest, improving their spatial resolution and diagnostic sensitivity. However, the incomplete angular sampling makes the reconstruction process susceptible to image artifacts and anisotropic spatial resolution. In this work, I show that combining total variation (TV) minimization with the maximum likelihood expectation maximization (MLEM) algorithm in the image reconstruction process can lead to greatly improved image quality for datasets with limited angular sampling range.

1. Introduction

Over the past decades, nuclear imaging has become one of the most powerful tools in an oncologist's arsenal against cancer. As a functional imaging modality, the metabolic information that is provided by positron emission tomography (PET) is complementary to the anatomical information that is provided by conventional imaging techniques such as computed tomography (CT) or magnetic resonance imaging (MRI), and is especially useful for finding cancer metastases inside the body.

As the field continues to develop, clinical investigations have begun focusing on other types of diagnostics that PET could be used for. More specifically, the ability to quantify properties such as metabolism, perfusion, and molecular concentration have applications in predicting cancer aggressiveness, looking at early-stage neural diseases such as Alzheimer's or Parkinson's, or even predicting long-term effects of concussions on football players. However, one important prerequisite for quantification of PET is accurate image reconstruction algorithms which can image the object of interest faithfully.

1.1. Limited-angle tomography

A common idea behind different tomography-based imaging modalities is the use of projection imaging over many angles. As shown in Fig. 1, each projection image

angle is composed of parallel lines of light which pass through the object and are attenuated by the density of the object that it has traversed. This information maps to a line in Fourier image frequency space. By rotating the angle of the light beam over all 360 degrees, the line that crosses through the Fourier space also rotates 360 degrees, covering the entire frequency space. Generally, sufficient sampling of data in the Fourier space is required in order to reconstruct the original object that was scanned.

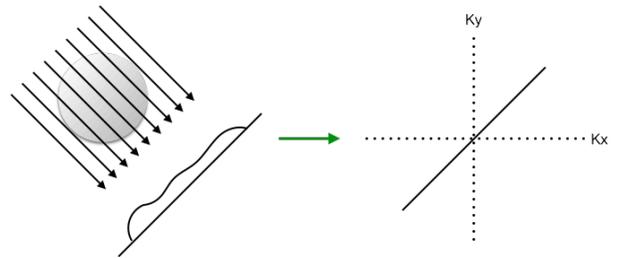


Figure 1: (Left) Illustration of projection imaging, showing an image displayed on a plane after being projected through an object. (Right) The same projection information displayed in Fourier space, showing that each projection angle produces a straight line of data in Fourier space.

Most PET systems are immobile ring-based geometries with a fixed field-of-view (FOV), which is sized to accommodate large patients. This places the detectors far from the body, reducing the achievable spatial resolution. For organ-specific imaging, a number of limited-angle tomography systems have been designed with dual parallel panels, which can be translated along one axis to be placed right up against the imaging area of interest. This can improve the best achievable spatial resolution of the system. However, limited-angle tomography results in undersampling of Fourier space, with missing data from the angles which were not scanned. This leads to image artifacts and degradation from the image creation process, which can impact clinical diagnostics [1]. A method to mitigate these artifacts can lead to greatly improved patient outcomes with these high-resolution limited-angle PET scanners.

1.2. Current reconstruction techniques

The projection dataset that is acquired from a PET scan is called a “sinogram”, shown in Fig. 2, which is composed of projection data as a function of angle. Re-forming an image of the original object from a sinogram is called “image reconstruction.”

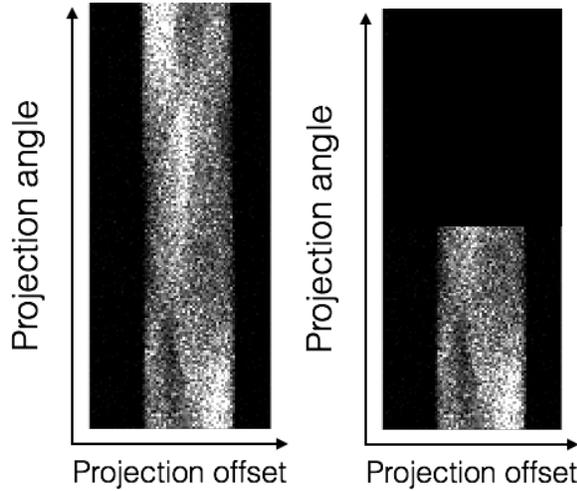


Figure 2: (Left) A sinogram of a full-angle tomography dataset, showing projection angle on the vertical axis and line integral offset on the horizontal axis. (Right) The same object scanned with limited-angle tomography, showing a lack of data over a certain number of projection angles.

Image reconstruction can be broadly split up into two different categories: analytical and probabilistic. Analytical reconstruction is a “faithful” reconstruction, in that it takes the sampled Fourier space data and does an inverse Fourier transform, with appropriate filtering in place. The formula that governs this is shown below in (1), where $P(\theta, k)$ is the projection data in Fourier space, and μ is the density of the original object that was imaged. This method is also known as “**Filtered Backprojection**” (FBP).

$$\mu(x, y) = \int_0^{2\pi} \int_0^{\infty} P(\theta, k') e^{2\pi i k' (x \cos \theta + y \sin \theta)} k' dk' d\theta \quad (1)$$

$$\mu = \operatorname{argmax}_{\mu} \{P(\vec{\mu} | \vec{p})\} \quad (2)$$

Probabilistic image reconstruction is based on the idea that, given the set of projection data, find the most likely object that could have created this set of data. The formulaic representation is shown in (2), where P represents the probability, \vec{p} represents the projection dataset and μ represents the object. The actual implementation of this method is called “**Maximum Likelihood Expectation Maximization**” (MLEM), which is an iterative method that iterates between the projection

sinogram space and the image μ space until convergence. It has become the de-facto industry standard for PET scans, despite the fact that it is not as “faithful” as analytical reconstruction, because of its robustness to noise in the projection space. This can be seen in Fig. 3.

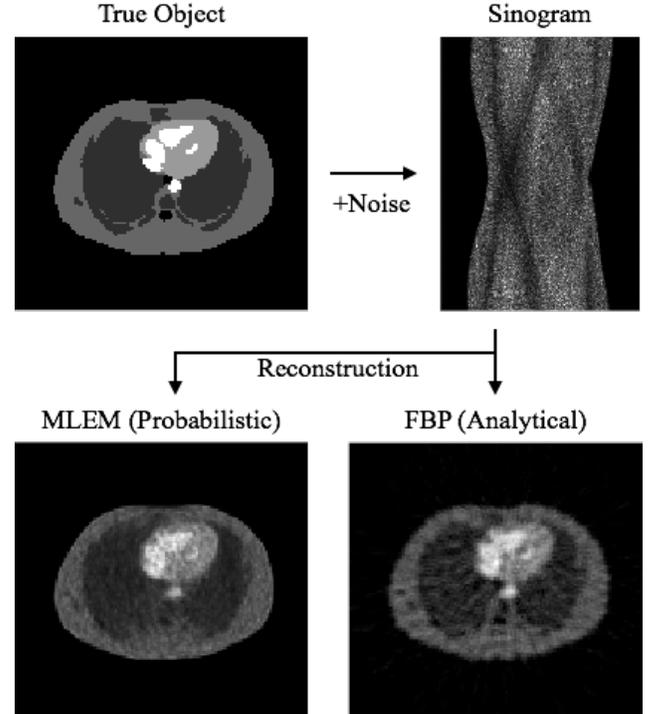


Figure 3: Example of full-angle tomography reconstruction process with noise. The probabilistic reconstruction image clearly shows better noise statistics than the analytical reconstruction.

1.3. Investigational Approach

When we limit the object’s sampling angle, severe distortions appear in both the analytical and probabilistic reconstruction, as shown in Fig. 4. The distortions and artifacts in the FBP analytical reconstruction can be explained by the undersampling in Fourier space. However, since the probabilistic MLEM reconstruction method is not directly dependent on an inverse Fourier transform, it should be possible to improve its performance. My hypothesis is that the undersampling in Fourier space increases the number of possible solutions which MLEM can converge to. From this paradigm, it would appear that the image degradation simply emerges as a result of a difference in the **convergence rate** of areas which are undersampled, versus areas which are well-sampled.

A number of works [2-4] have shown that it is possible to improve the image quality of limited-angle reconstruction in computed-tomography by using total-variation (TV) minimization. However, this has not

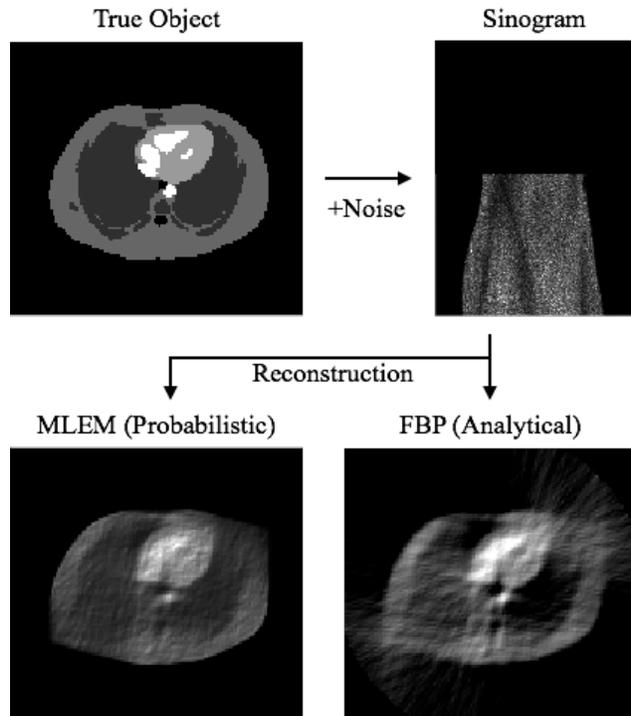


Figure 4: Example of limited-angle tomography reconstruction process. Image artifacts are clearly visible from the undersampling in Fourier space.

been previously investigated for PET and MLEM-based reconstruction algorithms for limited-angle tomography. There are several intrinsic reasons why the TV-minimization prior may create favorable results when incorporated in MLEM reconstruction:

1. Image areas with high contrast will have sharp edges.
2. Image areas with low contrast will naturally converge faster than high contrast areas.

By utilizing TV-minimization during the MLEM process, we ensure that we bring out edges that are present in under-sampled and high-contrast areas, therefore speeding up convergence in these areas and ensuring the entire image converges at roughly the same rate.

In this work, all the data that is used is from simulations in order to obtain a ground truth image, with which I can compare the effectiveness of the results. I obtain a set of projection data (sinogram) using a limited-angle tomographic simulation of the ground truth object, and reconstruct this limited-angle dataset with MLEM with different TV-minimization methods incorporated within. I will also compare the peak signal-to-noise (PSNR) and mean-squared error (MSE) of the reconstructed images with the original object.

2. TV-Minimization

To incorporate TV-minimization into MLEM, I took the output image after a certain number of iterations of MLEM, ran it through a TV-minimization algorithm, and sent the processed image back into the MLEM algorithm for more iterations, as shown in Fig. 5.

In this work, I focused on anisotropic TV-minimization, which is shown by equation (3). In this equation, x is the image pixel values, C is the blurring function (also called the “point spread function”), and b is the noisy image. The anisotropic nature means that the contributions of gradients in the x - and y - directions in the image are incorporated into the penalty function separately, instead of jointly. In my testing, I utilized an equal weight of edges in both directions.

$$x^{i+1} = \operatorname{argmin}_{x^i} \{ \| Cx^i - b \|_2^2 + \lambda \| \nabla x^i \|_1 \} \quad (3)$$

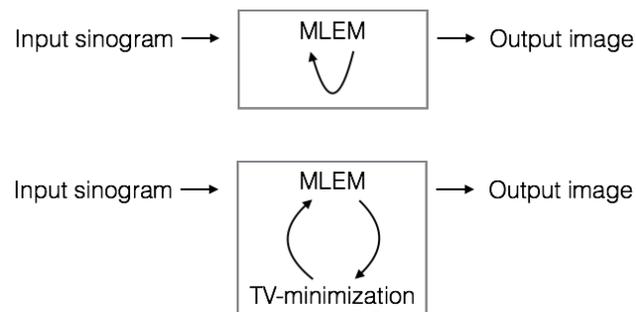


Figure 5: (Top) Original MLEM-based reconstruction. (Bottom) Incorporating TV-minimization into MLEM image reconstruction.

I investigated two different methods of performing TV-minimization, using different methods of point-spread function (PSF) modeling.

2.1. ADMM-based TV-minimization

The alternating direction method of multipliers (ADMM) is a method of optimizing a sum of two objective functions which can be linked together in a linear fashion. In the case of TV-minimization, the objective function is for image de-blurring (with a PSF present), which is shown the left side of the addition sign in (3), while the other is for TV-minimization, shown by the right side of (3).

The actual implementation of ADMM-based TV minimization is very complicated. It involves forming an augmented Lagrangian function, and performing sequential iterative updates to objective 1, 2, and the linear function that links them, as shown in equation (4). For more information, please refer to EE 367 course notes, or the original ADMM publication by Boyd et al. [5-6]

$$\begin{aligned}
L_\rho(x, y, z) &= f(x) + g(z) + y^T(Ax + Bz - c) + (\rho/2)\|Ax + Bz - c\|_2^2 \\
x^{k+1} &:= \underset{x}{\operatorname{argmin}} L_\rho(x, z^k, y^k) \\
z^{k+1} &:= \underset{z}{\operatorname{argmin}} L_\rho(x^{k+1}, z, y^k) \\
y^{k+1} &:= y^k + \rho(Ax^{k+1} + Bz^{k+1} - c)
\end{aligned} \tag{4}$$

2.2. Richardson-Lucy based TV-minimization

The Richardson-Lucy (RL) deconvolution method is based upon finding the maximum-likelihood solution of images that are dominated by Poisson noise. Since PET image reconstruction is based upon the emission of photons from nuclear decay, which is intrinsically a Poisson process, it makes RL-based TV-minimization a method that is worth trying. However, it is worth noting here that, even though the emission process is a Poisson process, the MLEM-based image reconstruction may not produce images with Poisson noise statistics.

The RL deconvolution process is based on the idea that at convergence, further iterations will not change anymore. Adding a total variation term in the log likelihood function allows us to perform RL-based TV-minimization. The iterative update equation is shown in (5). The derivation will not be done here, and can be found in EE 367 course notes [6].

$$\mathbf{x}^{(q+1)} = \frac{\mathbf{A}^T \left(\frac{\mathbf{b}}{\mathbf{A}\mathbf{x}} \right)}{\mathbf{A}^T \mathbf{1} - \lambda \left(\frac{\mathbf{D}_x \mathbf{x}}{|\mathbf{D}_x \mathbf{x}|} + \frac{\mathbf{D}_y \mathbf{x}}{|\mathbf{D}_y \mathbf{x}|} \right)} \cdot \mathbf{x}^{(q)} \tag{5}$$

3. PSF Modeling Techniques

Both the ADMM-based and RL-based TV-minimization methods mentioned above include a point-spread function (PSF) deconvolution method, in addition to TV-minimization. This is represented by the C term in equation (3) and the A term in equation (5). The PSF is a representation of how a singular impulse point in the input will be shaped by the system. Because an image can be thought of as an addition of many impulse points that are scaled to their pixel values and added together linearly, a system that is linear and spatially-invariant can be completely characterized by a single PSF.

Unfortunately, MLEM image reconstruction on a limited-angle tomographic dataset does not create a spatially-invariant PSF. The reason is that the spatial blurring and degradation is dependent on the direction of edges and other high-contrast features in the original object. Therefore, given a particular pixel location and exactly the same system matrix, the PSF that is associated with that pixel location will be dependent on the object that is being imaged. This can be seen in Fig. 6. Therefore, the PSF in this problem formulation is spatially-variant and object-dependent.

3.1. Point Perturbation Method

One method to find a PSF for this image system is called a point perturbation method. This is done by inserting a grid of points into the original object that is being simulated, performing image reconstruction, and then doing a pixel-wise subtraction of this point-perturbed image with the original unperturbed image. The resulting image will show the contributions of point impulses at different locations in the field-of-view on the final image.

In this work, two different phantoms were used: a circular phantom, and a Zubal phantom, as seen in Fig. 6. The point perturbation method was used on these two phantoms, resulting in a grid of PSFs in the bottom row of Fig. 6. Once we have found the grid of PSFs using the point perturbation method, it must be incorporated into the ADMM and RL deconvolution methods somehow. Different methods for incorporating PSF into deconvolution are described in the sections 3.2 to 3.4 below.

One very important point to note here is that the point perturbation method does not necessarily need to be applied in image-space, as is done in the simulation described above. If one has the projection data from a grid of points, this projection dataset can be added to the projection dataset of the object of interest, yielding equivalent results. The PSF modeling techniques that are described in this section can therefore be extended to real and non-simulated datasets in this fashion.

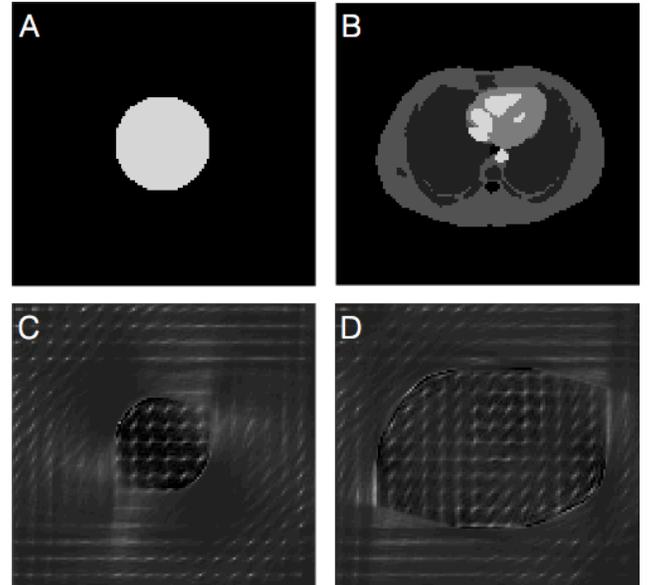


Figure 6: (A) Circular and (B) Zubal phantoms, which are simulated. (C,D) PSF grid for each phantom, found using the point perturbation method.

3.2. Average PSF

Firstly, each PSF in the grid is cropped out individually, making sure that every PSF is centered and all the PSFs are of the same spatial size. Afterward, I averaged all of the PSFs in the entire image to create an average PSF. This mean PSF is then used in the ADMM- and RL-deconvolution process.

3.3. Eigen-PSF using KL-decomposition

This method is detailed in source [7], and is called the Karhunen-Loève decomposition method. After each PSF is cropped out, they are all normalized so the integral of each PSF is equal to 1. After normalization, we have a set of PSF's, which we call P_i , where i is the total number of PSFs. We then calculate a covariance matrix C , where each element is the 2D covariance of two different PSFs, as shown in (6). We then calculate the eigenvectors of the covariance matrix, resulting in a set of i eigenvectors, as shown in (7). Starting with the eigenvectors corresponding to the highest eigenvalues, we define eigen-PSFs according to equation (8), each eigen-PSF being a linear combination of all the PSF's which are present in the system.

$$C_{ij} = Cov(P_i, P_j) \quad (6)$$

$$x_i = \text{eigenvector}(C) \quad (7)$$

$$PSF_i = \sum_{j=1}^N x_{ij} P_j \quad (8)$$

In this method, the number of non-zero eigenvalues show the number of “independent” PSF's there are in the image. In my investigation, I used the top 4 eigen-PSF's in my investigation in the ADMM and RL deconvolution, resulting in 4 images for each process (1 image per eigen-PSF), which were then scaled and added together using on the eigenvalues corresponding to those eigen-PSFs.

3.4. PSF with Block-Interpolation

In this method, I perform ADMM and RL deconvolution on the image using all of the PSFs which are present. For example, if there are 120 PSFs in the grid, I perform deconvolution 120 times each time with a different PSF in the grid, resulting in 120 different images. To combine these images, I perform pixel-wise interpolation between the images corresponding to the 4 spatially closest PSFs, scaled appropriate to account for the distance from the pixel to the surrounding PSFs.

4. Methods

All data used in this work are based on simulation data,

although as previously mentioned, all the steps outlined in this work can be adapted to real-world data. I used 2 phantoms in my simulations, which are the circular and Zubal phantoms shown in Fig. 6A and 6B. I simulated a 30 degree limited-angle tomographic dataset of these phantoms, which corresponds to 1/6 of the full angular data. I then tested the following combinations of deconvolution methods and PSF models on this projection dataset:

- A. ADMM with average PSF
- B. ADMM with eigen-PSF
- C. ADMM with PSF block-interpolation.
- D. RL with average PSF
- E. RL with PSF block-interpolation

For the ADMM methods, I optimized the parameters manually, looking for the optimal visual results. I used a constant ρ value of 10 for all iterations. I used a variable λ that varied linearly, as the number of iterations between MLEM and ADMM increased, from 0.2 to 0.01. I found that this led to the best convergence properties when compared to a constant λ value. For each ADMM instance, I iterated ADMM 100 times. I found that 4 iterations between MLEM and ADMM resulted in the best output image. Fewer iterations resulted in blurrier features, while more iterations resulted in noise amplification through the MLEM process.

For the RL methods, I used a λ of 0.01, and 10 iterations per RL instance. I found that 2 iterations between MLEM and RL led to the best results. More iterations led to severe “peaking” of image pixels around the periphery of the object, which may be caused by division of numbers close to zero in the denominator.

For each of these test cases, I quantified the peak SNR and mean square error when comparing the final output to the ground truth images.

5. Results and Discussion

Table 1 below shows the numerical results. The shaded rows, which represent the current image reconstruction techniques, show currently achievable results. Comparing the results from the different methods to MLEM alone, we see that method A and B both offer superior performance

	Circular Phantom		Zubal Phantom	
	PSNR	MSE	PSNR	MSE
FBP	2.30	0.59	4.17	0.38
MLEM	9.68	0.11	9.63	0.11
A	13.05	0.05	10.61	0.09
B	9.67	0.11	10.39	0.09
C	9.11	0.12	9.23	0.12
D	10.01	0.10	8.77	0.13
E	7.60	0.17	8.31	0.15

Table 1: Numerical results of the different methods that were tested. The shaded rows represent the current analytical (FBP) and probabilistic (MLEM) image reconstruction algorithms.

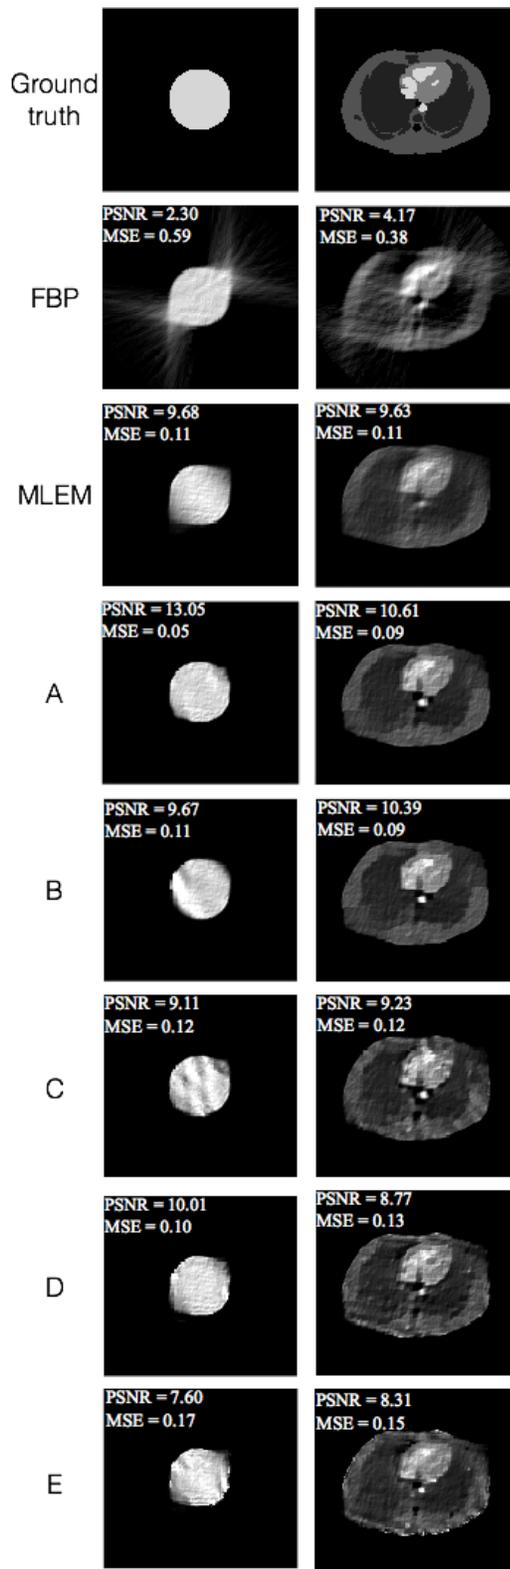


Figure 7: The reconstructed images using the different methods as outlined in part 4.

when compared to MLEM alone. The different reconstructions are shown in Fig. 7. We can visually see that method A is the best for both phantoms, while if we consider only the Zubal phantom, then method B also performs well.

One observation we can make is that the ADMM-based deconvolution methods result in less noisy images and smoother edges than RL-based methods. As both MLEM and RL methods tend to push certain pixels toward zero, the appearance of very small pixel intensities in the denominator could cause some pixels to spike in intensity during the MLEM-RL iteration process.

Another observation we can make is that, in both of the PSF block-interpolation methods (C and E), the results are blotchy and have weird artifacts in the images. This might be attributed to poor interpolation process and significant differences between neighboring PSFs. The dependence of PSF on the features in the object likely make this method a poor choice for dealing with limited angle tomographic images.

It is also interesting to note that method A, which is the “average PSF” method, achieves the best results. Since I have previously mentioned that the PSF is both spatially-variant and dependent on the object, it seems weird that an average PSF would do better than an eigen-PSF, for example.

6. Future Work

This current work will most likely be submitted to the next IEEE NSS/MIC conference. I will adapt this work to real-world data that is collected by our limited-angle tomography system, which is designed for breast cancer screening, to see whether or not it can improve image quality.

Improvements that can be made to this process include using a different variant of anisotropic TV norm. Currently, the gradients in x and y directions are considered separately but with equal contributions. Since we know that image quality degradation in limited-angle tomography is dependent on the orientation of high-contrast edges in the object versus the projection angle, it should be possible to improve our prior to account for the directionally-dependent edge degradation.

7. Acknowledgements

The author would like to acknowledge the helpful MATLAB packages developed by Dr. Jeff Fessler at the University of Michigan, for performing the simulations of limited-angle tomographic datasets. In addition, the author would like to thank the EE 367 course staff for being so knowledgeable, and making this course one of the best courses he’s taken at Stanford.

References

- [1] Berg WA, Madsen KS, Schilling K, et al. Comparative effectiveness of positron emission mammography and MRI in the contralateral breast of women with newly diagnosed breast cancer. *AJR Am J Roentgenol.* 2012;198:219-232.
- [2] Chen, Zhiqiang, et al. "A limited-angle CT reconstruction method based on anisotropic TV minimization." *Physics in medicine and biology* 58.7 (2013): 2119.
- [3] Isernhagen, C. F., et al. "Three-dimensional anisotropic regularization for limited angle tomography." *Current Directions in Biomedical Engineering* 1.1 (2015): 283-285.
- [4] Jonsson, Elias, Sung-cheng Huang, and Tony Chan. "Total variation regularization in positron emission tomography." *CAM report* (1998): 98-48.
- [5] Boyd, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." *Foundations and Trends® in Machine Learning* 3.1 (2011): 1-122.
- [6] EE 367 course notes, Winter quarter 2015-16.
- [7] Lauer, Tod. "Deconvolution with a spatially-variant PSF." *Astronomical Telescopes and Instrumentation*. International Society for Optics and Photonics, 2002.