Exercises on the Radon transform and the Filtered Back-Projection algorithm

Jakob Sauer Jørgensen, DTU Compute, `jakj@dtu.dk`

These exercises are intended to make you familiar with the Radon transform used in the forward problem of CT, i.e., modeling how CT data is generated, as well as the Filtered Back-Projection (FBP) algorithm for image reconstruction from CT data. Exercises 1 and 2 (for Monday morning and afternoon) use simulated data and will prepare you for doing Exercise 3 on reconstruction from a real CT data set Tuesday morning. The exercises are designed for MATLAB with commands stated in `typewriter` font.

Exercise 1: The Radon Transform

1.1 Phantoms and Sinograms

1. The sinogram is called so because the Radon transform of a single white pixel on a black background is a sinusoid curve in the sinogram. Generate such phantoms and inspect the corresponding sinograms. Example: `X = zeros(128); X(10,10) = 1; S = radon(X); imagesc(S)`

2. A popular test image in CT is the Shepp-Logan phantom. It is available in MATLAB using the command `phantom`. Generate and display the default Shepp-Logan phantom. Try different sizes, such as `N = 101, 201, 501, 1001`.

3. The Radon transform of an image can be computed using `radon`. Pick a range of angles such as `= 0:179` and generate and display the Radon transform (the sinogram) of the Shepp-Logan phantom. Each column of the sinogram corresponds to a 1D projection at a particular angle (try plotting individual columns to better see how intensity values varies as function of position). Try also using more angles (smaller angular increment). Try to match up features in the sinogram with features of the Shepp-Logan phantom.

4. Try extending the angular range up to 360°. Can you see any symmetries in the sinogram? Do we need 360° data to reconstruct the image?

5. Due to the many ellipses present in the Shepp-Logan phantom it can be difficult to fully understand the sinogram. A user-defined phantom consisting of a superposition of ellipses can be created, see `doc phantom`. For example a single centered disk in a 501-by-501 image can be created using `P = phantom([1, 0.2, 0.2, 0, 0, 0], 501);` Create this phantom and try to predict what its sinogram will look like. Verify using `radon`.

6. Repeat the experiment with a single ellipse while changing its parameters to get a larger disk, a disk with changed intensity, an ellipse, a rotated ellipse, and a translated ellipse.

7. Describe how the sinogram changes (for example in terms of amplitude, phase, thickness, intensity) as function of the ellipse position, size and orientation.

1.2 Linearity of the Radon Transform

The Radon transform is a linear transform. This means that a sinogram of an image can be decomposed into a sum of sinograms of the various objects in the image.

1. Verify analytically that the Radon transform is linear, i.e., that $\mathcal{R}(ax + by) = a\mathcal{R}(x) + b\mathcal{R}(y)$ for two images x and y and constants a and b.
2. Verify linearity numerically: Generate two different test images `X` and `Y` with a single ellipse in each. Generate and display the sinogram of each. Generate also the sinogram of a linear combination of `X` and `Y`, for example, `Z = 3*X + 2*Y`. Demonstrate that the sinogram of `Z` equals the same linear combination of the sinograms of `X` and `Y`.

1.3 The Lambert-Beer Law, Intensity/Absorption Data, and Noise

1. We use the Shepp-Logan phantom `X` of size 251×251 pixels..
2. Choose the angles `theta = 0:0.5:179.5` and compute and display the *absorption sinogram* $\mathcal{R}(\mathtt{X})$ using `radon`. The sinogram must be scaled to take the pixel size into account: Assuming physical image size of 1, divide the sinogram by $N = 251$:
   `Y = radon(X,theta)/251`.
3. From the Lambert-Beer law, $I = I_0 \exp\{-\mathcal{R}(\mathtt{X})\}$ determine the intensity data I. We assume an incident flux $I_0 = 10^4$ photons/detector bin. Display this *intensity sinogram* and use `colorbar` to show the range of values. Compare with the absorption sinogram, for example what do 0-values in the absorption sinogram correspond to in the intensity sinogram? It may also be helpful for comparison to plot single columns of the sinograms.
4. In practice, data are noisy due to the fact that the number of transmitted photons can be modelled as a Poisson process. Poisson noise can be simulated on the intensity data using the function `poissrnd` (from the Statistics toolbox) on I, which corresponds to taking samples from a Poisson distribution with parameter I. Alternatively, for large I_0 the Gaussian distribution with mean and variance I closely approximates the Poisson distribution and samples can therefore be generated using `randn`. Apply Poisson noise to the clean intensity sinogram I obtained for $I_0 = 10^4$ to determine noisy intensity sinogram I_{noisy} and compare with the clean intensity sinogram. You can compare either by plotting single columns of the sinogram or by viewing the entire sinogram as an image.
5. Determine also the noisy *absorption* sinogram by dividing by I_0 and log-transforming, and compare with the clean absorption sinogram $\mathcal{R}(\mathtt{X})$ obtained previously.
6. Try two different noise levels corresponding to two different incident intensities, e.g., $I_0 = 10^6$ and $I_0 = 10^2$. What happens to the data when the incident intensity is decreased?

1.4 Analytical vs Numerical Radon Transform of Ellipses

1. Implement the formula to compute the Radon transform of an ellipse at a fixed angle θ and for a vector with two elements with the coordinates $(x_{\text{c}}, y_{\text{c}})$. It may be convenient

to first write a function for the non-rotated, non-translated ellipse, and then a separate function to apply rotation and translation using the first function. Hint: Use `atan2` or `atan2d` for the $\tan^{-1}$.

2. Generate a test image with a single ellipse. Compute the "numerical" projection using `radon` and the "analytical" projection using your own function, and compare them. Try other ellipses to make sure your function works correctly. Due to the way `radon` selects the center of rotation you will see the best agreement if you use images of odd size, e.g., 501 rather than 500×500.

3. Compare the analytical projection with numerical projection obtained using increasingly fine discretizations (larger images). Do the numerical projections approach the analytical ones, as expected?

Exercise 2: Filtered Back-Projection

2.1 Reconstruction

1. Generate clean and noisy projection data as in the previous exercise using $I_0 = 10^4$ and size 251. Use 512 angles over 180°.

2. Compute an FBP reconstruction from the clean and noisy projection data using `iradon`. Check the help of `iradon` to specify inputs, in particular the `OUTPUT_SIZE` input must be set to 251 to get the correct image size. Other inputs can be set to their defaults.

3. Display the reconstructions and compare with the original Shepp-Logan image. Small differences can sometimes be seen clearer by displaying the difference images and/or by displaying only subregions of the full image.

2.2 Changing the Filter

1. By default `iradon` uses the Ram-Lak filter. Experiment with the effects on the reconstruction of using other filters, including the filter `'none'`, which corresponds to an unfiltered backprojection.

2. Similarly experiment with the effects of using other values between 0 and 1 for the input `frequency_scaling`. Smaller values introduce more low-pass filtering.

3. Increase the noise by reducing I_0. Experiment with filter and frequency scaling to get the best reconstruction.

2.3 Linearity of FBP

The FBP is a linear reconstruction algorithm meaning that the complete reconstruction can be obtained as the sum of reconstruction from parts of the data. This can be used in practice to compute the contribution to the reconstruction from each projection as soon as it is acquired, i.e., before the scan is complete.

1. Partition the full data set into four partial data sets, e.g., $[0°, 45°[$, $[45°, 90°[$, $[90°, 135°[$ and $[135°, 180°[$. You may also partition in other ways for example by interleaving the four sets. Make the FBP reconstruction from each partial data set, display and comment.
2. Now sum the four reconstructions and compare with the reconstruction from all data.

2.4 Limited Data

FBP requires many projections/a small angular increment and projections from the complete angular range between 0° and 180°.

1. The effect of *few projections* from full angular range ("few-view data"):
   Subsample the data so that only angles `theta = 0:5:175` are used, i.e., a reduction by a factor 10, compared to the first exercise. Make an FBP reconstruction and comment on the introduced artifacts. Try even fewer projections.
2. The effect of many projections but *incomplete angular range* ("limited-angle data"):
   Use only data for `theta = 0:0.5:139.5` to make FBP reconstruction and comment. Try further reducing the angular range. Describe which features of the image are lost.

The material on Wednesday afternoon will go into the mathematical description of the limited-angle artifacts and how they may be reduced.

2.5 The Region-of-Interest Problem

This exercise illustrates the region-of-interest problem (also called the interior problem), which occurs when the sample is not fully contained within the field of view. This can happen if the sample is too large or if one chooses to zoom in on a specific region inside the sample – the latter is the case in the real-data reconstruction exercises Tuesday afternoon. Applying FBP directly yields a so-called cupping artifact, which can be reduced by a simple trick.

1. Set up a 513×513 Shepp-Logan phantom, generate the sinogram using 1024 angles over 0° to 180° degrees and compute the FBP reconstruction.
2. Discard the top and bottom 300 rows of the sinogram to simulate that only data of a small region-of-interest was measured.
3. Compute an FBP reconstruction from the region-of-interest data, specifying that the reconstruction should be of size 513. Compare with the original phantom and a reconstruction from full data. Adjust the display color range using `caxis` such that all reconstruction have the same color range. Also plot the central vertical profiles through each reconstruction to better see the cupping artifact.
4. A simple trick of padding the data can reduce the region-of-interest artifact. Pad at the top of the sinogram with 100 copies its top row and at the bottom with 100 copies of its bottom row. Then recompute the FBP reconstruction and compare the image and central vertical profile with the previous reconstructions.
5. Play with the number of rows to cut and the number to pad by. Can you get a good reconstruction of the full image? Of the region of interest?

2.6 A More Realistic Test Image (If Time Permits)

The Shepp-Logan phantom is a model of a cross-section through a human head (showing the skull, brain and some features inside). It is, of course, a simplification which is convenient to work with because it consists of ellipses, for which analytical projections are known. A more realistic test image resulting from an actual scan can be loaded from `brain_radiopedia.mat`.

1. Repeat some of the reconstruction exercises of your choice with the brain test image, e.g., the region-of-interest and limited-data exercises.

Exercise 3: Reconstruction of Sand Stone Data

The data used in this exercise is courtesy of Henning Osholm Sørensen, DTU Physics. A sample of sand stone was scanned at the TOMCAT beamline at the Swiss Light Source synchrotron. Using a parallel-beam scan configuration with a detector of 2560×2560 pixels, a total of 1731 images were recorded of which, according to their index:

- 1–30 are dark fields,
- 31–130 are flat fields before inserting the sample,
- 131–1631 are the actual 1501 intensity images (of which the first and last are at $0°$ and $180°$ and the latter should be excluded), and
- 1632–1731 are flat field taken after removal of the sample.

A selection of raw data images and reconstruction slices (in TIFF format) are included in the exercise files. Furthermore, for convenience of this exercise, 4 separate data sets corresponding to particular detector rows (row numbers 270, 540, 810 and 1080) have been extracted and collected in the MAT-files `slice_0270_data.mat`, `slice_0540_data.mat`, `slice_0810_data.mat`, and `slice_ 1080_data.mat`. The goal of this exercise is to do FBP reconstruction on single slices of this data.

3.1 Loading and Visualization of the Data

1. Load (`imread`) and display examples of intensity, flat field, and dark field images. Add colorbars to see the range of values.
2. Load and familiarize yourself with `slice_1080_data.mat` as well as the reference image `BBii_1080.rec.16bit.tif`, of which we hope to be produce a comparable reconstruction of. Add colorbars and adjust the color range, e.g., using `caxis` for improved contrast. Displaying a histogram of the pixel values may be helpful to determine a good display color range.

3.2 Flat- and Dark-Field Correction

1. Apply flat- and dark-field correction to the intensity data: Compute the elementwise means of the dark fields and flat fields, to get a single dark field and a single flat field to use in the correction. Then apply this single dark and flat field for correcting all the intensity data.

2. Perform a negative log-transform to get from the transmission data to the absorption data.

3. Set the angle vector and compute an FBP reconstruction using `iradon`. Compare with the reference image (be prepared for some artifacts).

3.3 Center of Rotation Correction

The observed heavy artifacts are due to misalignment: `iradon` assumes that the projections are centered with respect to the center of rotation, which is not the case for the provided data set. In the meta data it is given that the center of rotation was located at 1323.61 detector bin widths of the total 2560 detector bins. The problem can be solved by padding the sinogram on the top or bottom to artificially shift the measured data in order to center it with respect to the center of rotation.

1. Determine how many rows to pad and whether to pre- or post-pad the sinogram (for example using a sketch and the actual position of the center of rotation or by trial and error). Pad with zeros, reconstruct using `iradon`, and compare with the uncorrected reconstruction and the reference image.

2. Pad instead with the first or last row of the sinogram and reconstruct. Does this improve the reconstruction? If so, why?

3.4 Region-of-Interest Correction

Now a big ring is seen around the field of view along with a "cupping" artifact, i.e., artificial increase in intensity towards the edges of the circular field of view. This is due to the region-of-interest/interior problem from Exercise 2.5, since for the given data the whole sample was not fully contained within the field of view. If you did not do Exercise 2.5 on the region-of-interest problem yesterday, you may want to do that first.

1. Apply symmetric padding using the first/last sinogram row to the top and bottom of the sinogram before doing reconstruction. Determine how many rows to pad in order to get rid of the ring and cupping artifacts. Compare reconstructions as images and using profile plots of the middle row to better see the effect of the correction.

3.5 Suggestions for Further Experiments (if time permits)

1. Play with the `iradon` settings, i.e., filter and frequency cut-off to try to improve reconstruction. The reference image is a reconstruction which used a Shepp-Logan filter, but a stronger filter may be better.

2. Try reconstructing some of the other slices and compare with the provided slice reconstructions.

3. Experiment with reconstruction from reduced data, e.g., the sparse-view data. How few projections suffice for obtaining a satisfactory reconstruction?