X-ray Tomography Project Work
Inverse Problems Course 2015

Henri Vuollekoski

INTRODUCTION

X-ray tomography is a technique for producing tomographic images of objects that facilitates the investigation of the object's internal structure in a non-destructive manner. This is particularly useful in medical imaging, but has applications in industry as well e.g. reverse engineering[1].

The main idea is to take several measurements of the attenuation of X-rays while they pass through the object of interest, e.g. the patient's head. Since different tissues inside the head attenuate the signal by differing factors, a sufficiently large collection of measurements from different directions along with some clever inversion mathematics (and some computational muscles) are enough to reconstruct an image of the head's internals[2].

Safety must be prioritized when working with X-rays that can cause radiation overdose and cancer for the patient as well as the person making the measurements, or plain bystanders[1].

METHODS AND MATERIALS

The object of interest was chosen to be a *single slightly customized peanut*, the objective being a 2D slice image of its internals. See Figure 1 for a photograph of the peanut.

The measurements were taken at the Industrial Mathematics X-ray Laboratory at the University of Helsinki, Finland. The measurement data consisted of 180 X-ray scans (N parallel beams) of the object from different angles (J directions), out of which a subset of 10 projections were chosen for further investigation. The reason for this is that with the data from all 180 projections a simple inverse Radon transform is enough to reconstruct a nice slice image of the peanut, see Figure 2. Assuming only 10 projections were available, such a simple approach would not work as well (see Figure 3), and a more sophisticated reconstruction method would be appreciated.

A square-shaped domain was split into a $NJ \times N^2$ grid of densities $f \in \mathbb{R}^{N^2}$ representing variations in the peanut's internal structure. This grid can be described by matrix $A \in \mathbb{R}^{NJ \times N^2}$ representing unit distances that the beam passes within each grid cell for each angle. With the measurement $m \in \mathbb{R}^{NJ}$, the inverse problem is then solving[2]

$$Af = m + \epsilon \tag{1}$$

with error ϵ. The measurement m corresponds to the Radon transform of density f, defined as the line integral[2]

$$Rf(s,\theta) = \int_{x \cdot \vec{\theta} = s} f(x) dx^{\perp} \tag{2}$$

where s is the beam's normal displacement and $\vec{\theta}$ is the unit vector $(\cos\theta, \sin\theta)^T$ in the direction of the beam. The inverse Radon transform can then be applied to give the density.

The problem (1) can be solved with generalized Tikhonov regularization[2]

$$T_\alpha(m) = \arg\min_{z \in \mathbb{R}^n} \{||Az - m||^2 + \alpha||z - f^*||^2\}, \tag{3}$$

where α is the regularization parameter, found using the L-curve method[3], and f^* is a good guess. The solution $T_\alpha(m)$ provides a good balance between having a small L^2 norm as well as giving a small residual for $AT_\alpha(m) - m$.

The Eq. (3) can further be modified to give a quadratic optimization problem[2]

$$\text{minimize} \quad \frac{1}{2} z^T Q z - b^T z \tag{4}$$

with $Q = A^T A + \alpha I$ and $b = A^T m$, which can be solved e.g. iteratively with the conjugate gradient method by adding negative gradient vectors.

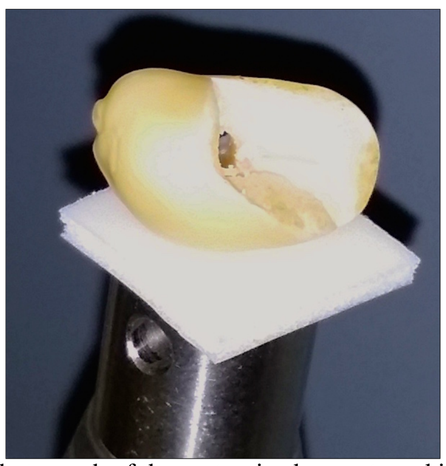

Figure 1: A photograph of the customized peanut used in this project.

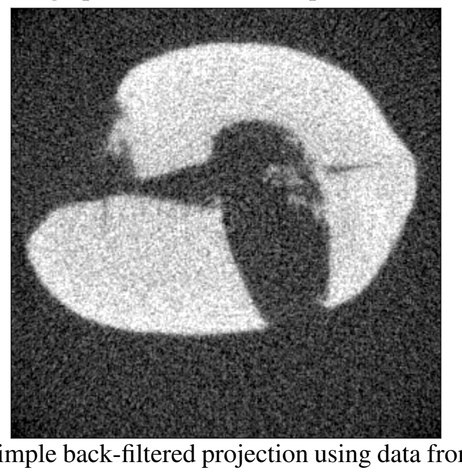

Figure 2: A simple back-filtered projection using data from 180 scans. L2-error defined as 0.

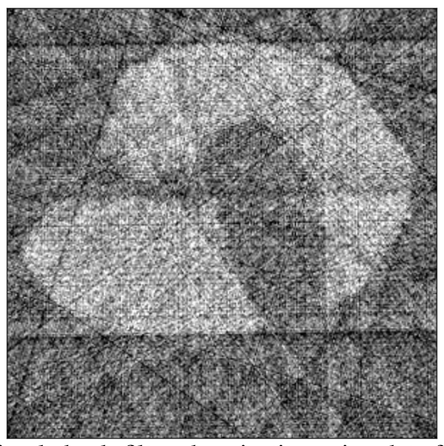

Figure 3: A simple back-filtered projection using data from 10 scans. L2-error of 127%.

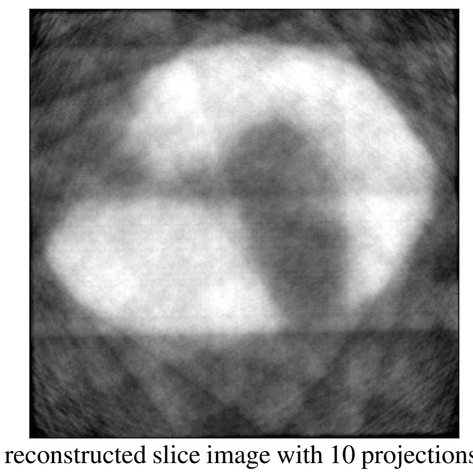

Figure 4: A reconstructed slice image with 10 projections and $\alpha \approx 3.8$ using classical Tikhonov regularization ($f^* = 0$). L2-error of 48%.

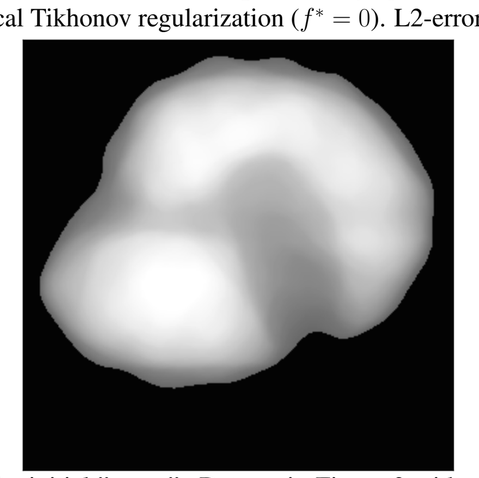

Figure 5: The initial "guess". Data as in Figure 3 with median-filter and values limited to $0.4 \times maximum$.

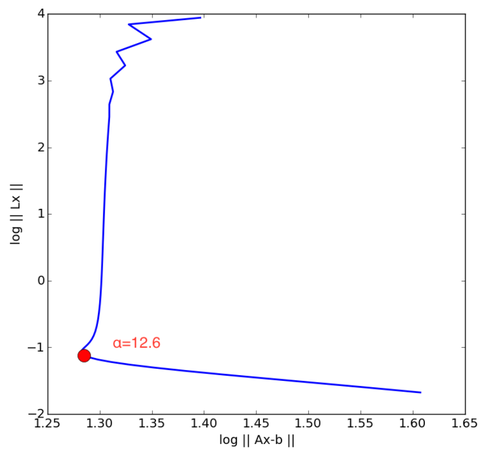

Figure 6: The terms in Eq. (3) (with $f^* = 0$ and $L = 1$) plotted with different values of the regularization parameter form an L-curve. The optimal value of $\alpha \approx 12.6$ is in the corner. The iteration is unstable with $\alpha \lesssim 0.03$.

RESULTS

A Python program was written to solve the inverse problem. First, the classical Tikhonov regularization (without initial guess) was tested, see Figure 4, which required 47 iterations until convergence (i.e. the squared residual was below 1E-9 and the iteration was stopped automatically). The L2-error against the data in Figure 2 was 48% (cf. 127% for the simple filtered back projection in Figure 3).

After that, the generalized Tikhonov regularization was applied with the initial guess shown in Figure 5, which is just the median-filtered back projection of the data with all values below $0.4 \times maximum$ set to zero. This was an attempt to reduce the effect of the star-like formations visible in the reconstructions.

The regularization parameter α was chosen based on the L-curve method, which tries to find an optimal balance between the terms in Eq. (3). The corner of the "L" coincides with its maximum curvature, which in Figure 6 corresponds to $\alpha \approx 12.6$.

The final reconstruction is shown in Figure 7, which was found after 32 iterations of the conjugate gradient method using $\alpha \approx 12.6$. The L2-error was 40%. The reconstruction was limited to non-negative values.

DISCUSSION

Due to the obvious similarity of Figures 2 and 7, it can be safely concluded that the implemented algorithms do function as intended. However, as the image is a bit blurred (no image-sharpening was applied), the fine details of the peanut, such as the tiny crack on the right-hand side, cannot be discerned with data from only 10 projections. Further testing (not shown here) suggests that they do become visible in the reconstruction only after using about 60 projections, at which point the filtered back projection is vastly superior in terms of detail and speed of processing. On the other hand, with only 6 projections the resulting reconstruction from either method is unacceptable.

The generalized Tiknohov regularization, with a good guess to start with, converged significantly faster than the classical Tikhonov regularization, and the L2-error is smaller (although the data is quite noisy). However, the initial guess (Figure 5) does not help much in reducing the star-like projection artifacts – the non-negativity constraint was more effective in this respect.

Finally, it may be argued whether the reconstruction in Figure 7 is actually any better than the simple and computationally much more efficient filtered back projection in Figure 3, as a human eye can see the same shape in both. After all, both of these figures could certainly have been tuned further in image processing software, but such actions were not taken here. It should be noted that had the density of the scanned object varied more, the comparative results could be different.

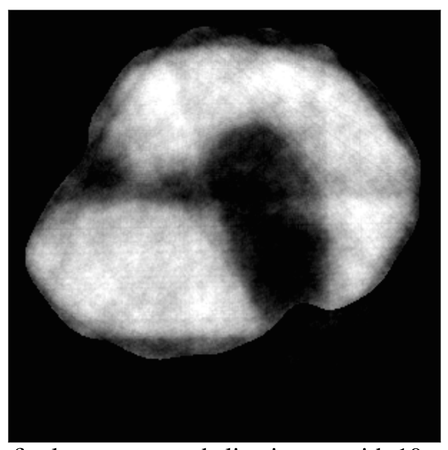

Figure 7: The final reconstructed slice image with 10 projections and $\alpha \approx 12.6$ using generalized Tikhonov regularization (f^* as in Figure 5) with naive non-negativity constraint. The solution converged after 32 iterations with an L2-error of 40%.

References

[1] `http://en.wikipedia.org/wiki/X-ray_computed_tomography`

[2] Mueller, J. L. and Siltanen, S., *Linear and Nonlinear Inverse Problems with Practical Applications*, SIAM 2012.

[3] Johnston, P. R. and Gulrajani, R. M, *Selecting the Corner in the L-Curve Approach to Tikhonov Regularization*, Transactions on Biomedical Engineering, 47, 9, 2000.