

RTG IPDE Summer School 2011

University of Washington - Seattle

François Monard*

July 11, 2011

This note summarizes what has been covered during five computer lab sessions at UW in June-July 2011 during the RTG IPDE Summer School on “X-Ray Tomography and Transport Theory”. The main topic is the Radon Transform, or more generally, generalized ray transforms on certain family of curves in two-dimensional space.

After some recalls on the usefulness of the Fast Fourier Transform for filtering signals in the fourier domain with good accuracy and reasonable cost in session 1, we go over the filtered-backprojection (FBP) algorithm in session 2, which is the most famous inversion formula for the two-dimensional Radon Transform. In session 3, we then cover the computation of the attenuated Radon transform, which is the theoretical backbone of SPECT (Single Positron Emission Computerized Tomography), as well as its numerical inversion. We then go over matters related to the transport of singularities from a Fourier-integral operators (FIO’s, a general class of operators to which the Radon transform and its adjoint belong) and wavefront sets standpoint in session 4. This perspective seems to be the most complete way to describe what happens to the singularities of an image (a.k.a. small inclusions and sharp edges) after being Radon-transformed. The transformation laws between the wavefront set of a function and that of its Radon transform allow us to understand very well what singularities are lost in limited data problems (e.g. limited angle and exterior problems), and justify the usefulness of lambda-tomography, which is a more “local” reconstruction procedure than the FBP algorithm in 2D. Finally in session 5, we study another type of curve geometry (that of the geodesics of the Poincaré disk, a typical 2d manifold of constant negative curvature) with which one can build a Radon transform and its adjoint and repeat the analyses made during sessions 2 to 4. Further, one can also intertwine geometries by using line integrals w.r.t a given geometry to produce a Radon transform, and then backproject the generated data using another geometry. Besides the artistic potential of these experiments, it is a great tool to understand the capabilities of FIO’s, as they make us able to describe analytically the distortions observed on images, provided that the expression of the curves is not too complicated.

One can also find on the course website the MatLab files that illustrate all these sessions. Detail of the instructions during the sessions are available on my page at www.columbia.edu/~fm2234/Teaching.html.

*Department of Applied Physics and Applied Mathematics, Columbia University, New York NY, 10027; fm2234@columbia.edu

Session 1: Fast Fourier Transform and applications

As a preliminary for the subsequent sections, we focus on a way of computing convolution operators and derivatives using the Fast Fourier Transform (FFT). The FFT is a fast way of computing the Discrete Fourier Transform of a vector of size N , whose expression and inversion formula are given below:

$$\begin{aligned}\tilde{x}_k &:= DFT(x)_k = \sum_{j=1}^N x_j \omega^{-(j-1)(k-1)}, \quad \omega := e^{\frac{2i\pi}{N}}, \quad 1 \leq k \leq N, \\ x_j &= IDFT(\tilde{x})_j = \frac{1}{N} \sum_{k=1}^N \tilde{x}_k \omega^{(j-1)(k-1)}, \quad 1 \leq j \leq N.\end{aligned}$$

These operators seem to require $\mathcal{O}(N^2)$ floating operations, however when N is a highly composite number (the best case being $N = 2^p$ for some p), rearranging these operations in a smart way reduce this number down to $\mathcal{O}(N \log N)$, a process known as the Cooley and Tuckey algorithm. Further, if we define the circular convolution $h = f \star g$ of two discrete vectors $f = \{f_j\}_{j=1}^N$ and $g = \{g_j\}_{j=1}^N$ as follows

$$h_j = \frac{1}{N} \sum_{k=1}^N f_k g_{[j-k]}, \quad [j-k] := j - k \pmod{N},$$

the seemingly $\mathcal{O}(N^2)$ operations that this process takes can also be reduced to $\mathcal{O}(N \log N)$ via FFT and the fact that

$$h = IFFT(FFT(f) \cdot FFT(g)).$$

Just as for a convolution, any operator that can be represented by a “Fourier multiplier”, including differentiation (ik), shift (e^{isk}), convolution by some g (\tilde{g}_k), or Hilbert transform ($-i \operatorname{sign}(k)$), can be computed using the previous method. In the x variable, this amounts to applying the same operator at the continuous level to a function that interpolates the data in the smoothest way possible while remaining the most band-limited. This interpolant is called the *spectral interpolant*. For data $f = \{f_m\}_{m=1}^N$ given on an equispaced subdivision $\{x_m\}_{m=1}^N$ of $[0, 2\pi]$, the spectral interpolant is given by

$$f(x) = \sum_{m=1}^N f_m S_N(x - x_m), \quad S_N(x) := \frac{\sin(\pi x/h)}{(2\pi/h) \tan(x/2)}. \quad (1)$$

The function S_N , a.k.a. the *periodic sine function*, is the unique trigonometric polynomial whose Fourier transform is 1 at each discrete frequency of the grid and zero elsewhere, so it makes sense that the convolution (1) leaves the discrete signal unchanged at the gridpoints. Fig. 1 gives a few examples of interpolations.

Passing to the Fourier domain, or equivalently using the spectral interpolant, is a process that has *spectral accuracy*, i.e. the order of convergence of such methods in terms of the stepsize h does not depend on the method itself, but on the regularity of the function that is being approximated, see [4] for an excellent account of these methods. See Fig. 2 for an example comparing spectral and finite-difference based differentiations.

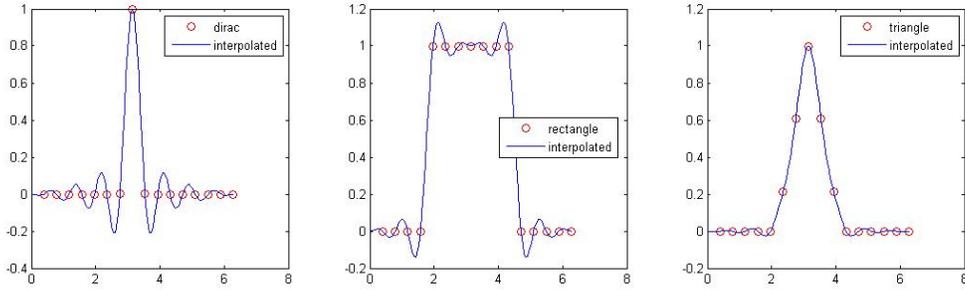


Figure 1: A few examples of functions initially defined on a 16-point grid and spectrally interpolated. The interpolant S_N itself appears on the left, since it interpolates a dirac δ on the grid.

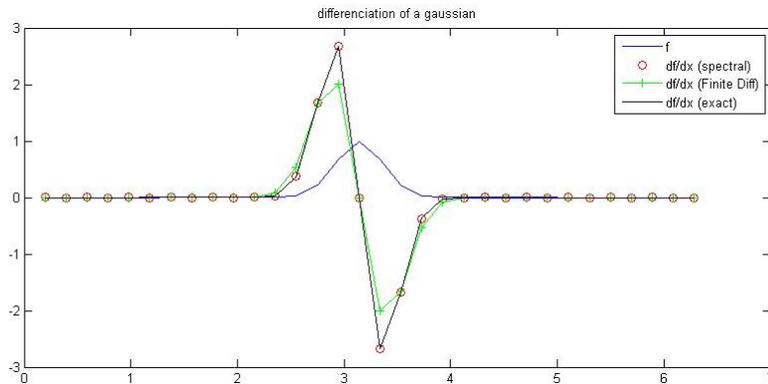


Figure 2: Comparison between spectral differentiation and centered finite difference differentiation (stencil: $[-1, 0, 1]/(2h)$). With the data discretizing a gaussian, the spectral method reaches machine precision with a very coarse discretization, while the second method converges much slower (order $\mathcal{O}(h^2)$ with h the stepsize).

Note that all these Fourier-based methods are only useful if the problem is inherently periodic, or if the data used is compactly supported away from the boundary points (which can be assimilated to the first case). In other cases, we may again refer the reader to [4] for more appropriate techniques.

For our purposes in this summer school, these FFT-based methods will be mostly used whenever one needs to apply a Hilbert transform, which happens every time one inverts the Radon transform or attenuated Radon transform. They may also come into play whenever we do image rotations if one wants to use rotations with spectral accuracy, however this part is optional.

Session 2: The Radon Transform and the Filtered-Backprojection algorithm

The purpose of this session is to implement numerically the formula

$$f = \frac{1}{4\pi} R^t \Lambda R f,$$

where R denotes the Radon Transform

$$Rf(s, \theta) = \int_{\mathbb{R}} f(s\boldsymbol{\theta} + t\boldsymbol{\theta}^\perp) dt,$$

Λ denotes the Riesz operator

$$\Lambda g = \mathcal{F}_{\sigma \rightarrow s}^{-1} |\sigma| \mathcal{F}_{s \rightarrow \sigma} g,$$

($\mathcal{F}_{s \rightarrow \sigma}$ denotes one-dimensional Fourier transform in the s variable), and R^t denotes the backprojection operator

$$R^t g(\mathbf{x}) = \int_{\mathbb{S}^1} g(\mathbf{x} \cdot \boldsymbol{\theta}, \theta) d\theta.$$

This inversion is called the *filtered backprojection* algorithm because we use a “filtered” version of the Riesz operator, whose main features will be explained later.

Radon Transform : We first compute the Radon transform by the following procedure. For each direction θ , we rotate the image so that its $\boldsymbol{\theta}$ -axis becomes aligned with the x -axis, so that the computation of its line integrals along that direction is reduced to a mere summation of entries along each row of the rotated image. The `rotateImage` function provided can either use image slantings and resamplings (see [1] and references therein), of spectral accuracy with complexity $\mathcal{O}(n^2 \log n)$, or it can just be a 4-point interpolation of accuracy $\mathcal{O}(n^{-1})$ with complexity $\mathcal{O}(n^2)$, with $n \times n$ the image size. For our purposes, it is enough to use the latter method, which is computationally cheaper. Examples of rotations as well as a computed Radon transform may be seen on Fig. 3.

Filtered backprojection : Secondly, we aim at computing a “regularized” version of the operator $R^t \Lambda$. The operator Λ is computed in the Fourier domain, since its Fourier multiplier is $|\sigma|$. This part is actually filtered (i.e. multiplied by a function with certain properties in the Fourier domain) by a function to be chosen among a few different ones. The backprojection step consists in discretizing the integral in $\theta \in \mathbb{S}^1$ by

$$\int_{\mathbb{S}^1} g(\mathbf{x} \cdot \boldsymbol{\theta}, \theta) d\theta \approx \frac{2\pi}{N_{th}} \sum_{i=1}^{N_{th}} g(\mathbf{x} \cdot \boldsymbol{\theta}_i, \theta_i),$$

for \mathbf{x} belonging to a point of the two-dimensional mesh and $\{\theta_i\}$ an equispaced subdivision of \mathbb{S}^1 . Computing $g(\mathbf{x} \cdot \boldsymbol{\theta}_i, \theta_i)$ is done by linear interpolation in the s -axis only, since we assume to know slices of the Radon Transform exactly at the directions $\{\theta_i\}$. See some reconstructions on Fig. 5 from the data that appears on Fig. 3.

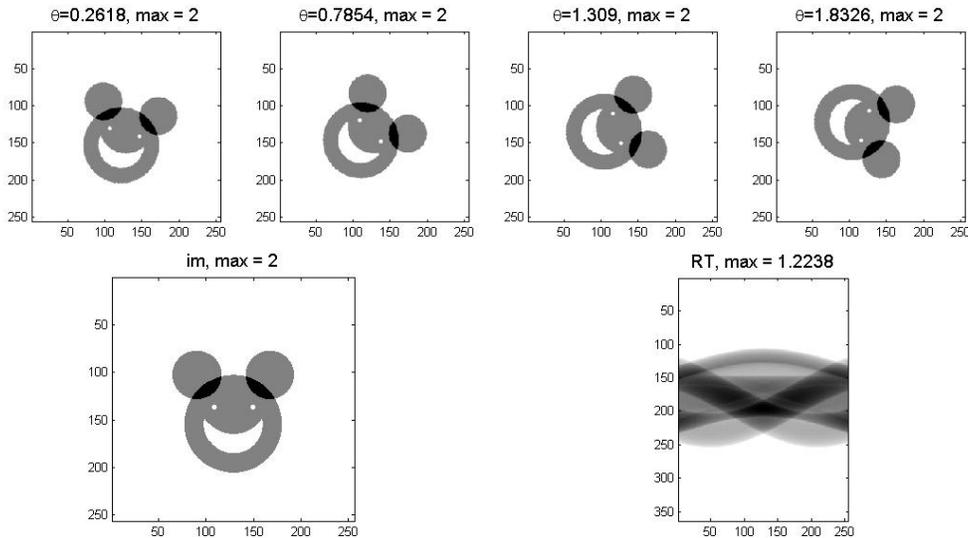


Figure 3: Top: examples of image rotations used to compute the Radon transform. Bottom: an image and its Radon Transform.

Choice of filter: We now study the impact of different filters and what the main features of these filters are. The first criterion is that the filter must be *low-pass* as a means of giving up the frequencies that have been buried into noise¹. To this end, one can play with the normalized cutoff bandwidth parameter $0 \leq d \leq 1$, an input argument of the `myIRadon` function, and notice that the smaller the d the smoother the reconstruction. The second criterion is that, because of the cutoff in bandwidth, the point-spread-function (PSF) associated to a filter (i.e., its inverse DFT) may have a strong second mode (i.e. the second bump following the first one). As can be seen for the Ram-Lak filter for instance (see Fig. 4), its second mode is pretty strong, and as a result, a point on the image will create artifacts in the vicinity of that point, an effect known as *ringing effect*. In medical imaging for instance, one would rather locate tumors at the right place rather than a bunch of metastases born out of numerical artifacts, so it is crucial that the filter chosen does not create artifacts. The ringing effect is lessened when using a Hamming filter for instance, since this filter can be seen to have a much smaller second mode.

¹We do not explore this route further here and during these sessions, although one should know that regularization techniques and the scale of ill-posedness are very recurrent themes in the field of inverse problems

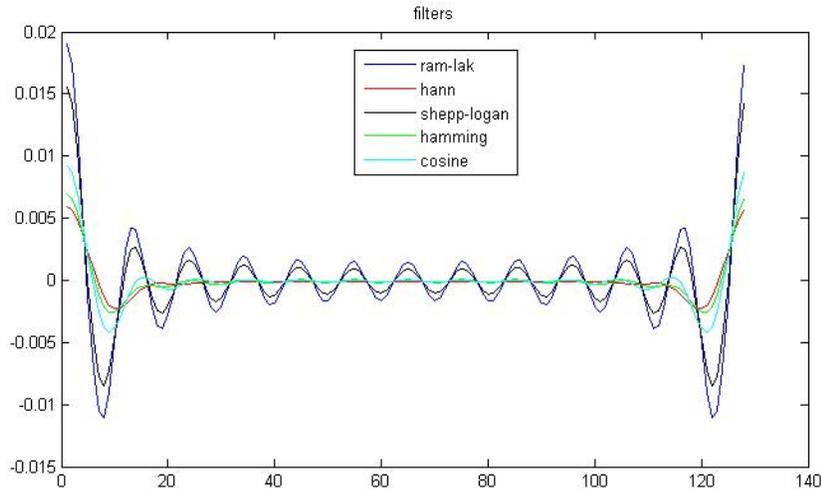


Figure 4: Point Spread Functions of the most common filters used by the FBP algorithm. Ram-Lak and Shepp-Logan have too slow a decay and a strong second mode, which is responsible for ringing effects, while cosine, Hann and Hamming lessen these effects due to the absence of their second modes.

Session 3: The attenuated Radon Transform and its inversion.

In this session, we compute a discretized version of the *attenuated Radon Transform* (AtRT) and then invert it numerically.

Computation of the AtRT: We first solve the direct problem, i.e. we compute

$$R_a f(s, \theta) := \int_{\mathbb{R}} f(s\boldsymbol{\theta} + t\boldsymbol{\theta}^\perp) e^{-D_a(s\boldsymbol{\theta} + t\boldsymbol{\theta}^\perp, \boldsymbol{\theta}^\perp)} dt.$$

Again, one may do so using image rotations: for each θ required, rotate the a and f maps so that their $\boldsymbol{\theta}$ -axis is aligned with the x -axis to get, say, images a_θ and f_θ . Then one may either compute $D_a(\cdot, \theta)$ by using cumulated sums and then compute the integral $\int_{-1}^1 f_\theta(t, y) e^{-D_a(t, \theta)} dt$ by summing entries along rows, or equivalently, solve the following ODEs along the x -axis:

$$\frac{du}{dx} + a_\theta u = f_\theta, \quad u(-1, y) = 0,$$

and store the values $u(1, y)$ that represent $R_a f(s, \theta)$ for a certain family of s 's. Fig. 6 contains the maps utilized in this experiment and Fig. 7 shows the computed AtRT, side-by-side with the non-attenuated RT.

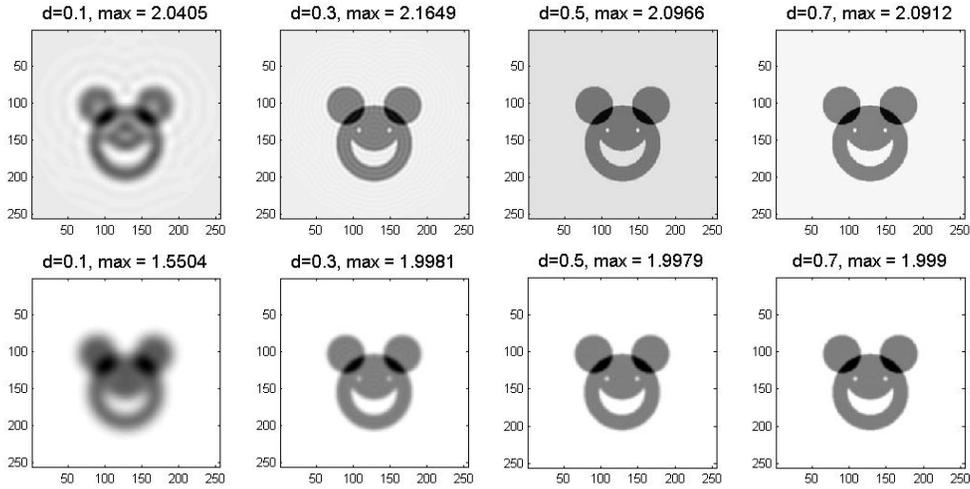


Figure 5: Effect of the cutoff parameter d and the filter. Top: using a Ram-Lak filter, the ringing effect is very apparent. Bottom: using a Hamming filter reduces these effects.

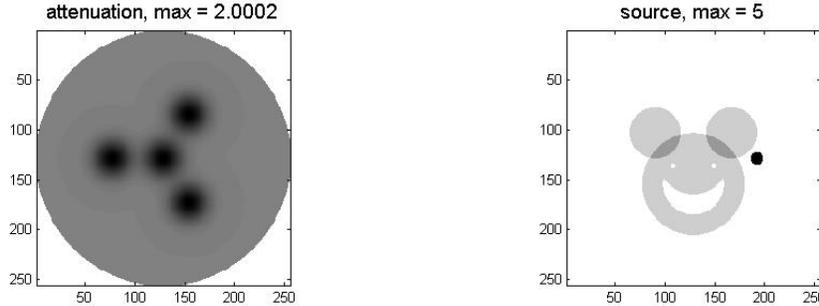


Figure 6: Attenuation and source maps.

Inversion: In the second part of this session, one inverts for f from knowledge of a and the data $g = R_a f$ by using a discrete version of the formula

$$f(x) = \frac{1}{4\pi} \Re \operatorname{div} \int_{\mathbb{S}^1} \boldsymbol{\theta} e^{D_a(\mathbf{x}, \boldsymbol{\theta}^\perp)} e^{-h(\boldsymbol{\theta}, \mathbf{x} \cdot \boldsymbol{\theta})} H(e^h R_a f) d\boldsymbol{\theta}, \quad \text{where}$$

$$h(\boldsymbol{\theta}, t) := \frac{1}{2}(I + iH)Ra(\boldsymbol{\theta}, s).$$

Following [3], a way of implementing it is to first compute the real part analytically in the first place. To this end, we write $h = h_1 + ih_2$, and calculations yield

$$f(x) = \frac{1}{4\pi} \operatorname{div} \int_{\mathbb{S}^1} \boldsymbol{\theta} e^{D_a(\mathbf{x}, \boldsymbol{\theta}^\perp)} g_a(\boldsymbol{\theta}, \mathbf{x} \cdot \boldsymbol{\theta}) d\boldsymbol{\theta},$$

$$g_a := e^{-h_1} (\cos h_2 H(e^{h_1} \cos h_2 R_a f) + \sin h_2 H(e^{h_1} \sin h_2 R_a f)).$$

Although we have not considered filtering issues during this session, one could very well replace the Hilbert Transform by a “regularized” version that includes any of the filters that have been used in the previous section.

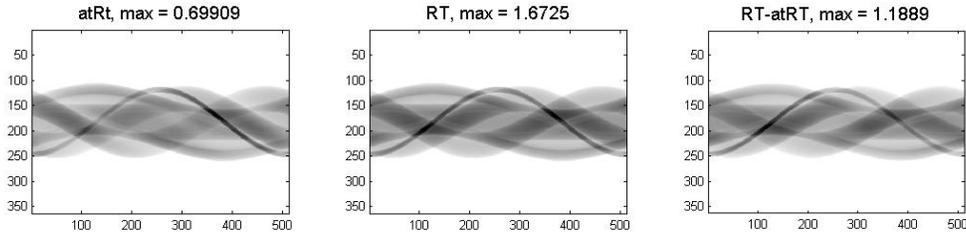


Figure 7: Computed AtRT side-by-side with the regular RT. Right: difference of both, the effect of the attenuation is clearly noticeable.

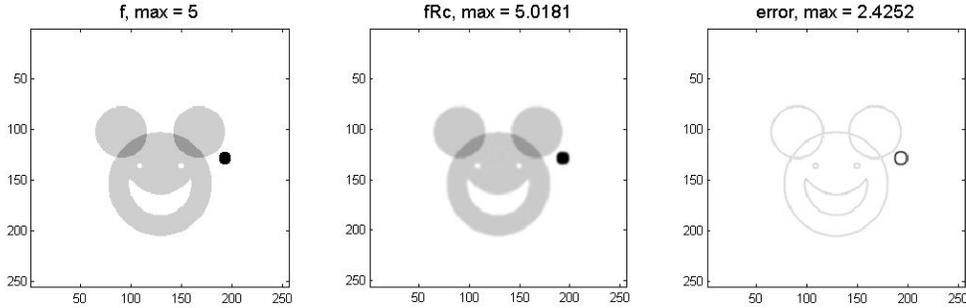


Figure 8: Inversion of the attenuated Radon transform. Error is only located at the sharp edges due to slight regularization during the inversion.

Session 4: Local tomography and limited data problems: recovering singularities

Local tomography: As was seen in morning class, the inversion of the Radon transform is a *local* formula in odd dimensions of space and a *non-local* formula in even dimensions. In particular in 2d, it is non-local. This means that in order to reconstruct a function f at a point x_0 , the knowledge of values of Rf at *all lines* is necessary, and not just at those who pass near x_0 . The term that is responsible for that in the reconstruction formula is the hilbert transform, which is a non-local operator. A way of getting this locality back is to only worry about *singularities* of f , and this is the motivation for lambda-tomography, where one computes

$$(-\Delta)^{\frac{1}{2}}f = c\Delta R^t Rf$$

instead of f itself by backprojecting the data and then computing the laplacian. This operation is now local, as it no longer involves a Hilbert transform or a $(-\Delta)^{\frac{1}{2}}$. The interest is that the singularities of $(-\Delta)^{\frac{1}{2}}f$ are located in the same place as the singularities of f (this is because $(-\Delta)^{\frac{1}{2}}$ is an *elliptic* pseudo-differential operator). As a result, if the image is mostly made of areas of constant values separated by jump discontinuities, the interfaces will still appear very clearly on the reconstructions. See Fig. 9 for results.

Incomplete data: So far, we have seen exact reconstruction formulas when we have at our disposal full measurement data, that is $Rf(s, \theta)$ for (s, θ) belonging to a fine enough

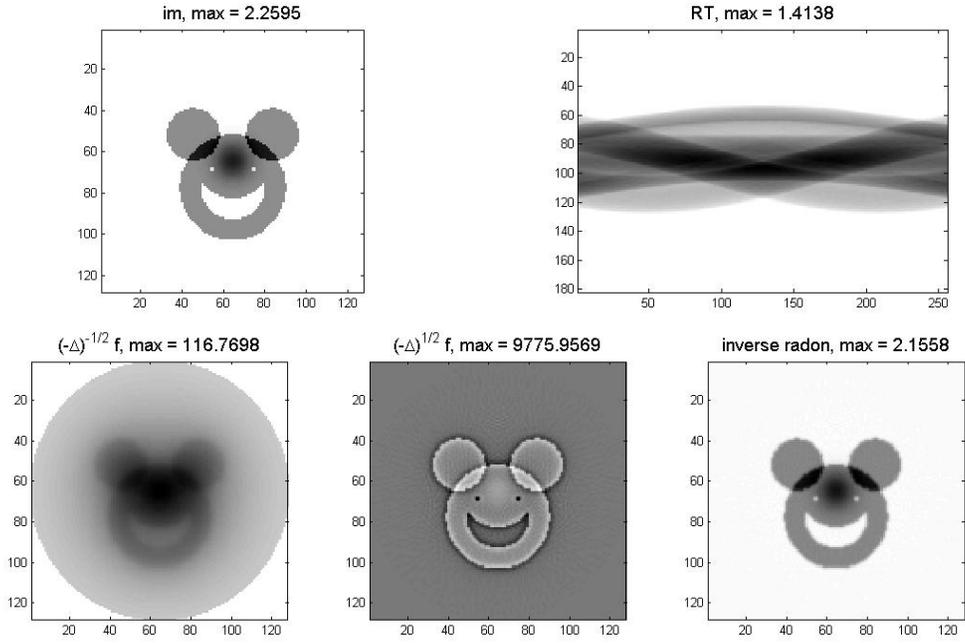


Figure 9: Lambda tomography. Top: initial image f and its Radon Transform Rf . Bottom, Left: backprojection $R^t Rf$, the singularities are back into place but they are smoothed out. Middle: Lambda tomography algorithm $\Delta R^t Rf$, the singularities are now more visible. Right: regular inversion

discretization of the whole $(-1, 1) \times [0, \pi]$ (we assume here that our object is compactly supported in the unit disk $\{\mathbf{x} \in \mathbb{R}^2, |\mathbf{x}| \leq 1\}$ and we also disregard directions $\theta \in [\pi, 2\pi]$ because of the redundancy $Rf(s, \theta + \pi) = Rf(-s, \theta)$). If we now only have half of the data, say $(-1, 1) \times [0, \frac{\pi}{2}]$, we observe that some parts in the reconstruction are completely blurred. These parts correspond to the singularities that were no longer present in the data.

In the case of the Radon transform, there is a one-to-one correspondence between the wavefront set of a function f and the wavefront set of its Radon transform Rf . Precisely, these relations are given by²

$$\begin{aligned} WF(Rf) &= \{((x \cdot \hat{\theta}, \theta); a(e_s - x \cdot \theta^\perp e_\theta)), (x; a\hat{\theta}) \in WF(f)\}, \\ WF(R^t g) &= \{(s\hat{\theta} - r\hat{\theta}^\perp, a\theta), ((s, \theta); a(e_s + r e_\theta)) \in WF(g)\}. \end{aligned} \quad (2)$$

We apply this to two types of limited data problems :

limited angle problem : Start by computing the Radon data of a given image and do a filtered projection to this data, using only, say, the directions $\theta \in [0, \frac{\pi}{2}]$. One can see that singularities whose direction was not in this range is completely missed in the reconstruction. This is clear from (2), since we see that these singularities are simply not present in the data in this case, see an example on Fig. 10.

²Note here that I use the Radon geometry notation and not the X-Ray geometry.

Exterior problem: In this case, we assume to only have the Radon Transform of f for $|s| \geq a > 0$. While not even trying to reconstruct the part of f that is in the disk $\{|\mathbf{x}| \leq a\}$ (we simply have no microlocal information on that in the data), we see that even objects outside of this area are smeared. This is because some of their singularities can only be seen by some lines $L(s, \theta)$ for $|s| < a$, which cannot happen here. Again, equation (2) provides the exact answer to what singularities can or cannot be reconstructed in this case. Results may be found on Fig 11.

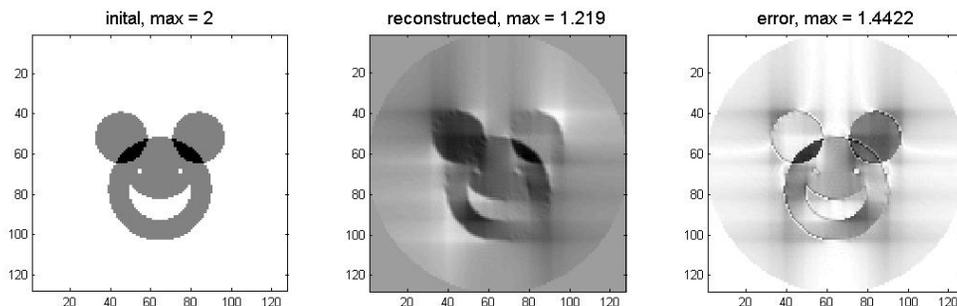


Figure 10: Limited angle problem. Backprojection using only directions $\theta \in [0, \frac{\pi}{2}]$. All singularities of direction $\theta \in [\frac{\pi}{2}, \pi]$ are blurred.

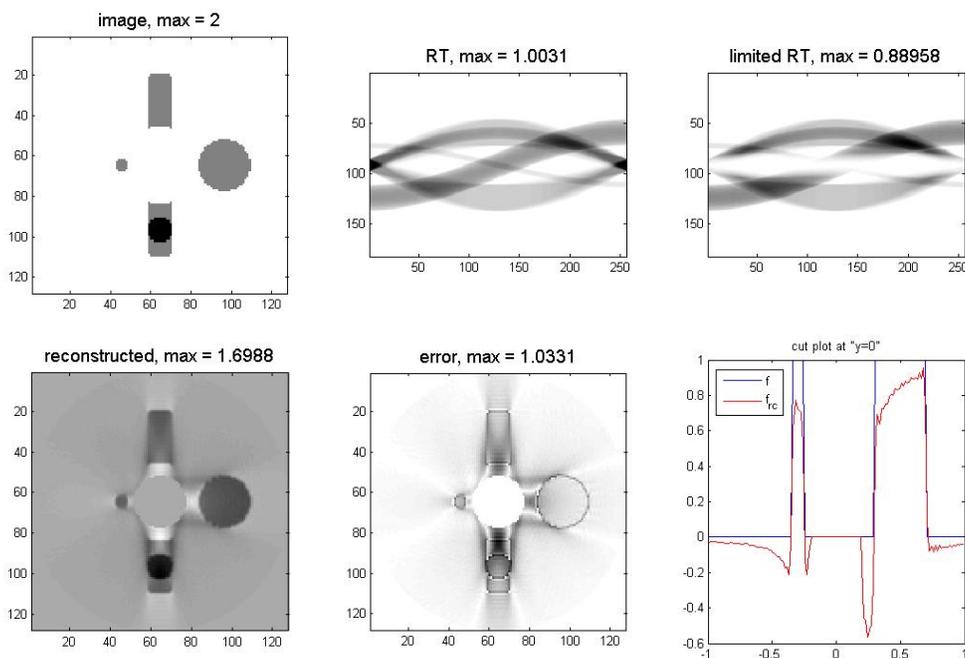


Figure 11: Exterior problem.

Session 5: Cartesian and hyperbolic geometries, fun with FIO's.

The hyperbolic geometry

During this session, we explore another possible family of curves that is not the straight line geometry, namely the curves that are geodesics of the hyperbolic disk $\{|x| < 1, x \in \mathbb{R}^2\}$ with the metric $g_{ij}(x) = \frac{1}{1-|x|^2} \delta_{ij}$. These curves are either arc circles intersecting the unit disk at right angles, or straight lines passing through the origin. A way of parameterizing these curves is given by, for given $(s, \theta) \in (-1, 1) \times \mathbb{S}^1$, $\gamma(\cdot, s, \theta)$ is the unique geodesic that passes through the point $s\theta$ whose tangent is colinear to θ^\perp . After working out the geometry, we find that for $s \in (-1, 1), s \neq 0$,

$$\gamma(t, s, \theta) = \rho(s)\theta + r(s)\alpha(\theta + \pi + t\alpha_0), t \in (-1, 1),$$

where we have defined $\rho(s) := \frac{1}{2} \left(\frac{1}{s} + s \right)$, $r(s) := \frac{1}{2} \left(\frac{1}{s} - s \right)$ and $\alpha_0(s) = \sin^{-1} \frac{1}{\rho(s)}$, and $\alpha(\beta) = (\cos \beta, \sin \beta)$ for any $\beta \in \mathbb{S}^1$. For $s = 0$, we define

$$\gamma(t, s, \theta) = -t\theta^\perp, t \in (-1, 1).$$

One may show that despite the $\frac{1}{s}$ factors, we rigorously have $\lim_{s \rightarrow 0} \gamma(t, s, \theta) = -t\theta^\perp$, which parameterizes the line segment properly.

Now in order to compute the corresponding hyperbolic Radon transform defined as

$$R_h f(s, \theta) := \int_{-1}^1 f(\gamma(t, s, \theta)) |\partial_t \gamma| dt, \quad (3)$$

we create a function `geodesic` that, for given $(s, \theta) \in (-1, 1) \times \mathbb{S}^1$, returns a few sample points

$$(x_i, y_i) = \gamma(t_i, s, \theta), \quad i = 1, \dots, N$$

of the curve $\gamma(\cdot, s, \theta)$. Then the expression (3) may be computed by means of a Riemann sum, each value of which is interpolated from the given data $\{f_{ij} = f(x_i, y_j)\}$, where for $1 \leq k \leq N$, $x_k = y_k = -1 + 1/N + 2k/N$. Note here that

$$|\partial_t \gamma| = r(s)\alpha_0(s)$$

does not depend on t . Figure 12 shows both euclidean (straight line) and hyperbolic Radon Transforms.

Backprojection in hyperbolic geometry

Secondly, in order to backproject the hyperbolic Radon data, we need a function $s(\mathbf{x}, \theta)$ that for given \mathbf{x}, θ , returns the unique geodesic $\gamma(\cdot, s, \theta)$ that passes through \mathbf{x} . In order to compute this, we split cases according to the value of $\mathbf{x} \cdot \theta$. If $\mathbf{x} \cdot \theta = 0$, then clearly $s = 0$. For $\mathbf{x} \cdot \theta \neq 0$, expressing the fact that $\mathbf{x} \in \gamma(s(\mathbf{x}, \theta), \cdot, \theta)$ yields

$$r^2 = |\mathbf{x} - \rho\theta|^2 = (\mathbf{x} \cdot \theta - \rho)^2 + (\mathbf{x} \cdot \theta^\perp)^2.$$

This yields the equation

$$1 + |\mathbf{x}|^2 = 2\rho(\mathbf{x} \cdot \boldsymbol{\theta}) \Leftrightarrow \frac{1}{s} + s = \frac{1 + |\mathbf{x}|^2}{\mathbf{x} \cdot \boldsymbol{\theta}} := 2\rho(\mathbf{x}, \theta),$$

the solutions of which are

$$s = \rho \pm \sqrt{\rho^2 - 1}.$$

The product of these two solutions is 1 thus we pick the unique solution of magnitude less than 1 and we arrive at

$$s(\mathbf{x}, \theta) = \begin{cases} \rho - \text{sign}(\rho)\sqrt{\rho^2 - 1} & \text{if } \mathbf{x} \cdot \boldsymbol{\theta} \neq 0, \\ 0 & \text{if } \mathbf{x} \cdot \boldsymbol{\theta} = 0, \end{cases} \quad \text{where } \rho(\mathbf{x}, \theta) := \frac{1 + |\mathbf{x}|^2}{2\mathbf{x} \cdot \boldsymbol{\theta}}.$$

Computing $t(\mathbf{x}, \theta)$ as well (although it's not necessary for implementing a backprojection) one must fulfill the condition

$$\rho\boldsymbol{\theta} + r\boldsymbol{\alpha}(\theta + \pi + t\alpha_0) = \mathbf{x},$$

which upon dotting with $\boldsymbol{\theta}^\perp$ gives

$$t(\mathbf{x}, \theta) = \frac{1}{\alpha_0} \sin^{-1} \left(\frac{\mathbf{x} \cdot \boldsymbol{\theta}^\perp}{r} \right) = \frac{\sin^{-1} \left(\frac{\mathbf{x} \cdot \boldsymbol{\theta}^\perp}{r} \right)}{\sin^{-1} \frac{1}{\rho}}.$$

Now the back projection operator in this case is given by

$$R_h^t g(\mathbf{y}) = \int_{\mathbb{S}^1} g(s(\mathbf{y}, \theta), \theta) w(\mathbf{y}, \theta) d\theta.$$

Provided that we have a function `hyperS` that returns $s : (\mathbf{x}, \theta) \mapsto s(\mathbf{x}, \theta)$, we are thus able to compute this backprojection via, again, Riemann sums and interpolations.

Lambda tomography in hyperbolic geometry: Although this geometry has exact (though, again, non-local) reconstruction formulas, we only care about making sure that our backprojection function brings the singularities back to their right places. To this end, we implement a method similar to lambda tomography, that is, we backproject the hyperbolic Radon transform before applying a laplacian in order to make the singularities more visible. The results may be found on Fig. 14 and it can be seen that singularities are successfully back to where they initially were.

Intertwining geometries

Now what happens when we intertwine geometries, that is when we apply a ray transform from a given geometry followed by a backprojection from another geometry. Namely, we will see what happens when to the singularities of f when we compute $\Delta R^t R_h f$ and $\Delta R_h^t R f$, see Fig. 15 for results.

Let us denote the following two functions

$$s_1(\mathbf{x}, \theta) = \mathbf{x} \cdot \boldsymbol{\theta}, \quad s_2(\mathbf{x}, \theta) = \rho - \text{sign}(\rho)\sqrt{\rho^2 - 1}, \quad \rho := \frac{1 + |\mathbf{x}|^2}{\mathbf{x} \cdot \boldsymbol{\theta}}.$$

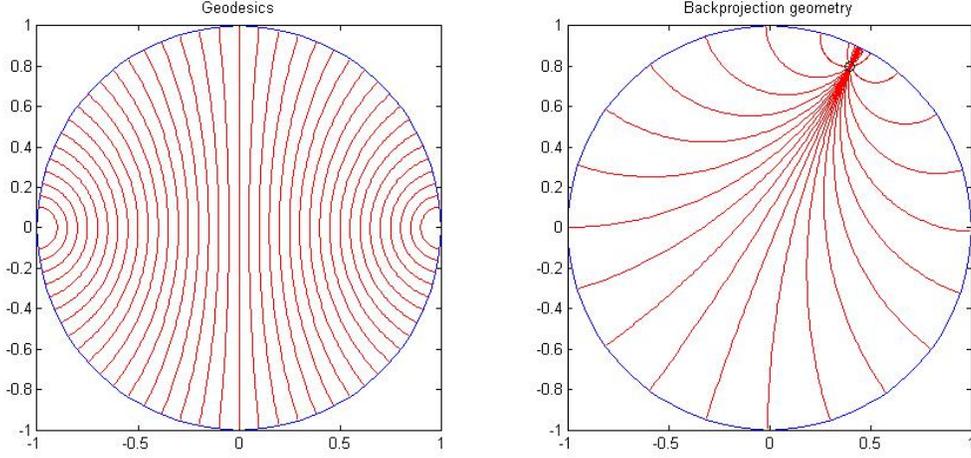


Figure 12: The hyperbolic geometry. Left: curves $\gamma(\cdot, s, \theta)$ for $\theta = \frac{\pi}{2}$ and $s = -0.9 : 0.05 : 0.9$. Right: geodesics that pass through a point that is fixed a priori

Then one can show that up to some smooth positive weight functions w_a, w_b , the compositions $R^t R_h f$ and $R_h^t R f$ may be written, for f compactly supported inside D

$$R^t R_h f(\mathbf{x}) = \int_D \int_{\mathbb{R}^2} e^{i(s_1(\mathbf{x}, \hat{\zeta}) - s_2(\mathbf{y}, \hat{\zeta}))|\zeta|} f(\mathbf{y}) w_a(\mathbf{x}, \mathbf{y}, \hat{\zeta}) \frac{d\zeta}{|\zeta|} d\mathbf{y},$$

$$R_h^t R f(\mathbf{y}) = \int_D \int_{\mathbb{R}^2} e^{i(s_1(\mathbf{x}, \hat{\zeta}) - s_2(\mathbf{y}, \hat{\zeta}))|\zeta|} f(\mathbf{x}) w_b(\mathbf{x}, \mathbf{y}, \hat{\zeta}) \frac{d\zeta}{|\zeta|} d\mathbf{x}.$$

In order to study how the wavefront set transforms after applying either of these operators, we are left to analyse the phase function

$$\phi(\mathbf{x}, \mathbf{y}, \zeta) = |\zeta|(s_1(\mathbf{x}, \hat{\zeta}) - s_2(\mathbf{y}, \hat{\zeta})).$$

We must first describe the set $\{\nabla_{\zeta} \phi = 0\}$. Expressing this condition yields the following two equations

$$s_1(\mathbf{x}, \theta) = s_2(\mathbf{y}, \theta) \quad \text{and} \quad \partial_{\theta} s_1(\mathbf{x}, \theta) = \partial_{\theta} s_2(\mathbf{y}, \theta), \quad (4)$$

where one may consider that $\hat{\zeta} = (\cos \theta, \sin \theta)$. Now when the same family of curves is used for both transform and adjoint transform, then we have $s_1 \equiv s_2$ and equations (4) imply $x = y$. That is, as expected, singularities are back to their initial locations. If $s_1 \neq s_2$, these equations becomes much more complicated and certainly not imply $x = y$. This means that the singularities will not be repropagated to where they initially came from, and we obtain all sorts of (artistic) artifacts, as may be seen from the pictures.

It may be an interesting exercise, though not straightforward, to work out the expression of the wavefront set transformation that characterizes such distortions.

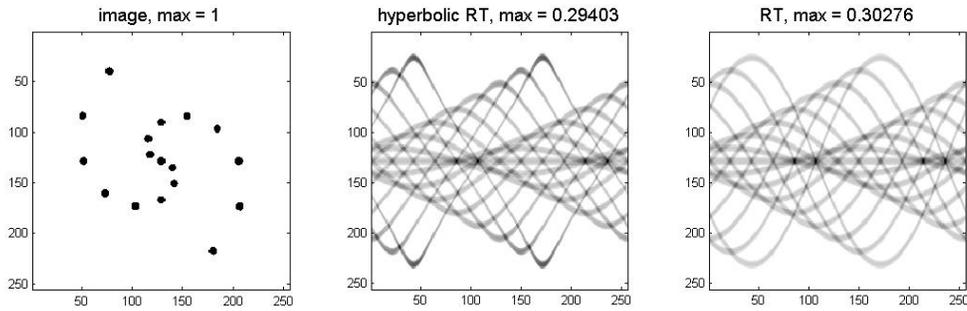


Figure 13: The euclidean and hyperbolic Radon Transforms (RT) side-by-side. Because the hyperbolic curves are almost straight near the origin, RT and hRT look similar near $s = 0$. They start differing a lot near $|s| = 1$: this is because a point of norm near 1 in the disk still generates a sinusoid in euclidean geometry while its track looks more like a sawtooth curve in hyperbolic geometry.

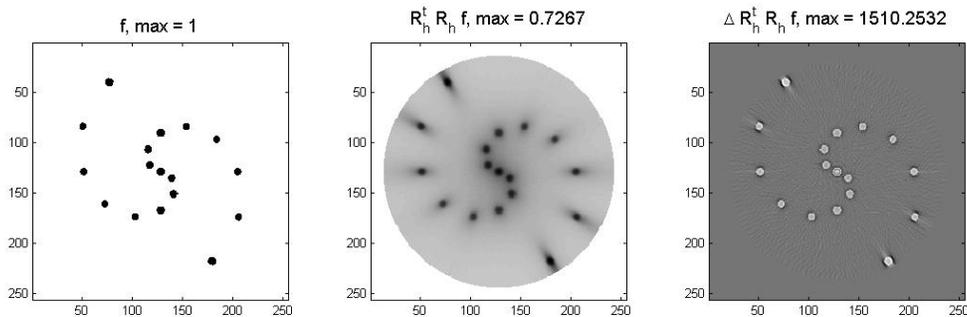


Figure 14: Lambda tomography in hyperbolic geometry: the singularities are back into place after backprojection.

References

- [1] F. MONARD, G. BAL, *An accurate solver for forward and inverse transport*, Journal of Computational Physics, Vol. 229, Issue 13, July, 2010.
- [2] F. NATTERER, *The mathematics of Computerized Tomography*, Wiley, New York, 1986.
- [3] F. NATTERER, *Inversion formulas for the attenuated Radon Transform*, Inverse Problems **17** (2001) 113-119.
- [4] L. TREFETHEN, *Spectral Methods in MatLab*, SIAM.

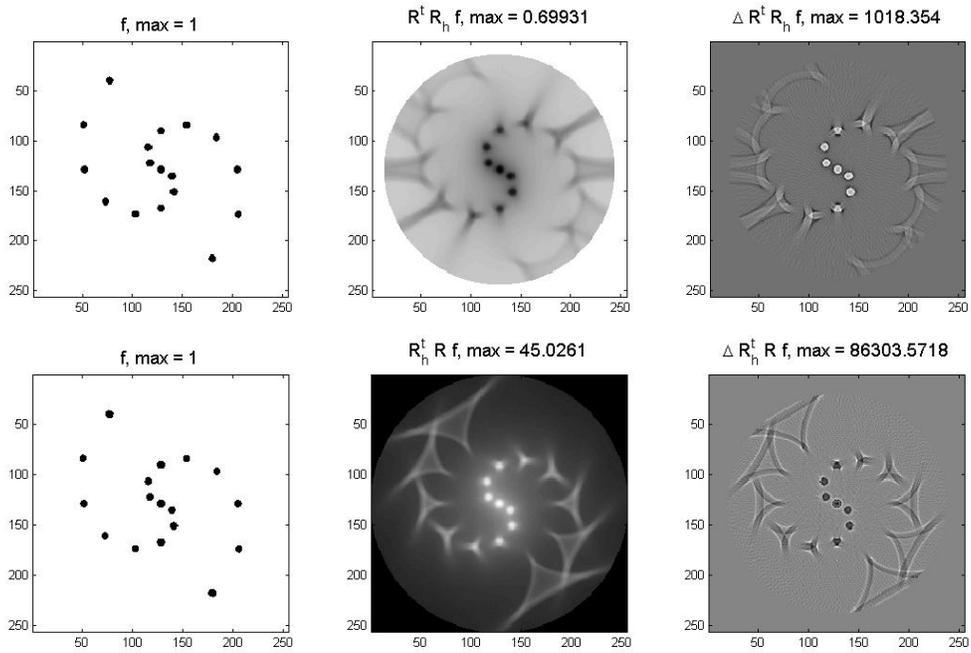


Figure 15: Intertwining geometries. Top: regular backprojection on hyperbolic Radon data. Bottom: hyperbolic backprojection on regular Radon data. In each case, a Laplacian is applied to make the singularities more visible.