

Computed Tomography for Region-of-Interest Problems with Limited Data

Nicolai André Brogaard Riis - s113200

Jacob Frøsig - s113255

DTU



Kongens Lyngby 2017

Technical University of Denmark
Department of Applied Mathematics and Computer Science
Richard Petersens Plads, building 324,
2800 Kongens Lyngby, Denmark
Phone +45 4525 3031
compute@compute.dtu.dk
www.compute.dtu.dk

Summary (English)

In this thesis the mathematical model for reconstructing cross-sectional images of deep sea oil pipes is studied for limited region-of-interest X-ray measurement data. This is motivated by a research and development project in collaboration with the Digital X-ray inspection division within the company FORCE Technology. The images are used to detect defects in the pipes such as cracks and corrosion.

The first part of the thesis provides *insight* into a mathematical model describing region-of-interest X-ray tomography. For this model it is shown exactly which singularities of a measured object are (or are not) visible in the data using a framework derived from microlocal analysis. This provides an expectation of the challenges in reconstructions from limited data.

The second part of the thesis studies *reconstruction algorithms* for the region-of-interest model. Firstly, the expected challenges of reconstructing using this model are verified numerically and additional challenges, when using standard algorithms, are shown. With the aim of overcoming these challenges a weighted frame-based sparsity penalty in a variational formulation is used to incorporate prior knowledge of the measurement geometry and object. This method is shown to include only significant details of the object that are visible in the data and is well-represented by the frame.

In the third and last part of the thesis this insight is used on *real measurement data* provided by FORCE from a prototype set-up. The expected challenges of ROI are shown to hold for real data. Hence, an exterior measurement geometry is proposed as an alternative to ROI yielding more singularities of the object in the data. The weighted frame-based methods are shown to provide reliable reconstructions on this type of data.

Summary (Danish)

I denne afhandling undersøges en matematisk model til at rekonstruere et tværsnitsbillede af undervands-olierør for ufuldstændige data med begrænset rekonstruktions område (ROI). Afhandlingen er motiveret af et R&D projekt i samarbejde med den digitale røntgen inspektions afdeling hos FORCE Technology. Tværsnitsbillederne bliver brugt til at detektere defekter i olierørene, såsom revner og korrosion, således at de kan repareres proaktivt.

Den første del af afhandlingen giver *indsigt* til en matematisk model der beskriver ROI røntgen tomografi. Det er vist for denne model, ved hjælp af mikrolokal analyse, præcis hvilke singulariteter af et målt objekt der er, eller ikke er, synlige i dataene. Dette giver en forventning til hvilke udfordringer, der vil forekomme i rekonstruktioner for ROI modellen.

Den anden del af afhandlingen undersøger *rekonstruktionsalgoritmer* beregnet til ROI modellen. Først er de forventede udfordringer verificeret numerisk og det er vist at, når standard algoritmer er benyttet, forekommer der yderligere udfordringer. Ved målsætningen om at overkomme de ovenstående udfordringer, udvikles en vægtet *frame-based* rekonstruktionsalgoritme, der kan inkorporere viden om målegeometrien og objektet. Denne metode viser sig kun at inkludere signifikante detaljer af objektet fra dataene, som repræsenteres godt i basen (frame), hvilket resulterer i gode rekonstruktioner.

I den tredje og sidste del af afhandlingen bruges indsigten og algoritmerne på rigtige måledata givet af FORCE fra en prototype måleopstilling. For rigtig data vises de ovenstående udfordringer i ROI sig også gældende og en alternativ målegeometri er da foreslået, som måler flere singulariteter af objektet i dataene. Den vægtede *frame-based* algoritme giver gode rekonstruktioner på denne type af data.

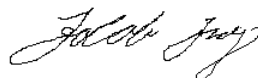
Preface

This thesis was prepared at DTU Compute in fulfilment of the requirements for acquiring a M.Sc. in Mathematical Modelling and Computation at The Technical University of Denmark (DTU). The thesis is jointly written by students Nicolai André Brogaard Riis and Jacob Frøsig representing a workload of 30 ECTS points for each student. The work was carried out from September 1st 2016 to February 21st 2017 under the supervision of Professor Per Christian Hansen, Associate Professor Yiqiu Dong and Ph.D. candidate Rasmus Dalgas Kongskov, whom all are from DTU Compute. A significant body of the work was done in collaboration with FORCE Technology, where the two students were employed and co-supervised by Ph.D. Torben Klit Pedersen and Ph.D. Arvid Böttiger.

Lyngby, 21. February 2017



Nicolai André Brogaard Riis
s113200



Jacob Frøsig
s113255

Acknowledgements

We wish to express our gratitude towards a number of people that have helped shape the direction of the thesis.

First and foremost we thank our supervisors Per Christian, Yiqiu and Rasmus, without whom this project would not have been possible. We are especially grateful for the weekly (sometimes hours long) meetings discussing every bit of detail in the project, in addition to the rich feedback on drafts for the thesis. Per Christian has followed our progress from the 4th semester at DTU. He has given us the opportunity for a great university degree with many special courses and has even provided us with detailed guides for our vacation travels, as well as a myriad of other things, for which we are eternally grateful.

We are also grateful to our co-supervisors at FORCE Technology: Torben and Arvid in addition to our colleagues Jan, Peter, Finn, Lars and Lars. The project was raised to a new level by the possibility for testing on real measurement data. We have had many great talks on the practical aspects of X-ray measurement data acquisition and various other topics.

We were very fortunate that both Jürgen Friel from OTH Regensburg and Todd Quinto from Tufts University were visiting Denmark at the time this project was carried out. Both are experts in the continuous model of limited data tomography and have had an enormous impact on the direction of the project. We are extremely grateful for the fruitful discussions, providing insightful answers to some of the most difficult questions raised in the thesis. The atmosphere at our meeting have always been relaxed and filled with humour, while maintaining an advanced level of academic content.

We also wish to thank many of the faculty members at DTU compute. Jacob Lemvig for his knowledge on shearlets, Martin Skovgaard Andersen for his insight into optimisation algorithms and Kim Knudsen with comments on the functional analysis aspects of the project.

Finally, we would like to thank Josephine and Margit for their proof reading of the thesis and general understanding of our absence caused by long days at DTU.

List of symbols

The following is a list of symbols used in the thesis. It is not a complete list of all symbols, but rather a list of the most commonly used ones. If a symbol is not in the list, it is clearly defined before use.

Symbol	Description
A	Matrix representation of the discrete Radon transform.
A_Ω	Matrix representation of the discrete ROI Radon transform.
A_E	Matrix representation of the discrete exterior Radon transform.
α	Regularisation parameter.
\mathbf{b}^δ	Vector representation of noisy discrete sinogram.
\mathbf{b}_Ω^δ	Vector representation of noisy discrete region-of-interest sinogram.
\mathbf{b}_E^δ	Vector representation of the noisy discrete exterior sinogram.
C^∞	Continuous and infinitely differentiable functions.
C_c^∞	C^∞ with compact support.
$C_{\mathcal{R}}$	Canonical relation for Radon transform.
χ_Ω	Region-of-interest based mask.
\mathcal{C}_Ω	Information based mask.
\mathbf{c}	Vector representation of frame coefficients.
c_μ	The μ th frame coefficient.
\mathcal{D}	Set of distributions.
\mathcal{D}_c	Set of distributions with compact support.
D_M	Dilation operator.
η	Additive white noise.
\mathbf{e}	Vector with additive white noise.
\mathbf{e}^δ	Vector with additive white noise with relative noise level δ .
\mathcal{F}	Fourier transform.
\mathcal{F}^{-1}	Inverse Fourier transform.
f	Attenuation coefficient of continuous object.
$g(\theta, s)$	Sinogram.
$g^\delta(\theta, s)$	Noisy sinogram.
g_Ω^δ	Region-of-interest sinogram with noise.

Symbol	Description
\mathcal{H}	2D Haar wavelet system.
I	Intensity of X-ray after having passed through an object.
I_0	Intensity of X-ray before having passed through an object.
i	Imaginary unit $\sqrt{-1}$.
L^1	The space of absolutely integrable functions.
L^2	The space of square integrable functions.
Λ	Filtering in frequency domain.
ℓ	Line describing the trajectory of an X-ray.
m	Number of discrete projections.
\mathcal{M}	Discrete index space.
M	Number of frame elements or dilation matrix.
n	Number of pixels in $N \times N$ image.
MI	Mutual information measure.
MI_Ω	Mutual information measure on region-of-interest.
Ω	The region-of-interest.
$p(\ell(\theta, s))$	Projection along the line $\ell(\theta, s)$.
Φ	Frame system.
φ	Frame element.
ϕ	1D or 2D Haar wavelets.
ψ	1D or 2D Shearlets.
\mathcal{P}	Pseudo-differential or Fourier integral operator
$p(x, y, \xi)$	The symbol of pseudo-differential operators.
ΨDO	Pseudo-differential operator.
$\phi(x, y, \xi)$	Phase function of Fourier integral operator.
$\mathcal{R}f$	The Radon transform of f .
\mathcal{R}^*	Continuous back projection operator.
\mathcal{R}^{-1}	Inverse Radon Transform
\mathcal{R}_Ω	Region-of-interest Radon transform.
RE	Relative error measure.
RE_Ω	Relative error measure on region-of-interest.
S^n	Unit sphere in \mathbb{R}^n .
\mathcal{S}	Schwartz space.
$\text{ssupp}(f)$	Singular support set of f .
$\Sigma(f)$	Frequency set of f .
$\Sigma_x(f)$	Local frequency set of f at x .
SH	Shearlet system.
\mathcal{T}_Φ	Analysis operator for frame system Φ .
\mathcal{T}_Φ^*	Synthesis operator for frame system Φ .
T_t	Translation operator.
V_Ω	Measure of weighted proportional support.
$\text{WF}(f)$	The wavefront set of f .
\mathbf{x}	Vector representation of discrete object.
\otimes	Outer product.
$\hat{\cdot}$	Fourier transform of \cdot .

Contents

Summary (English)	i
Summary (Danish)	iii
Preface	v
Acknowledgements	vii
List of symbols	ix
1 Introduction	1
1.1 Computed tomography	3
1.2 Region-of-interest tomography	5
1.3 Aim of the thesis	6
1.4 Structure of the thesis	6
2 Overview of computed tomography	7
2.1 Modelling X-ray tomography	7
2.1.1 Pencil beam attenuation	8
2.1.2 Continuous model	10
2.1.3 Discrete model	11
2.1.4 Implementation of model, simulation & testing	14
2.1.5 An immediate reconstruction approach	15
2.2 Inverse problems & regularisation	16
2.2.1 Inverse problems	17
2.2.2 Ill-posedness of inverse problems	17
2.2.3 Regularisation	18
2.2.4 Applying regularisation to computed tomography	19
2.3 Microlocal analysis	25

2.3.1	Singular support & wavefront Set	25
2.3.2	Pseudo-differential operators	28
2.3.3	Fourier integral operators	30
2.3.4	Application to tomography	32
3	Region-of-interest tomography	35
3.1	The region-of-interest model	35
3.2	An immediate reconstruction approach	37
3.3	Reflections	40
3.3.1	Added artefacts	40
3.3.2	Missing structure	42
4	Frame-based variational formulation	43
4.1	Framework	43
4.1.1	Frames	44
4.1.2	Wavelets	45
4.1.3	Shearlets	46
4.2	The variational formulation	51
4.2.1	Choice of parameters	52
4.2.2	Fast Iterated Soft-Thresholding Algorithm (FISTA)	54
4.3	Weighted wavelet sparsity penalty	57
4.3.1	Numerical experiments	58
4.4	Weighted shearlet sparsity penalty	64
4.4.1	Numerical experiments	65
4.5	Reflections	70
4.5.1	Choice of frame	70
4.5.2	Artefact removal	70
5	Computed tomography on deep sea oil pipes	73
5.1	Prototype set-up	73
5.2	Gathering & preprocessing of data	75
5.3	Comparing forward simulation with obtained data	77
5.4	Reconstructions	79
5.5	Exterior tomography	81
5.6	Reflections	88
5.6.1	Combining measurement data	88
5.6.2	Increased detector size for exterior measurements	91
6	Conclusion & future work	93
A	Appendix	95
A.1	Alternating Direction Method of Multipliers (ADMM)	95
A.2	Additional theory	97
	Bibliography	101

Introduction

Main author: Nicolai André Brogaard Riis.

Co-author: Jacob Frøsig.

This thesis is motivated by an industrial research and development project in collaboration with the Digital X-ray inspection division within the company FORCE Technology¹. The objective of the project is to inspect underwater oil pipes for defects that are not visible from visual inspection of the pipe. Today these defects are typically found using a diver with an ultrasound device to create an image of the pipe. However, because of the physical limitations of ultrasound, it is often necessary to remove part of the pipes outer layers to get a good enough image. This fact, combined with difficulties of having a diver far below sea level, makes investigation of such pipes an expensive and time consuming endeavour. The proposal from FORCE is to replace the diver and ultrasound with a remotely controlled X-ray inspection device, illustrated in Figure 1.1, generating accurate images of the pipe in real time and relay them to a vessel above for processing.

To generate an accurate image reconstruction of a pipe using X-rays, it is necessary to combine several *projections* into one image as done in *Computed Tomography* scanners (CT-scanners) today. Because the pipe consist of material that have high absorption coefficients, a powerful source with a narrow beam is necessary. Thus the beam never illuminates the whole pipe as illustrated in Figure 1.1. This is in contrast to regular CT where the beam always illuminates the entire object. This poses challenges, as we discover in the thesis. Before going further, we set the stage for the thesis and explain key concepts of CT in a broad perspective.

¹Force Technology: Digital X-ray inspection






Figure 1.1: Illustration of the X-ray inspection device used on a section of an oil pipe. Illustrated on the image is the area illuminated by X-rays from the source. Note that the pipe is not fully illuminated by the narrow beam. This illustration is graciously provided by FORCE.



(a) Projection data generated from measurements of an oil pipe.



(b) A CT reconstruction from the measured projection data.

Figure 1.2: CT reconstruction: The measured projection data are used to reconstruct a 2D cross-section of an oil pipe.

1.1 Computed tomography

Computed tomography (CT) has evolved into an indispensable imaging tool in both clinical and industrial applications. It is the method of *reconstructing* an image of a particular object from measurements of its "shadows". The measurements are made by detecting the number of photons (from an X-ray source) that pass through various points of the object, thus determining the "shadow" that is cast by the object from a particular direction. An example of a CT reconstruction on a pipe is illustrated in Figure 1.2b. The image is reconstructed from the measurement data shown in Figure 1.2a. Note that the colours are inverted such that brighter areas are material that lets fewer photons through.

CT has a wide array of applications ranging from the well-known CT-scanners in medical imaging, to synchrotron X-ray tomography used for fundamental research in material science and engineering. Since its invention in the 1960s and 1970s, for which Allan M. Cormack and Godfrey Hounsfield was awarded the Nobel prize in 1979, the field of CT has been subject to much active research.

In physics there are several ways to model the interaction between X-rays and matter, depending if one models the X-ray as single particles, waves or something else. These physical models are approximations to reality and depending on the application one model might be better suited than other models. In mathematical modelling one takes such a physical model and strips away the complexity even further leaving only the most essential part of the interaction

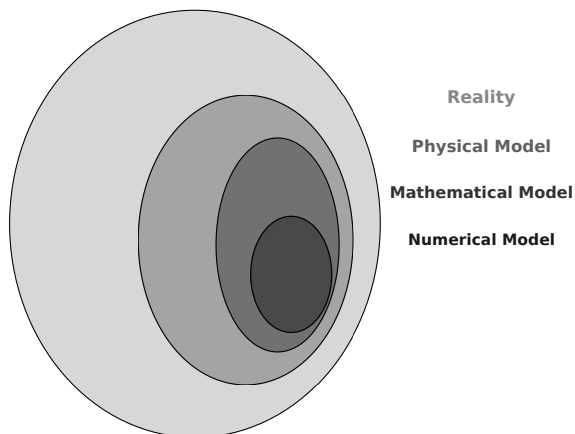


Figure 1.3: Illustration of the approximation hierarchy made by modelling reality. Each step approximates the complexity in the previous, thus simplifying the model further.

behind. Finally, when considering most real problems it becomes necessary to use computers for calculations. This is done by making a numerical model, approximating the mathematical model. The diagram in Figure 1.3 illustrates the idea of modelling reality in this way. In the end it is remarkable that we can get any usable results from real data using this hierarchy of approximations. It speaks to the importance of modelling real problems in the best way possible, which is an essential concept of this thesis.

From the point-of-view of mathematics, CT has traditionally been modelled by line integrals, simplifying the physical X-ray model by a single intensity pencil beam disregarding scattering, beam-hardening and other physical phenomena. The corresponding mathematical problem is then to reconstruct an unknown function from knowledge of its line integrals. This problem was studied by Johann Radon in 1917. Radon derived an explicit inversion formula for the line integral operator, thus providing an analytical solution to this mathematical model.

The challenge is that the inversion formula derived by Radon assumes the data in the model to be complete, i.e., the line integrals are continuous and available from all directions. In any real application this assumption does not hold and reconstructions typically show artefacts when the data becomes too limited. Additionally, real CT measurements are corrupted by noise from different sources such as measurement errors or background radiation. This is not included in the model and causes serious problems since the inversion formula is unstable, i.e, small perturbations (noise) in the data can cause huge errors in

the reconstructed image. Despite this, the analytical inversion formula is still the main tool when reconstruction images in CT. To avoid the above issues, the typical approach is to change the *filtering* step to minimise the impact of noise corruption. In addition, data is sometimes made less incomplete by sampling several times from the same data, or simply by interpolating in between measured data points. However, this approach does not always reach the desired result because of dose limits, measurement geometry or other physical restrictions. This challenge has given rise to a large body of theory that essentially use a new model of CT, that takes the discreteness and noise in measurement data into account, using algorithms that are more stable and/or utilise prior knowledge of the object to get better image reconstructions with worse quality data.

1.2 Region-of-interest tomography

The X-ray inspection device proposed by FORCE fits into the category of *Region-of-interest* (ROI) tomography. In ROI tomography the object under study is only partially illuminated by X-rays from all directions. This is caused by a measurement geometry that rotates the source and detector around the centre of an object. This means we have an interior region that is illuminated from all source and detector positions and an exterior region that is not. The interior region is called the region-of-interest. This is typically the case when the span of the rays, or the length of detector, is too small. This was illustrated on the rendering of the measurement set-up proposed by FORCE in Figure 1.1.

Indeed, if one is interested in the entire object, ROI tomography is a problem of limited data, since each measured direction does not carry information about the entire object. In fact, the quality of the information at a given point is worse the further away the point is from the ROI.

As we discuss in Chapter 3, using the regular CT model for ROI CT is not feasible, and depending on which algorithm one uses to create reconstructions, strong artefacts in the images are generated. This calls for insight into the ROI CT model and suited reconstruction methods.

1.3 Aim of the thesis

The goal of the project is three-fold: 1) Gain insight into the inherent difficulties and limitations of the ROI problem proposed by FORCE. 2) Study current state-of-the-art methods for ROI CT, determining which are best suited for this particular problem. 3) Apply the insights to real measurement data. To this end, microlocal analysis is used to determine which features of an object are visible from ROI-measurements, showing what we can expect to see in reconstructions. Furthermore, we study a general class of reconstruction methods that regularise by decomposing the object sparsely into a chosen frame. Prior knowledge of the sampling method and object can be incorporated by choosing suited frames and penalising weights.

1.4 Structure of the thesis

The thesis is structured as follows:

Chapter 1 is the introduction.

Chapter 2 gives an overview of the mathematical model for computed tomography, describes regularisation techniques for reconstructing images in CT and uses microlocal analysis to describe the propagation of singularities in tomographic transforms.

Chapter 3 develops a mathematical model for Region-of-Interest tomography and illustrates how typical CT reconstruction algorithms fare on this model. Furthermore, comments are made on which features are visible in ROI data using microlocal analysis.

Chapter 4 describes weighted frame-based reconstruction algorithms, which show promising results on the Region-of-Interest tomography model.

Chapter 5 presents results of using the algorithms developed in the thesis on real CT data generated from prototype tests on oil pipes by FORCE.

Chapter 6 concludes on the thesis as a whole and presents future work.

CHAPTER 2

Overview of computed tomography

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

This chapter gives an overview of the principles behind the mathematics of computed tomography. First the mathematical model of X-rays tomography is studied, showing how it can be viewed as an inverse problem. Regularisation techniques are then used to form standard reconstruction algorithms for this model. Finally the model is analysed using microlocal analysis to determine how singularities propagate in the related tomographic transforms.

2.1 Modelling X-ray tomography

Main author: Jacob Frøsig

Co-author: Nicolai André Brogaard Riis

The main purpose of this section is to develop a general mathematical model describing tomographic imaging to be considered throughout the thesis. To this end, we start by explaining how the attenuation of X-rays through objects are modelled as an integral over a straight line using Beer's Law.




Figure 2.1: Illustration of the pencil beam model describing the trajectory of an X-ray through an object.

2.1.1 Pencil beam attenuation

We model the attenuation of X-rays through an object as a collection of pencil beams. The intensity attenuation of these pencil beams is modelled by a measure describing the proportion of photons passing through the object. The simplified physical set-up is illustrated in Figure 2.1, where the pencil beam, indicated in red, follows a trajectory through an object described by its attenuation coefficients denoted f . The trajectory is uniquely determined by the signed distance to the origin, $s \in \mathbb{R}$ and the normal unit vector, $\theta \in S^1$ as

$$\ell(\theta, s) = \{\theta s + \theta^\perp t \in S^1 \times \mathbb{R} \mid t \in \mathbb{R}\},$$

where θ^\perp is the unit vector perpendicular to θ .

A physical mechanism leading to attenuation of intensity through a homogeneous object at a point, x , is modelled by a single attenuation coefficient, say $f(x) = f \in \mathbb{R}$. The X-ray intensity, I , after passing a distance of Δx through an object following the straight line, ℓ , is determined by

$$I(x + \Delta x) = I(x) - f(x)I(x)\Delta x. \quad (2.1)$$

Rearranging (2.1) and taking the limit reveals

$$\lim_{\Delta x \rightarrow 0} \frac{I(x + \Delta x) - I(x)}{\Delta x} = \frac{dI}{dx} = -f(x)I(x). \quad (2.2)$$

With the aim of isolating the proportional intensity, we separate the variables on both sides of the right hand equality in (2.2) revealing $dI/I(x) = -f dx$ and then integrating both sides to get

$$\int \frac{1}{I(x)} dI = -f \int dx,$$

or

$$\ln |I| = -f \cdot x + C.$$

Since the intensity is a non-negative quantity, $|I| = I$, taking the exponential of both sides gives $I(x) = e^{-fx+C}$. Hence, by denoting the initial intensity as $I(0) = I_0$ we have

$$I(x) = I_0 e^{-fx},$$

also known as Beer's Law. Throughout, we model the attenuation by Beer's law, excluding the scattering effect along with several other physical phenomena to simplify our mathematical model. The intensity coefficient, f , depends on type of material and density. The assumption of having one attenuation coefficient through an object reflects most real objects quite poorly and we are motivated to include a location dependent $f(x)$ to our model. Hence we describe the final intensity of a beam having passed through $\ell(\theta, s)$ by

$$I(\theta, s) = I_0 e^{-\int_{\ell(\theta, s)} f(t) dt},$$

or equivalently

$$-\ln \left(\frac{I(\theta, s)}{I_0} \right) = \int_{\ell(\theta, s)} f(t) dt.$$

Here the left hand side is calculated from the intensities, I_0 and I of the beam along the line, $\ell(\theta, s)$, before and after having passed through the object respectively. We denote

$$p(\ell(\theta, s)) = \int_{\ell(\theta, s)} f(t) dt \quad (2.3)$$

as a *projection* along the line $\ell(\theta, s)$. Throughout this thesis, we use (2.3) as the model for a single beams attenuation through an object.

2.1.2 Continuous model

The Radon transform, \mathcal{R} , describes all projections for a given object as follows:

DEFINITION 2.1 The 2D Radon transform $\mathcal{R}f : S^1 \times \mathbb{R} \rightarrow \mathbb{R}$ of a function $f \in \mathcal{S}(\mathbb{R}^2)$ is defined by

$$(\mathcal{R}f)(\theta, s) := \int_{\ell(\theta, s)} f(x) \, d\sigma(x) = \int_{\mathbb{R}} f(\theta s + \theta^\perp t) \, dt = p(\ell(\theta, s)).$$

Here $\mathcal{S}(\mathbb{R}^2)$ is the Schwartz-space of \mathbb{R}^2 described in Definition A.1. The *sino-gram*, subsequently denoted by g , is the collection of all projections.

Given the Radon transform, we form the first mathematical model for our problem.

MODEL 2.2 (CONTINUOUS CT) Let \mathcal{R} be the Radon transform from Definition 2.1. We model the attenuation of X-rays through an object described by its attenuation coefficients, f , as follows

$$g(\theta, s) = (\mathcal{R}f)(\theta, s) \quad \text{for } (\theta, s) \in S^1 \times \mathbb{R}.$$

Here $g(\theta, s) = -\ln(I(\theta, s)/I_0)$ is an attenuation measure where I_0 and $I(\theta, s)$ are the intensities of an X-ray, on the line $\ell(\theta, s)$, before and after passing through the object, respectively.

A naive approach to retrieve f would be a *back projection* given by:

DEFINITION 2.3 Given a function $g \in L^1(S^1 \times \mathbb{R})$, we define the back projection, \mathcal{R}^* , as

$$(\mathcal{R}^*g)(x) = \int_{S^1} g(\theta, x \cdot \theta) \, d\theta.$$

However, the back projection does not form an inverse for the Radon transform and we look to the theory developed by Radon: With the aim of constructing a two-dimensional function from its line integrals, Radon derived an inverse for the operator in Definition 2.1 without considering the underlying practical applications:

THEOREM 2.4 For $f \in \mathcal{S}(\mathbb{R}^2)$ and the 2D Radon transform, $\mathcal{R}f(\theta, s)$, $\theta \in S^1$, $s \in \mathbb{R}$, we have the inversion formula:

$$\begin{aligned} f(x) &= (\mathcal{R}^{-1}\mathcal{R}f)(x) = (\mathcal{R}^* \Lambda_s \mathcal{R}f)(x) \\ &= \frac{1}{2} (2\pi)^{-3/2} \int_{S^1} \int_{-\infty}^{\infty} \mathcal{F}_s \mathcal{R}f(\theta, \sigma) e^{i\sigma \langle x, \theta \rangle} |\sigma| d\sigma d\theta, \end{aligned}$$

where $i = \sqrt{-1}$ and $\Lambda_s = \mathcal{F}_s^{-1} \cdot |\mathcal{F}_s|$ is a filtering in Fourier domain giving the inversion formula its alias Filtered Back Projection (FBP). The theoretical justification of this inversion formula is given in Appendix A.2. The definition of the 1D Fourier transform, \mathcal{F}_s , is included in Appendix A.2.

Given complete data as in Model 2.2, we have the related inverse and we can reconstruct f perfectly. However, in practice, measurement data is never perfect and a more realistic model would be to include the possibility of noise in the data. Thus, we introduce the following new model:

MODEL 2.5 (CONTINUOUS CT WITH NOISE) Let \mathcal{R} be the Radon transform from Definition 2.1. We model the attenuation of X-rays through an object described by its attenuation coefficients, f , as follows

$$g^\delta(\theta, s) = (\mathcal{R}f)(\theta, s) + \eta \quad \text{for } (\theta, s) \in S^1 \times \mathbb{R}$$

Here $g^\delta(\theta, s) = g(\theta, s) + \eta$ is the measured sinogram with Gaussian distributed additive noise $\eta \sim N(0, \sigma^2)$ with zero mean and variance $\sigma^2 \in \mathbb{R}$.

For this model, we have no direct inversion formula to retrieve f . In fact \mathcal{R}^{-1} in Theorem 2.4 is unbounded, as shown in various fashions, e.g., see IV3 in [1]. This means that small difference between g^δ and g can lead to large differences between $\mathcal{R}^{-1}g^\delta$ and $\mathcal{R}^{-1}g = f$.

2.1.3 Discrete model

In practice we sample finite and imperfect measurements caused by physical acquisition limitations. Hence, our data shows as a finite number of projections in a discrete sinogram. Furthermore, for large data sets, we require computers for the calculations. Since computers are most efficient on discrete information, we are motivated to construct a model using both a discrete object and sinogram.

The discrete object is approximated on a pixel grid stored as a vector, $\mathbf{x} = [x_1, x_2, \dots, x_n] \in \mathbb{R}^n$, which in 2D is displayed as an $N \times N$ image where $n = N^2$. This is illustrated in Figure 2.2.

We approximate a projection in (2.3) by summing over the attenuation coefficients, x_j , multiplied with the euclidean distance travelled by the beam through x_j . Storing all projections by a single index, i , gives us the i th projection as

$$b_i = \sum_{j=1}^n a_{i,j} x_j, \quad (2.4)$$

where $a_{i,j}$ is the length of the i th beam in the j th pixel. The element $a_{i,j}$ is then stored in the i th row and j th column in what we call the *system matrix*, denoted A . From this we establish the following model:

MODEL 2.6 (DISCRETE CT WITH NOISE) We model the attenuation of X-rays through an object described by its attenuation coefficients, \mathbf{x} , as follows:

$$\mathbf{b}^\delta = A\mathbf{x} + \mathbf{e}, \quad \text{for } \mathbf{b}^\delta \in \mathbb{R}^m, \mathbf{x} \in \mathbb{R}^n \text{ and } A \in \mathbb{R}^{m \times n}.$$

Here b_i^δ is a measured projection from (2.4) with added Gaussian distributed white noise, $\mathbf{e} \in \mathbb{R}^m$. The system matrix, A , is thus the discrete approximation to the Radon transform for a given measurement set-up. When using this model we consider the back projection of A as its transposed denoted A^T .

When visually investigating a discrete sinogram, \mathbf{b}^δ , it is considered as a matrix with columns consisting of projections collected from one position of the source and detector. Remember that, in practice, a source emits more than one beam. Generally, two 2D projection sample fashions are considered, i.e., *fan- and parallel-beam*. Figure 2.3 illustrates the difference between the two sampling methods. The dots displayed above the set-up indicate projection values contained in one column of the sinograms, corresponding to the given ray. Throughout, we only consider the fan-beam sampling method, as it relates to the practical application that is considered in the thesis.




Figure 2.2: Illustration of how a single beam is modelled to go through the pixels indicated by dark grey in the discrete object, $\mathbf{x} \in \mathbb{R}^{81}$.




Figure 2.3: Illustrations of fan-beam (left) and parallel-beam (right) projections for a source and detector set-up, showing the projection values in a specific column of the sinogram.

2.1.4 Implementation of model, simulation & testing

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

In this section we describe the testing platform used in the thesis to run synthetic and real tests on the discrete CT models.

We require a method of generating the system matrix, A , that defines the fan-beam sampling method. To this end, we need to define physical parameters related to an actual measurement set-up, to run on real measurement data. We use the `fanbeamtomo.m` function from the AIR Tools package [2] as a starting point for generating the system matrix. We modify the function to use a linear detector and make it possible to define the set-up from 6 physical parameters and a given grid size N . The parameters are shown in Table 2.1. Note here the grid represents an object with $n = N^2$ pixels.

Table 2.1: Default parameters used to generate the system matrix, A , using a modified `fanbeamtomo` function with a linear detector. These parameters are used for synthetic tests in Chapter 2.1 and 2.2.

Grid size (N)	256
Number of source locations	180
Number of detector pixels	256
Domain size	46 [cm] \times 46 [cm]
Source to centre distance	59 [cm]
Source to detector distance	100 [cm]
Detector length	90 [cm]

In addition to the modified `fanbeamtomo` function, we use the ASTRA Toolbox [3] to generate a matrix free version of the system matrix using the same parameters as in Table 2.1. The details of how these functions are made has been left out of the thesis.

To simulate real measurement data, we create synthetic phantoms on a finer grid, $\mathbf{x} \in \mathbb{R}^{9n}$, to avoid *inverse crime*, multiplying the matrix free version of the system matrix, $A \in \mathbb{R}^{m \times 9n}$, with the phantom creating a sinogram, $\mathbf{b}^\delta \in \mathbb{R}^m$ as follows:

$$\mathbf{b}^\delta = A\mathbf{x} + \mathbf{e}^\delta.$$

Here \mathbf{e}^δ is added white noise with a *relative noise level* δ given by

$$\mathbf{e}^\delta = \delta \|\mathbf{b}\|_2 \frac{\mathbf{e}}{\|\mathbf{e}\|_2}, \quad (2.5)$$

where $\mathbf{e} \in \mathbb{R}^m$ such that $\mathbf{e} \sim N(\mathbf{0}, 1)$.

Throughout, we compare images by two different measures when we know the ground truth, i.e., the *relative error (RE)* and *mutual information (MI)* as defined below

DEFINITION 2.7 For two images $X, Y \in \mathbb{R}^{N \times N}$ we define the *relative error (RE)* and *mutual information (MI)* as

$$RE(X, Y) := \frac{\|X - Y\|_F}{\|Y\|_F}, \quad MI(X, Y) := H_2(X, Y) - H_1(X) - H_1(Y).$$

Here

$$\begin{aligned} H_2(X, Y) &:= - \sum_{j,k} PDF_{X,Y}(j, k) \log(PDF_{X,Y}(j, k)), \\ H_1(X) &:= - \sum_j \left[\left(\sum_k PDF_{X,X}(j, k) \right) \log \left(\sum_k PDF_{X,X}(j, k) \right) \right], \\ PDF_{X,Y}(j, k) &:= \frac{HIST_{X,Y}(j, k)}{\sum_{l,p} HIST_{X,Y}(l, p)}, \end{aligned}$$

and $HIST_{X,Y}$ is the joint intensity histogram of the images X, Y .

We note that RE is an error measure comparing intensity values at a fixed position and MI is a similarity measure comparing the distribution of intensity values. We include MI to give a different measure of similarity that do not take the scale or position of intensities into account, but rather the *information* in the image.

2.1.5 An immediate reconstruction approach

One approach to reconstruct the discrete image from corrupted data, \mathbf{b}^δ with parameters as in Table 2.1, would be to use the discrete version of the inversion formula in Theorem 2.4 approximating the back projection by A^T . Denoting Λ as a discrete filtering operator, $\mathcal{F}^{-1}|\cdot|_{\mathcal{F}}$ on \mathbf{b}^δ , we have the reconstruction formula:

$$\mathbf{x} = \frac{1}{2}(2\pi)^{-3/2} A^T \Lambda \mathbf{b}^\delta. \quad (2.6)$$

This reconstruction is motivated by the similarities of the discrete and continuous model. Just as with the continuous model, the filter, Λ , amplifies high

frequencies in the sinogram. This reveals a concerning issue since white noise contains high frequent elements. This amplification is illustrated in Figure 2.4 where adding 10% relative noise, defined in (2.5), to the sinogram before reconstructing shows a severe difference compared to adding the same relative level of noise to the image directly. This shows that Model 2.6 preserves some of the properties from the continuous model. The challenges of reconstructing from a noisy sinogram as in Model 2.6 has been under the scope of researchers for many years and is treated in the following sections.

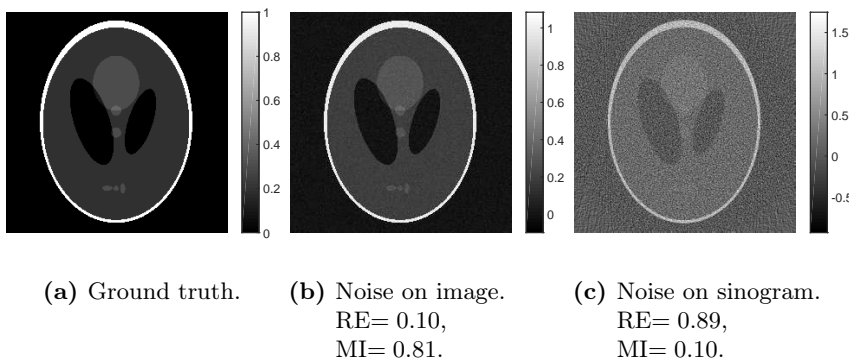


Figure 2.4: Illustration of noise amplification using the discrete filtered back projection in (2.6). The centre image has added 10% relative noise directly on it and the right is a reconstruction of a sinogram with added 10% relative noise. It is clear that the noise is amplified in the reconstruction.

2.2 Inverse problems & regularisation

Main author: Nicolai André Brogaard Riis.

Co-author: Jacob Frøsig.

This section is meant to give a brief overview to the field of inverse problems, to which CT belongs. By studying CT in a more general setting, we are able to gain significant insights into how to tackle the challenges in CT reconstruction as described in Section 2.1. We describe the notion of *ill-posedness* in inverse problems and how regularisation is used to obtain meaningful solutions to these types of problems. Finally, we give give examples of typical regularisation methods and show how they fare for reconstruction in CT problems.

2.2.1 Inverse problems

The term *inverse problem* generally tends to describe the framework used in mathematics to obtain information on features or parameters for an object or system that is not directly observable. The information is obtained by processing a measurable property that is affected by the object. The goal of solving an inverse problem, is to approximate the features or parameters of the object that best match with the measured property.

A general mathematical statement of an inverse problem is given data, g , find the model parameters, f , such that,

$$\mathcal{A}(f) = g. \quad (2.7)$$

Here \mathcal{A} is a model operator that describes the relation between the model parameters, f , and some given measurement data, g . The model operator is known for a specific model, but is for real data an approximation to the actual phenomena generating the data. Note that finding the attenuation coefficients in the CT Models 2.2, 2.5 and 2.6 are inverse problems.

To clarify: We call the process of constructing g given f the *forward problem* and finding f given g the *inverse problem*. Generally, if an approximation of f is known, it is a stable process to calculate g . On the other hand if an approximation of g is known, and we want to find f , it can be an unstable process if the model operator induces *ill-posedness* to the inverse problem.

2.2.2 Ill-posedness of inverse problems

Ill-posedness is one of the key challenges for most inverse problems, and lie at the heart of why inverse problems such as CT are still studied to the extent that they are today. When discussing the class of ill-posed problems, it is beneficial to first define the compliment; *well-posed* problems. For the purposes of this thesis, we restrict our attention to the case when the model operator \mathcal{A} in (2.7) is linear, as is the case with the CT models. In the early 20th century Hadamard [4] gave three conditions for a linear problem to be well-posed:

1. **Existence:** A solution to the problem exists.
2. **Uniqueness:** The solution is unique.
3. **Stability:** The solution depends continuously on the data.

If any of these conditions are not satisfied, we say that the problem is *ill-posed*.

The stability condition can be interpreted as follows: Small changes in data must only give a small change in solution.

For the discrete CT Model 2.6, added noise to the measurements can move the data out of the range of A , meaning no solution to the problem exists if one simply tries to find \mathbf{x} such that $A\mathbf{x} = \mathbf{b}^\delta$. Likewise if A is not injective, then several solutions can exist and so condition 2 is not satisfied. Finally, as we have seen in Chapter 2.1, small changes in the data need not translate to small changes in the solution, but rather to arbitrarily large changes. When facing ill-posed inverse problems such as CT, it is essential to overcome these challenges to get any meaningful results from the solutions. A good reference on (more general) discrete inverse problems is [5].

2.2.3 Regularisation

To overcome the challenges of ill-posed inverse problems, it is most common to introduce regularity on the solution by solving a modified version of the problem that is well-posed by construction. This is known as regularisation. For example; one can satisfy the existence condition by rephrasing (2.7) into one of fitting the solution to the data:

$$f^* = \operatorname{argmin}_f \{ \|\mathcal{A}(f) - g\|^2 \}.$$

To deal with stability, we can add an additional term that enforces the norm of the solution to be small in some linear operator \mathcal{L} :

$$f^* = \operatorname{argmin}_f \{ \|\mathcal{A}(f) - g\|^2 + \|\mathcal{L}f\|^2 \}.$$

Depending on the norm, this additional term can also enforce uniqueness, say for the 2-norm, when the nullspaces of \mathcal{L} and \mathcal{A} have a trivial intersection.

In general, regularisation uses a *data fidelity* term $T(f)$, and a *regularisation term* $R(f)$ for the regularised problem:

$$f_\alpha^* = \operatorname{argmin}_f \{ T(f) + \alpha R(f) \}. \quad (2.8)$$

Here α is a regularisation parameter that needs to be chosen appropriately balancing both terms. The data fidelity term is a measure of the fit between the given data and how the data would look if the solution and forward model was used to generate it. The regularisation term is how we incorporate additional information about our object. Say that we know, a priori, the solution should not have large individual values. Then a 2-norm regularisation term $\|f\|_2^2$ can be a

good way to incorporate this prior. Depending on the regularisation parameter, α , we can enforce more or less of this "smoothness" prior.

The choice of regularisation parameter deserves a story of its own and is indicative of the parameter tuning necessary for most, if not all, regularisation methods. In general, there are many guides to choose such a parameter automatically. However, all of them depend on the problem and it is typical to try several guides as well as judging manually (e.g. by looking at the reconstructed images in the case of CT) which parameter to choose. We do not go into detail on the guides in this thesis, but refer to [5] for more details on the choice of regularisation parameter.

2.2.4 Applying regularisation to computed tomography

We turn our attention to the discrete CT Model 2.6. Recall that added noise, even if relatively small, made the discretised inversion formula of Theorem 2.4 produce severely deteriorated solutions caused by the ill-posedness of the problem. Hence, we are motivated to apply regularisation to the problem.

We test on the Shepp–Logan phantom shown in Figure 2.5. The generated fan-beam data can be seen as a sinogram in Figure 2.6. The data is generated from fan-beam projections using the physical parameters in Table 2.1. We now describe three methods that add regularisation to the problem and show how they perform on this set-up.




Figure 2.5: Shepp–Logan phantom used for simulating measurement data. The interior of the blue circle indicates the area which is illuminated from all source positions by X-rays defined for the set-up in Table 2.1. In this case the entire object is illuminated from all source positions.

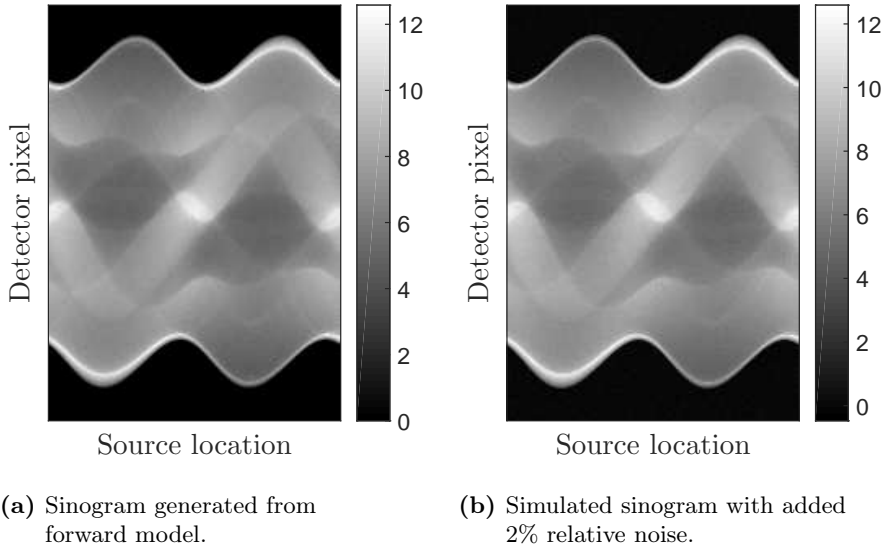


Figure 2.6: Sinograms generated from the object in Figure 2.5 using the measurement set-up in Table 2.1.

2.2.4.1 Tikhonov regularisation

Tikhonov regularisation takes the form of the optimisation problem in (2.8) with a squared 2-norm data fidelity and regularisation term as follows:

$$\mathbf{x}^*|_{\alpha} = \underset{\mathbf{x}}{\operatorname{argmin}} \{ \|\mathbf{A}\mathbf{x} - \mathbf{b}^{\delta}\|_2^2 + \alpha \|\mathbf{x}\|_2^2 \}.$$

The method is quite popular as a regularisation technique for the discrete CT Model 2.6, because of the combination of an intuitive smoothness prior, $\|\mathbf{x}\|_2^2$, and the fact that it can be rewritten, and hence implemented, as a least squares problem (see e.g. [5]).

The need for parameter tuning for Tikhonov is apparent in the regularisation parameter, α . If the parameter is chosen too small, i.e. $\alpha \rightarrow 0$, we get a solution corrupted by noise. On the other hand, if α is chosen too large, i.e. $\alpha \rightarrow \infty$, we get a solution that is smoothed to zero. The idea is thus to choose a regularisation parameter in-between that gives the best reconstructed image. For this simulated problem, we know the real solution and so we can compare with the relative error and mutual information measures as described in Definition 2.7. This is illustrated in Figure 2.7 for a number of regularisation parameters. From this we see that choosing $\alpha = 2$ gives the lowest error. In Figure 2.8 we

see a comparison of the image quality in the extreme cases of α too low or too high compared to $\alpha = 2$; just right.




Figure 2.7: Error measures (see Definition 2.7 p. 15) depending on regularisation parameter for Tikhonov regularisation in Section 2.2.4.1 on the data in Figure 2.6. We see that there is an optimal choice of regularisation parameter for both error measures around 2 for this particular problem.




Figure 2.8: Reconstructed images using Tikhonov regularisation on the simulated data in Figure 2.6 with regularisation parameters as shown.

2.2.4.2 Semi converging iterative methods

The semi converging methods arise from iterative regularisation methods that let the number of iterations play the role of regularisation parameter. Typically these methods converge towards an undesired solution. The interesting thing, however, is on the way to this solution some of the methods exhibit *semi-convergence*. That is, they approach some solution, x^* , that is the best approximation to the ground truth possible for the method and then diverge again towards the undesired solution. For the purpose of this thesis, we state the simplest of these methods, namely the Landweber iterative method which takes the form:

$$\mathbf{x}^{[k+1]} = \mathbf{x}^{[k]} + sA^T(\mathbf{b}^\delta - A\mathbf{x}^{[k]}),$$

where s is a step size parameter that must satisfy $0 < s < 2\|A^T A\|_2^{-1}$.

There is much to say on the iterative methods that exhibit semi-convergence (see e.g. [5]), but for the purpose of this thesis, we simply show how the iteration number acts as a regularisation parameter. In Figure 2.9 we see the change in error measures as the iteration number increases. Indeed the method shows semi-convergence and we see that choosing to stop iterating around the 100th iteration gives the best solution. In Figure 2.10 the reconstructed images are shown for three choices of iteration numbers.

2.2.4.3 Filtered back projection

Finally, we consider the discrete inversion formula (2.6). By modifying the filter to not amplify high frequent noise, we add regularisation to the problem. Denote by Λ_Γ a modified filtering step such that $\Lambda_\Gamma = \mathcal{F}^{-1}\Gamma(\cdot)\mathcal{F}$, for some frequency filter $\Gamma : \mathbb{R} \rightarrow \mathbb{R}$. The regularised solution is then found by

$$\mathbf{x} = \frac{1}{2}(2\pi)^{-3/2}A^T\Lambda_\Gamma\mathbf{b}^\delta. \quad (2.9)$$

When the filter is $\Gamma(\cdot) = |\cdot|$, we have the discrete version of the inversion formula (2.6). In Figure 2.11 we see an illustration of some standard filters that can be used to replace the $|\cdot|$ filter in Fourier domain thus adding regularisation. The idea is to let the filter go to zero as the frequency increases, thus removing the noise that is represented by high frequencies. In Figure 2.12 we see reconstructed images using the different filtering methods. We see that the filtering acts as regularisation and dampens the noise corruption in the image. Note that, since the discrete Fourier transform in frequency domain is band limited, the $|\cdot|$ filter also regularises the solution to some extent.




Figure 2.9: Error measure depending on iteration number for Landweber method in Section 2.2.4.2 on the data in Figure 2.6. The optimal stopping iteration for the Landweber method is seen to be around iteration 100 for this problem.

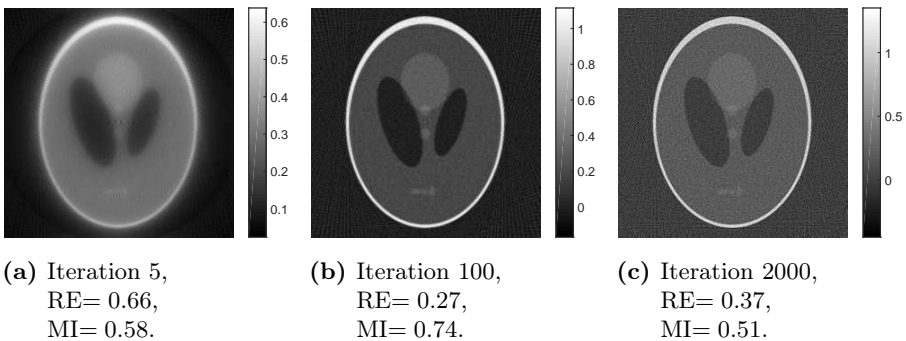


Figure 2.10: Reconstructed images using the Landweber method on the simulated data in Figure 2.6 with iterations as shown.

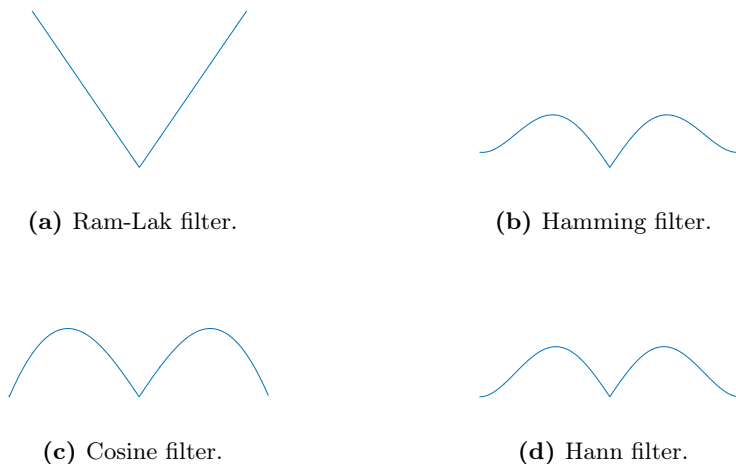


Figure 2.11: Illustration of different filters for the filtered back projection formula in (2.9). The Ram-Lak filter is from the inversion formula, and the other filters remove high frequent elements in Fourier domain before back projection; a form of regularisation.

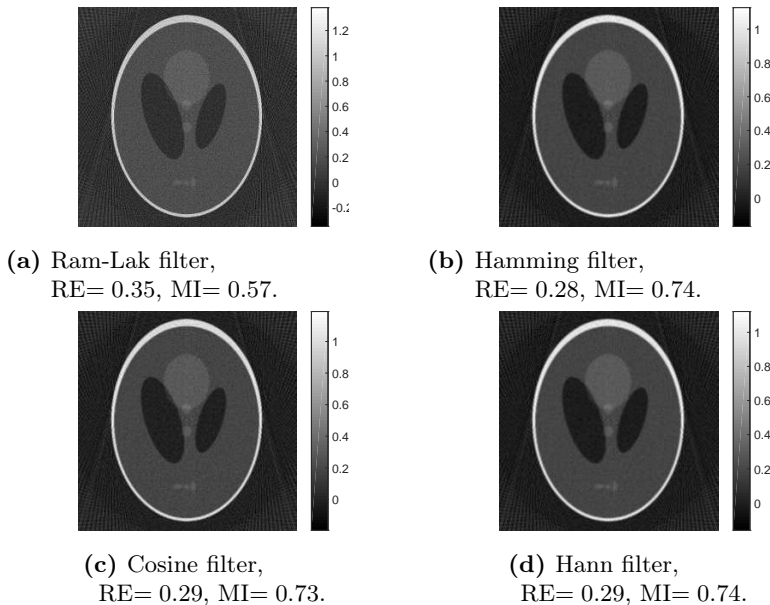


Figure 2.12: Reconstructed images using the discrete filtered back projection formula in (2.9) with the filters from Figure 2.11.

2.3 Microlocal analysis

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

Microlocal analysis (MLA) can be used to determine which singular features we can expect to recover in a range of continuous tomography problems. Here singular features describe edges arising from transitions in material. Many objects are partly characterised by their material boundaries and hence it is useful to describe how these so-called singularities propagate in general tomographic transforms. As an example: If we can determine which singularities in the sinogram are (or are not) generated by the object and which are added artefacts, we can develop methods that take this insight into account.

To fully understand MLA in tomography, one requires knowledge of wavefront sets, pseudo-differential operators and Fourier integral operators. The next three sections (2.3.1-2.3.3) establish this large body of theory briefly, leading up to the actual propagation theorems in Section 2.3.4. The reader can skip the first three sections and go straight to the latter if he/she is only interested in the resulting theorems on how singularities propagate. We have, however, decided to include the theory for completeness. We base this section on the excellent work in [6].

2.3.1 Singular support & wavefront Set

To determine how material boundaries (edges) propagate in tomographic transforms, we must first agree on a useful mathematical definition on what exactly is meant by this. We characterise material boundaries by *singularities* as elements of the so called *wavefront set*.

Singularities are composed by two components: The first being spatial location and the second being direction. The location information is included in the *singular support set* and the directional information in the *localised frequency set*. The product space of these is exactly the *wavefront set*. In Figure 2.13 we see an illustration of the wavefront set for an object with smooth boundaries. Here the boundary is the singular support and the localised frequency set are the normals at each point on the boundary. Before we get ahead of ourselves, we must define the above more formally.

Denote the set of distributions by $\mathcal{D} = (C_c^\infty)'$. The *singular support set* of a distribution, f , is defined as follows.




Figure 2.13: The wavefront set of an object, f , is the singular support set (red boundary) paired with directional information as illustrated for three points on boundary.

DEFINITION 2.8 Let $f \in \mathcal{D}(X)$, $X \subset \mathbb{R}^n$. The singular support of f , denoted $\text{ssupp}(f)$, is the complement of the largest open set in X on which f is C^∞ smooth.

The singular support are all points on which f has some non-smoothness. We are interested in describing how these points propagate for different transforms. To do this we need to include directional information for each point. Directional information can be described in frequency domain for each singular point using the notion of *rapid decay*:

DEFINITION 2.9 A function $f : \mathbb{R}^n \rightarrow \mathbb{C}$ is said to be *rapidly decaying at infinity* if for every $N \geq 0$, there exists a C_N such that

$$|f(\xi)| \leq C_N(1 + \|\xi\|)^{-N}, \quad \text{for all } \xi \in \mathbb{R}^n. \quad (2.10)$$

Rapid decay is related to the smoothness of a distribution by the following theorem:

THEOREM 2.10 (P. 252 IN [7]) A compactly supported distribution $f \in \mathcal{D}_c(X)$ is in $C_c^\infty(X)$ if and only if its Fourier transform is rapidly decaying at infinity.

REMARK. We use the integral *notation* for the Fourier transform of a distribution f . When doing integral manipulations, we assume that $f \in L^1 \subset \mathcal{D}$.

The above theorem gives directional information of the singularities in the following sense: Let \mathcal{V}_η be an open conic neighbourhood of some non-zero point, η , we define the frequency set of f as follows:

$$\Sigma(f) := \{\eta \in \mathbb{R}^n \setminus \{\mathbf{0}\} : (2.10) \text{ is not satisfied for } \xi \text{ in any } \mathcal{V}_\eta\}.$$

Note here that $\Sigma(f) = \emptyset$ if and only if $f \in C_c^\infty$.

The frequency set of f , $\Sigma(f)$, is exactly the frequency directions for which f is not C^∞ smooth. However, this information is global for all points in the singular support set. To obtain local directional information, we must assign directions for each location.

To this end we localise our domain using the following lemma:

LEMMA 2.11 (LEMMA 8.1.1 ON P. 253 IN [7]) *If $\phi \in C_c^\infty(\mathbb{R}^n)$ and $f \in \mathcal{D}_c(\mathbb{R}^n)$ then*

$$\Sigma(\phi f) \subset \Sigma(f).$$

For an open set X the *local* frequency set, $\Sigma_x(f)$, for some $x \in X$ and $f \in \mathcal{D}(X)$ is then defined for some set of test functions $\{\phi_k\}$ as follows:

$$\Sigma_x(f) := \bigcap_{\substack{\phi_k \in C_c^\infty \\ \phi_k(x) \neq 0}} \Sigma(\phi_k f).$$

If $\phi_k \in C_c^\infty(X)$, $\phi_k(x) \neq 0$ and $\text{supp } \phi_k \rightarrow \{x\}$ for $k \rightarrow \infty$, then by Lemma 2.11, we can *microlocalise* from the frequency set such that,

$$\Sigma(\phi_k f) \rightarrow \Sigma_x(f).$$

This implies $\Sigma_x(f) \neq \emptyset$ if and only if $x \in \text{ssupp}(f)$.

We combine the singular support and the local frequency set into the *wavefront set* by the following definition:

DEFINITION 2.12 The wavefront set of $f \in \mathcal{D}(X)$, denoted $WF(f)$, is defined by

$$WF(f) := \{(x, \xi) \in X \times (\mathbb{R}^2 \setminus \{\mathbf{0}\}) : \xi \in \Sigma_x(f)\}$$

The projection in X is $\text{ssupp}(f)$.

We now have the wavefront set describing the location and (local) direction of singularities. The aim going further is to show how the wavefront set propagates in tomographic transforms. To achieve this goal, we describe how it propagates for pseudo-differential and Fourier integral operators and later relate these to tomographic transforms.

2.3.2 Pseudo-differential operators

In this section an enlargement of the differential operator class, i.e., the pseudo-differential operators is introduced. The class shows general properties regarding propagation of wavefront sets which is stated as a theorem later on.

To motivate pseudo-differential operators consider the following linear partial differential operator on the form

$$(Pf)(x) := \sum_{|v| \leq m} a_v(x) D_x^v(f(x)), \quad (2.11)$$

with v as a multi-index: $v = (v_1, \dots, v_n)$, $|v| := v_1 + \dots + v_n$, $x^v := x_1^{v_1} x_2^{v_2} \dots x_n^{v_n}$ for x in n -dimensional space and

$$D_x^v := (-i)^{|v|} \frac{\partial^{v_1}}{\partial x_1^{v_1}} \dots \frac{\partial^{v_n}}{\partial x_n^{v_n}}.$$

Recall that differentiation is equivalent to multiplication in Fourier domain by the following lemma.

LEMMA 2.13 *For a compactly supported function $f \in C_c^\infty(\mathbb{R}^n)$*

$$\widehat{D_x^v f}(\xi) = \xi^v \hat{f}(\xi).$$

PROOF. Applying the Fourier Transform to $D_x^v u$ and integrating by parts $|v|$ times, we get (recall definition of D has $(-i)$ and that u has compact support):

$$\begin{aligned} \widehat{D_x^v u}(\xi) &= (2\pi)^{n/2} \int e^{-ix \cdot \xi} D_x^v u(x) \, dx \\ &= (2\pi)^{n/2} \int (-i)^{|v|} \xi^v e^{-ix \cdot \xi} (-i)^{|v|} u(x) \, dx \\ &= \xi^v \hat{u}(\xi). \end{aligned}$$

□

Consider the operator in (2.11) applied to a compactly supported function, $f \in L^1(\Omega)$. Applying the Fourier and Inverse Fourier Transform, we get

$$\begin{aligned}
 (Pf)(x) &= \sum_{|v| \leq m} a(x) \mathcal{F}^{-1} \mathcal{F} D_x^v f(x) \\
 &= \sum_{|v| \leq m} a(x) \mathcal{F}^{-1} \xi^v \hat{f}(\xi) \\
 &= \sum_{|v| \leq m} a(x) (2\pi)^{-n/2} \int e^{ix \cdot \xi} \xi^v \hat{f}(\xi) \, d\xi \\
 &= \sum_{|v| \leq m} a(x) (2\pi)^{-n} \int \int e^{i(x-y) \cdot \xi} \xi^v f(y) \, dy \, d\xi.
 \end{aligned}$$

Rearranging, we get

$$(Pf)(x) = (2\pi)^{-n} \int e^{i(x-y) \cdot \xi} p(x, \xi) f(y) \, dy \, d\xi, \quad (2.12)$$

with $p(x, \xi) = \sum_{|v| \leq m} a(x) \xi^v$.

This Fourier representation of the differential operator helps us define a larger class of operators that act like the differential operator, namely pseudo-differential operators. The requirement is that $p(x, \xi)$ in (2.12) is an amplitude:

DEFINITION 2.14 Let $X \subset \mathbb{R}^n$ be an open subset. An *amplitude* of order m is a function that satisfies the following properties:

1. $p(x, y, \xi) \in C^\infty(X \times X \times \mathbb{R}^n \setminus \{\mathbf{0}\})$,
2. For every compact set K and for multi-index α, β, γ ,

(a) there is a constant $C = C(K, \alpha, \beta, \gamma)$ such that

$$|D_\xi^\alpha D_x^\beta D_y^\gamma p(x, y, \xi)| \leq C(1 + \|\xi\|)^{m - |\alpha|} \text{ for } \|\xi\| > 1, \text{ and}$$

(b) $P(x, y, \xi)$ is locally integrable for x and y in K and $\|\xi\| \leq 1$.

The definition of a pseudo-differential operator is then as follows.

DEFINITION 2.15 Let $X \subset \mathbb{R}^n$ be an open subset. A pseudo-differential operator (Ψ DO), \mathcal{P} , is an operator of the form:

$$(\mathcal{P}f)(x) = \frac{1}{(2\pi)^n} \int e^{i(x-y) \cdot \xi} p(x, y, \xi) f(x) \, dx \, d\xi,$$

where its *symbol*, $p(x, y, \xi)$, is an amplitude on $(X \times X \times \mathbb{R}^n \setminus \{\mathbf{0}\})$. We say the operator \mathcal{P} is of order m if its symbol has order m . \mathcal{P} is elliptic of order m if for each compact set $K \subset X$, there is a constant $C_k > 0$ such that for x and y in K and $\|\xi\| \geq C_k$

$$|p(x, y, \xi)| \geq C_k^{-1}(1 + \|\xi\|)^m.$$

Now that we have defined Ψ DO's, we investigate how they propagate wavefront sets.

THEOREM 2.16 (THEOREM 14 P. 878 IN [6]) *If \mathcal{P} is a Ψ DO, then \mathcal{P} satisfies the following pseudo-local property:*

$$\text{ssupp}(\mathcal{P}f) \subset \text{ssupp}(f) \text{ and } \text{WF}(\mathcal{P}f) \subset \text{WF}(f).$$

In addition, if \mathcal{P} is elliptic, then

$$\text{ssupp}(\mathcal{P}f) = \text{ssupp}(f) \text{ and } \text{WF}(\mathcal{P}f) = \text{WF}(f).$$

This tells us that if \mathcal{P} is a Ψ DO, any singularity in $\mathcal{P}f$ arises from, and will have the same structure as, singularities in f if they exist. In other words, \mathcal{P} introduces no new singularities. Furthermore, if \mathcal{P} is elliptic, the singularities of \mathcal{P} and $\mathcal{P}f$ are the same. This indicates that if we know the singularities of $\mathcal{P}f$, we know the singularities of f .

2.3.3 Fourier integral operators

In this section we introduce an enlargement of the Ψ DO class described in the previous section, i.e, the Fourier integral operators. Even though this class do not retain the same strong singularity propagation properties, there is still a story to be told for specific FOI's. Hence, some general definitions are established.

To consider Fourier integral operators (FIO) we first define a *phase function*:

DEFINITION 2.17 Let $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$ be open subsets. A real-valued function $\phi \in C^\infty(X \times Y \times \mathbb{R}^N \setminus \{\mathbf{0}\})$ is called a phase function if

1. ϕ is positive homogeneous of degree 1 in ξ . That is $\phi(x, y, r\xi) = r\phi(x, y, \xi)$ for all $r > 0$.
2. $(\partial_x \phi, \partial_\xi \phi)$ and $(\partial_y \phi, \partial_\xi \phi)$ do not vanish for all $(x, y, \xi) \in X \times Y \times \mathbb{R}^N \setminus \{\mathbf{0}\}$.

Now, we are able to define the FIO's as follows:

DEFINITION 2.18 Let $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$ be open subsets. A Fourier integral operator (FIO), \mathcal{P} , is defined as

$$\mathcal{P}f(x) = \int e^{i\phi(x,y,\xi)} p(x,y,\xi) f(y) dy d\xi,$$

where ϕ is a nondegenerate phase function and $p(x,y,\xi) \in C^\infty(X \times Y \times \mathbb{R}^N)$ is an amplitude satisfying:

For every compact set $K \subset X \times Y$ and for every multi-index α, β, γ , there exists a constant $C = C(K, \alpha, \beta, \gamma)$ such that

$$|D_\xi^\alpha D_x^\beta D_y^\gamma p(x,y,\xi)| \leq C(1 + \|\xi\|)^{m-|\alpha|} \text{ for all } x,y \in K \text{ and for all } \xi \in \mathbb{R}^N.$$

Note how the enlargement from Ψ DO's to FIO's becomes apparent by letting $\phi(x,y,\xi) = (x-y) \cdot \xi$.

In the following, we introduce notation and definitions to investigate the FIO's. Firstly, for any FIO we introduce a related auxiliary manifold and canonical relation:

DEFINITION 2.19 (AUXILIARY MANIFOLD & CANONICAL RELATION)

Let $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$ be open subsets. For a given FIO we define the auxiliary manifold by

$$\Sigma_\phi := \{(x,y,\xi) \in X \times Y \times \mathbb{R}^N \setminus \{\mathbf{0}\} : \partial_\xi \phi(x,y,\xi) = 0\},$$

and the canonical relation by

$$C := \left\{ \begin{pmatrix} x & \partial_x \phi(x,y,\xi) \\ y & -\partial_y \phi(x,y,\xi) \end{pmatrix} : (x,y,\xi) \in \Sigma_\phi \right\}.$$

Later, we use the following operator for the canonical relation on two sets to describe singularities propagation for FIO's.

DEFINITION 2.20 Let $X \subset \mathbb{R}^m$ and $Y \subset \mathbb{R}^n$ be open subsets and let $C \subset T^*X \times T^*Y$ and $A \subset T^*Y$. Define

$$C \circ A = \left\{ (x, \xi dx) : \exists (y, \eta dy) \in A \text{ with } \begin{pmatrix} x & \xi dx \\ y & \eta dy \end{pmatrix} \in C \right\}.$$

Here we use the notation T^*X , to define the set of point and directions of X .

2.3.4 Application to tomography

We now have the tools to investigate how singularities (edges, boundaries) propagate in tomographic transforms. To this end we will first study the combined projection and back projection operator $\mathcal{R}^*\mathcal{R}$. That is, for an object, what does the process of *measuring* the object using the Radon Transform and then reconstructing using back projection (with no filtering) do to singularities of the object. It turns out this operator is an elliptic pseudo differential operator:

THEOREM 2.21 (P. 875 IN [6]) $\mathcal{R}^*\mathcal{R}$ is an elliptic pseudo-differential operator with symbol $4\pi/\|\xi\|$. Hence $\text{WF}(\mathcal{R}^*\mathcal{R}f) = \text{WF}(f)$ by Theorem 2.16.

That is the wavefront set is maintained for $\mathcal{R}\mathcal{R}^*$.

The Radon transform relates our object to the obtained data. Hence, propagation of singularities over this transform is of interest to us. Let $\theta(\varphi) = (\cos(\varphi), \sin(\varphi))$ and $\theta(\varphi)^\perp = \omega(\varphi + \pi/2)$. We then have the following theorem.

THEOREM 2.22 (THEOREM 17 P. 884 IN [6]) The Radon transform, \mathcal{R} , is an elliptic FIO on $\mathcal{D}_c(\mathbb{R}^2)$ associated with the canonical relation for:

$$C_{\mathcal{R}} = \left\{ \begin{array}{l} \left(\begin{array}{cc} (\varphi, s) & \alpha(-x \cdot \theta^\perp(\varphi)d\varphi + ds) \\ x & \alpha\theta(\varphi)dx \end{array} \right) \\ : (\varphi, s) \in [0, 2\pi] \times \mathbb{R}, x \in \mathbb{R}^2, \alpha \neq 0, x \cdot \theta(\varphi) = s \end{array} \right\},$$

and $\text{WF}(f) = C_{\mathcal{R}} \circ \text{WF}(\mathcal{R}f)$.

We collect our knowledge of propagations of singularities in the following corollaries:

COROLLARY 2.23 Let $f \in \mathcal{D}_c(\mathbb{R}^2)$ and let \mathcal{R} be the 2D Radon transform, then we have the following relations

$$\text{WF}(\Lambda\mathcal{R}^*\mathcal{R}f) = \text{WF}(f) \quad \text{for any elliptic } \Psi\text{DO } \Lambda.$$

PROOF. By Theorem 2.21 we have that $\text{WF}(\mathcal{R}^*\mathcal{R}f) = \text{WF}(f)$ and since Λ is an elliptic ΨDO , we then have that $\text{WF}(f) = \text{WF}(\mathcal{R}^*\mathcal{R}f) = \text{WF}(\Lambda\mathcal{R}^*\mathcal{R}f)$. \square

The corollary shows that any reconstruction method defined as $\Lambda\mathcal{R}^*\mathcal{R}f$ retains the wavefront set if Λ is an elliptic ΨDO . In [6] this is shown to be the case

for the FBP lambda tomography reconstruction formulas. This is not studied further in this thesis. However, the propagation of the Radon transform is studied since it has implications for limited data tomography. From Theorem 2.22 we get the following corollary.

COROLLARY 2.24 (COROLLARY 1 P. 885 IN [6]) *Let $f \in \mathcal{D}_c(\mathbb{R}^2)$ then*

- *If $(x_0, \xi_0 dx) \in \text{WF}(f)$, then*

$$\left(\begin{array}{c} \varphi_0, x_0 \cdot \theta(\varphi_0) \\ t(-x_0 \cdot \theta^\perp(\varphi_0) d\varphi + ds) \end{array} \right) \in \text{WF}(\mathcal{R}f),$$

where φ_0 is chosen such that $\xi_0 = t\theta(\varphi_0)$ for some $t \neq 0$.

That is, the singularities in the object shows as a singularity in the sinogram if a line through the point with direction orthogonal to the singularity is measured.

- *Let $(\varphi_0, s_0) \in [0, 2\pi] \times \mathbb{R}$ and $q \in \mathbb{R}$. If*

$$\left(\begin{array}{c} \varphi_0, s_0 \\ t(-q d\varphi + ds) \end{array} \right) \in \text{WF}(\mathcal{R}f),$$

for some $t \neq 0$, then $(x_0, \xi_0 dx) \in \text{WF}(f)$, where $x_0 = s_0\theta(\varphi_0) + q\theta^\perp(\varphi_0)$ and $\xi_0 = t\theta(\varphi_0)$.

That is, given a singularity in the sinogram we can find the corresponding singularities that generated it.

The corollary has implications for limited data tomography. A singularity $(x_0, \xi_0 dx) \in \text{WF}(f)$ is visible from limited 2D Radon data if and only if the line through x_0 perpendicular to ξ_0 is measured.

Region-of-interest tomography

Main author: Jacob Frøsig.

Co-author: Nicolai André Brogaard Riis.

In this chapter we introduce the region-of-interest (ROI) tomography problem and relate it to the problem described by FORCE. To understand the underlying challenges of the problem, we describe a suitable model for ROI tomography and illustrate how it is caused by the measurement geometry. Using the new model, we investigate how previous mentioned reconstruction methods perform. In the latter part, we interpret the results and apply MLA to characterise challenges of ROI tomography.

3.1 The region-of-interest model

The ROI problem is a limited data problem. The limitation lies in the fact that not all parts of the object are illuminated by X-rays from every source position. In Figure 3.1 the ROI, denoted by Ω , is shown in red for a measurement set-up. As the source moves around the object (indicated by dashed lines and arrows) Ω is the only part of the object that is illuminated from all positions. Since this area can be seen as the interior of the object, the problem is sometimes also called *the interior problem*. The problem of limited data occurs if one is interested in more than the interior of the object.

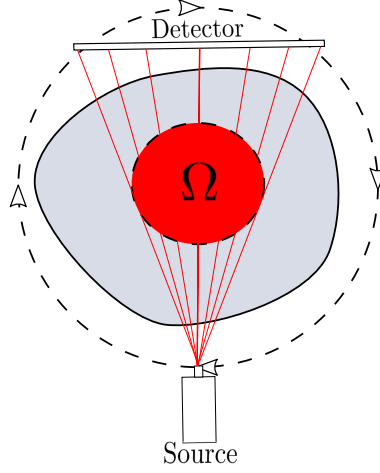


Figure 3.1: Illustration of a ROI measurement geometry. This illustrates how only the ROI, indicated in red, is illuminated from all source positions.

Recall that the measurement set-up proposed by FORCE has a source with a narrow beam that do not illuminate the whole pipe from any source position. Hence the problem can be interpreted as a ROI tomography problem.

The continuous forward operator in ROI tomography differs from Model 2.5 since the data is limited to lines intersecting the ROI. Hence, we define the ROI tomography model as follows:

MODEL 3.1 (CONTINUOUS ROI-TOMOGRAPHY WITH NOISE)

Let \mathcal{R} be the Radon transform from Definition 2.1. We model the attenuation of X-rays through the ROI, Ω , of an object described by its attenuation coefficients, f , as follows

$$g_{\Omega}^{\delta}(\theta, s) = (\mathcal{R}_{\Omega}f)(\theta, s) + \eta \quad \text{for } (\theta, s) \in S^1 \times \mathbb{R},$$

where

$$(\mathcal{R}_{\Omega}f)(\theta, s) = \begin{cases} (\mathcal{R}f)(\theta, s), & \text{for } \ell(\theta, s) \cap \Omega \neq \emptyset, \\ \text{Undefined}, & \text{otherwise.} \end{cases}$$

Here g_{Ω}^{δ} is the measured ROI sinogram with Gaussian distributed additive noise $\eta \sim N(0, \sigma^2)$ with zero mean and variance $\sigma^2 \in \mathbb{R}$.

The undefined information should be interpreted as follows: The forward model should not include any information for positions that are not measured in an ROI set-up. This is effortlessly included in the discrete forward model (2.6) since it only includes projections for measured positions. Despite this, to mark its importance in later chapters, we make space to define the discrete ROI model as follows:

MODEL 3.2 (DISCRETE ROI-TOMOGRAPHY WITH NOISE) We model the attenuation of X-rays through the ROI, Ω , of an object described by its attenuation coefficients, \mathbf{x} , as follows:

$$\mathbf{b}_\Omega^\delta = A_\Omega \mathbf{x} + \mathbf{e}, \quad \text{for } \mathbf{b}_\Omega^\delta \in \mathbb{R}^m, \mathbf{x} \in \mathbb{R}^n \text{ and } A_\Omega \in \mathbb{R}^{m \times n}.$$

Here $(\mathbf{b}_\Omega^\delta)_i$ is a measured ROI projection from (2.4) with added Gaussian distributed white noise, $\mathbf{e} \in \mathbb{R}^m$ and the system matrix, A_Ω , contains rays intersecting Ω . The back projection is assumed to be A_Ω^T .

In Chapter 2.2 we discussed different regularisation methods to reconstruct an object from noisy data. Hence, we are motivated to try these methods on the problem formulated by Model 3.2 as we go through the next section.

3.2 An immediate reconstruction approach

In this section we try to reconstruct the Shepp–Logan phantom from a ROI-sinogram simulated using a system matrix generated by the parameters in Table 3.1. To simulate ROI data, we reduce the detector size such that we create an ROI illustrated in Figure 3.2. We keep the rest of the parameters the same as in Section 2.2.4 for comparison. We use three different reconstruction algorithms, FBP with a Hamming filter, Tikonov regularisation and the Landweber method.

To have an objective comparison measure, we use the relative error and mutual information as described in Section 2.1.4. However, we are mostly interested in the quality of the reconstructions near the ROI. Hence, we only consider the difference between the ground truth and the reconstruction in a disc with radius 1.3 times larger than the ROI giving us

$$\begin{aligned} RE_\Omega(X, Y) &:= RE(\chi_{\tilde{\Omega}} X, \chi_{\tilde{\Omega}} Y), \\ MI_\Omega(X, Y) &:= H_2(\chi_{\tilde{\Omega}} X, \chi_{\tilde{\Omega}} Y) - H_1(\chi_{\tilde{\Omega}} X) - H_1(\chi_{\tilde{\Omega}} Y), \end{aligned}$$

where χ_{Ω} is a mask with ones on a disc with radius 1.3 times the radius of ROI and zeros otherwise.

Table 3.1: Default parameters used to generate the system matrix, A_{Ω} , using a modified `fanbeamtomo` function with a linear detector. These parameters are used for the synthetic ROI set-up in Chapters 3 and 4.

Grid size (N)	256
Number of source locations	180
Number of detector pixels	256
Domain size	46 [cm] \times 46 [cm]
Source to centre distance	59 [cm]
Source to detector distance	100 [cm]
Detector length	41.1 [cm]

The ROI sinograms are shown in Figure 3.3. These are similar to the sinograms for full data, where the top and bottom are cut off. This is due to the structure of the sinograms where the upper and lower part correspond to the outermost projections for each source position.




Figure 3.2: Shepp-Logan phantom used for simulating ROI measurement data. The interior of the blue circle indicates the ROI.

We investigate the reconstructions of this data using the above mentioned reconstruction methods choosing the best regularisation parameters selected by minimising the ROI relative error and mutual information measures as done in Section 2.2.4. These reconstructions are displayed in Figure 3.4, where we observe a ring artefact about the ROI as the most pronounced structure in all reconstructions. This is not a singularity from the ground truth, but rather

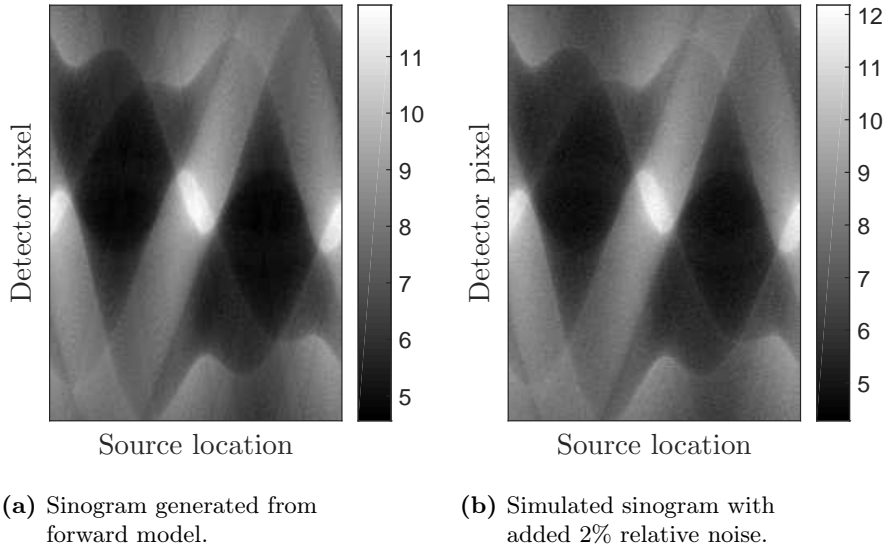


Figure 3.3: Sinograms generated from the object in Figure 3.2 using the measurement set-up in Table 3.1.

one that is added by the reconstruction. The attenuation coefficients do not fit the ground truth either. Finally most of the singularities outside the ROI have disappeared in the reconstructions.

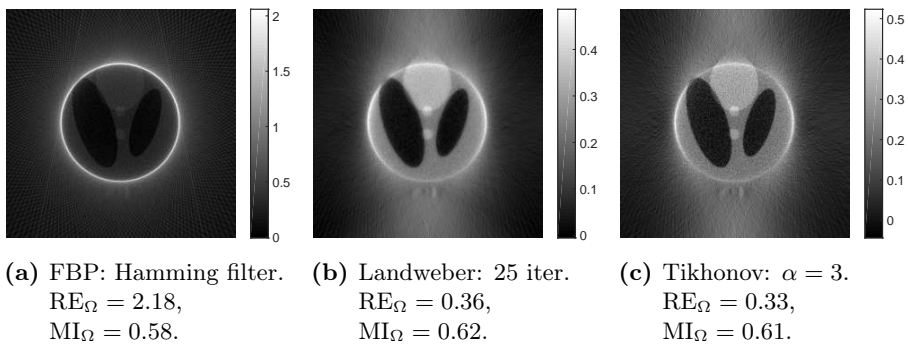


Figure 3.4: Comparison of standard reconstruction methods on the ROI data in Figure 3.3. The methods add a ring artefact on the ROI boundary and most singularities from the ground truth outside the ROI have disappeared.

3.3 Reflections

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

From studying the ROI model, we separate the challenges of ROI into two overarching types, that is, added artefacts and missing singularities in the reconstructions.

3.3.1 Added artefacts

In the previous section we illustrated some of the present challenges in ROI CT to be commented on. Consider the continuous ROI-tomography Model 3.1. When using filtered back projection from Theorem 2.4 on the ROI-sinogram, we make a hidden assumption by back projecting the filtered sinogram, $g_{\Omega,\Lambda} = \Lambda \mathcal{R}_\Omega(\theta, s)$ as follows

$$(\mathcal{R}^* g_{\Omega,\Lambda})(x) = \int_{S^1} g_{\Omega,\Lambda}(\theta, x \cdot \theta) \, d\theta.$$

To clarify, we implicitly set $g_{\Omega,\Lambda}(\theta, s) = 0$ when $\ell(\theta, s)$ does not intersect the ROI to well-define the integral. This assumes our forward operator to be

$$\mathcal{R}_\Omega f(\theta, s) = \begin{cases} (\mathcal{R}f)(\theta, s), & \text{for } \ell(\theta, s) \cap \Omega \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

So to speak, it assumes no attenuation outside the ROI enforcing all attenuation to be inside the ROI. This could explain the ring artefact in the reconstructions, as all attenuation is pressed to the boundary of the ROI.

However, we used the discrete model assumptions in Model 3.2 for our forward operator to avoid setting unknown projections to zero, but rather excluding them from the model. Yet the same artefact occurs in almost all reconstructions for varying parameters and methods in objects that have attenuation outside the ROI. In Figure 3.5 we showcase some of the reconstructions on *noise free* sinograms from the Shepp–Logan phantom using Tikhonov regularisation. In all cases singularities are added to the reconstruction, and except for an under-determined system with no regularisation, this singularity appears as the one expected from the continuous model. This calls for methods taking the possibility of added artefacts into account.




Figure 3.5: Reconstructed images using Tikhonov regularisation on a noise free version of the ROI sinogram in Figure 3.3b for varying grid size (under- and overdetermined) and regularisation parameters. Note the added artefacts and missing structure outside the ROI in the reconstruction even for noise free data and varying grid size.

3.3.2 Missing structure

Using the paradigm from MLA, we can determine which singularities are visible in the sinogram and hence which we expect to reconstruct. By Corollary 2.24 we find that any singularity with direction pointing towards the centre of rotation is invisible in a sinogram constructed by an ROI measurement set-up as in Figure 3.1. This is due to the fact that no rays passes through the singularity orthogonal to its direction.

This is shown to hold for the discrete set-up as well. In the reconstructions shown in Figures 3.4 and 3.5, any singularity of the ground truth that lie outside the ROI is visible only if measured by a ray passing through orthogonal to the direction of the singularity.

It is not clear which methodology is best suited to obtain these missing singularities. Two possibilities are to modify the data or add strong priors to the reconstruction algorithms. Regardless of the method, additional information, in addition to the data, is required. Another possibility is to change the measurement set-up, such that the singularities are directly obtained in the measurements.

Frame-based variational formulation

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

In this chapter we introduce a frame-based variational formulation for solving general CT problems. The idea is that prior knowledge of the objects structure and the measurement geometry can be incorporated into a weighted sparsity penalty term by a frame decomposition. Frame decompositions are known as strong noise reduction techniques for images assuming they are sparsely represented by the frame. The weights can be chosen to penalise certain frame elements depending on location and size. The method is tested with wavelets and shearlets on the region-of-interest tomography problem to determine if the method is suited for this type of problem.

4.1 Framework

In this section we introduce the necessary framework to develop frames for the purpose of describing objects by their frame decomposition. The aim is to present two frames, i.e., wavelets and shearlets, which are well-suited to represent piecewise constant objects and objects having singularities on curves respectively. To realise this, some fundamental definitions of frames are introduced, followed by the development of wavelets and shearlets, taking inspiration from the structure in [8]. As we go through the theory, figures help illustrate key properties of the frames.

4.1.1 Frames

We introduce some fundamental definitions and desired properties for a generalisation of orthonormal basis in terms of frames:

DEFINITION 4.1 A sequence, $\Phi = \{\varphi_\mu\}_{\mu \in \mathcal{M}}$, where \mathcal{M} is a discrete index space, is called a *frame* for a separable Hilbert space, X , if

$$\exists A, B > 0 : \quad A\|f\|^2 \leq \sum_{\mu \in \mathcal{M}} |\langle f, \varphi_\mu \rangle|^2 \leq B\|f\|^2 \quad \text{for all } f \in X.$$

For the different frames, we introduce the following related operators.

DEFINITION 4.2 The *Analysis operator* for $\Phi = \{\varphi_\mu\}_{\mu \in \mathcal{M}}$ is given by

$$\mathcal{T}_\Phi : X \rightarrow \ell^2(\mathcal{M}), \quad f \mapsto \{\langle f, \varphi_\mu \rangle\}_{\mu \in \mathcal{M}} = \{c_\mu\}_{\mu \in \mathcal{M}}.$$

The *Synthesis operator* for Φ is given by

$$\mathcal{T}_\Phi^* : \ell^2(\mathcal{M}) \rightarrow X, \quad \{c_\mu\}_{\mu \in \mathcal{M}} \mapsto \sum_{\mu \in \mathcal{M}} c_\mu \varphi_\mu.$$

Here, the analysis operator describes how to decompose an object into frame coefficients $\{c_\mu\}_{\mu \in \mathcal{M}}$. The synthesis operator then describes how to reconstruct an object that is represented by coefficients for a given frame.

The inverse of the analysis operator is the synthesis operator if an isometry between X and $\ell^2(\mathcal{M})$, expressed by the synthesis and analysis operator, exists, i.e., if $\mathcal{T}_\Phi^* \mathcal{T}_\Phi = I_X$. The way we achieve this is by introducing the following properties.

DEFINITION 4.3 A *tight frame* for a separable Hilbert space, X , is a sequence of functions, $\Phi = \{\varphi_\mu\}_{\mu \in \mathcal{M}}$, satisfying Definition 4.1 for $A = B$ giving us the Parseval-type relation for any $f \in X$:

$$\|f\|_X^2 = \sum_{\mu \in \mathcal{M}} |\langle f, \varphi_\mu \rangle|^2,$$

motivating the alias *Parseval frame*. Furthermore, for a tight frame we get the familiar reproducing formula (see e.g. p. 10 in [8]):

$$f = \mathcal{T}_\Phi^* \mathcal{T}_\Phi f = \sum_{\mu \in \mathcal{M}} \langle f, \varphi_\mu \rangle \varphi_\mu.$$

Tight frames give a generalisation of orthonormal basis, maintaining the important reproducing formula. We are now interested in finding specific tight frames to give a sparse representation of certain objects.

4.1.2 Wavelets

We motivate and introduce *wavelets* being the frontrunner of efficiently encoding piecewise regular functions in $L^2(\mathbb{R}^2)$. Throughout, we make use of the following operators.

DEFINITION 4.4 For functions in $L^2(\mathbb{R}^2)$ we define the *translation operator* as

$$T_t : L^2(\mathbb{R}^2) \rightarrow L^2(\mathbb{R}^2), \quad \phi \mapsto \phi(\cdot - t),$$

and the *dilation operator* as

$$D_M : L^2(\mathbb{R}^2) \rightarrow L^2(\mathbb{R}^2), \quad \phi \mapsto |\det M|^{-1/2} \phi(M^{-1}\cdot)$$

given a dilation matrix, $M \in \mathbb{R}^{2 \times 2}$.

These operators are quite simple yet fundamental for generalising a basis from frames. We readily define the general wavelet systems as:

DEFINITION 4.5 We consider the set of functions defined by

$$\Phi_\phi = \{\phi_{t,M} = T_t D_M^{-1} \phi = |\det M|^{1/2} \phi(M(\cdot - t)) : (M, t) \in G \times \mathbb{R}^2\}$$

where D_M and T_t are given as in Definition 4.4, and G is a subspace of invertible matrices. If the system, Φ , forms a tight frame on $L^2(\mathbb{R}^2)$, we call ϕ a wavelet.

Note that a wavelet is simply a function for which a certain collection of translated and scaled versions form a tight frame. Wavelets are far from being unique and many different wavelets have been developed with different properties. In this thesis we consider the *Haar wavelet*:

DEFINITION 4.6 The 1D Haar wavelet is defined by

$$\phi(t) = \begin{cases} 1 & 0 \leq t < \frac{1}{2} \\ -1 & \frac{1}{2} \leq t < 1 \\ 0 & \text{otherwise,} \end{cases}$$

and the 1D Haar scaling function is defined by

$$\phi_s(t) = \begin{cases} 1 & 0 \leq t < 1, \\ 0 & \text{otherwise.} \end{cases}$$

We can extend the Haar wavelet in two dimensions in the following way:

DEFINITION 4.7 The 2D Haar wavelets are defined by

$$\begin{aligned}\phi_H &= \phi_s \otimes \phi, \\ \phi_V &= \phi \otimes \phi_s, \\ \phi_D &= \phi \otimes \phi,\end{aligned}$$

and the 2D Haar scaling function is defined by

$$\phi_S = \phi_s \otimes \phi_s.$$

Here \otimes denotes the product $(u \otimes v)(x_1, x_2) = u(x_1)v(x_2)$ and the subscripts H, V and D denote the horizontal, vertical and diagonal wavelets respectively. These four 2D wavelets are illustrated in Figure 4.1.

We can then define the 2D Haar wavelet system as

$$\mathcal{H} = \Phi_{\phi_H} \cup \Phi_{\phi_V} \cup \Phi_{\phi_D} \cup \Phi_{\phi_S}. \quad (4.1)$$

Note that the wavelet system using the 1D Haar wavelets form an orthonormal basis for $L^2(\mathbb{R})$, (see e.g. p. 161 in [9]). Similarly the 2D wavelet system, $\Phi_{\mathcal{H}}$, form an orthonormal basis for $L^2(\mathbb{R}^2)$ and hence a tight frame (see e.g. Chapter 10 in [10]).

These simple wavelets tend to represent piecewise regular functions well and is one of the two frames considered throughout the thesis.




Figure 4.1: Illustration of 2D Haar wavelets on a $[0, 1]^2$ grid. Here white represents the value 1 and black (-1) .

4.1.3 Shearlets

In the previous subsection we introduced the Haar wavelet as a well performing frame for piecewise regular functions. In this subsection we form a frame tailored

for structures with singularities around smooth curves such as the pipe from FORCE. In Figure 4.2 we illustrate how a frame with directional information is able to represent a smooth boundary with fewer frame elements compared to a frame with no directional component such as wavelets.

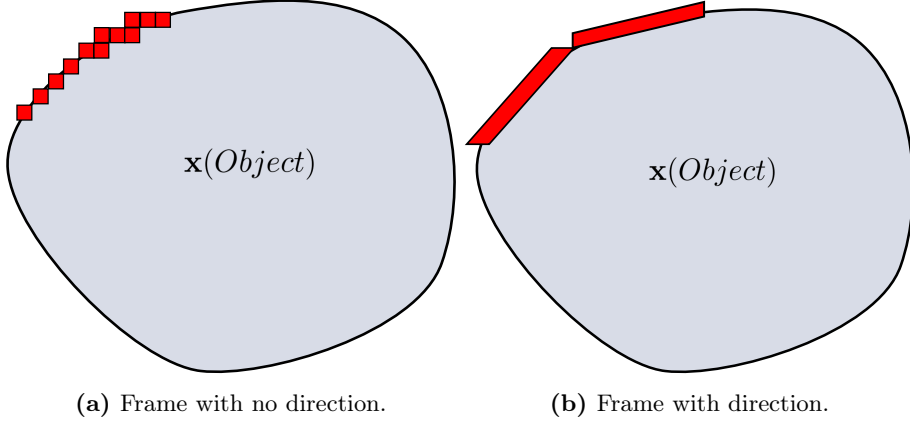


Figure 4.2: Comparing the number of frame elements necessary to represent a singularity on a smooth curve for frames with or without directional information.

With this in mind we introduce *shearlets* as follows.

DEFINITION 4.8 For $\psi \in L^2(\mathbb{R}^2)$ the continuous shearlet system, SH_ψ , is defined by

$$\text{SH}_\psi = \{\psi_{a,s,t} = T_t D_{A_a} D_{S_s} \psi : a > 0, s \in \mathbb{R}, t \in \mathbb{R}^2\}.$$

Here T_t , D_{A_a} and D_{S_s} are given as in Definition 4.4 where

$$A_a = \begin{bmatrix} a & 0 \\ 0 & a^{1/2} \end{bmatrix}, \quad S_s = \begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix}, \quad \text{for } a, s > 0.$$

The difference between the shearlet system and the wavelet system is the use of the shearing operator before applying dilation. This enables the shearlets to have a direction, making them more suited for general functions with smooth singularities.

Currently, the way to construct a tight shearlet frame is by enforcing certain properties in the frequency domain. Because of the Fourier isometry of $L^2(\mathbb{R}^2)$, it is equivalent to constructing a tight frame in $L^2(\mathbb{R}^2)$ and the Fourier representation of it. The *classical shearlet* system is a common way of doing this.

It is defined using the Fourier transform of a wavelet and a bump function as follows:

DEFINITION 4.9 Let $\psi \in L^2(\mathbb{R}^2)$ be defined by

$$\hat{\psi}(\xi) = \hat{\psi}(\xi_1, \xi_2) = \hat{\psi}_1(\xi_1) \hat{\psi}_2\left(\frac{\xi_2}{\xi_1}\right),$$

where $\psi_1 \in L^2(\mathbb{R})$ is a discrete wavelet satisfying the discrete Calderon condition given by

$$\sum_{j \in \mathbb{Z}} |\hat{\psi}_1(2^{-j}\xi)|^2 = 1, \quad \xi \in \mathbb{R} \text{ a.e.},$$

with $\hat{\psi}_1 \in C^\infty(\mathbb{R})$ and $\text{supp}(\hat{\psi}_1) \subseteq [-\frac{1}{2}, \frac{-1}{16}] \cup [\frac{1}{16}, \frac{1}{2}]$, and $\psi_2 \in L^2(\mathbb{R})$ is a bump function in the sense that

$$\sum_{k=-1}^1 |\hat{\psi}_2(\xi + k)|^2 = 1, \quad \xi \in [-1, 1] \text{ a.e.},$$

satisfying $\hat{\psi}_2 \in C^\infty(\mathbb{R})$ and $\text{supp}(\hat{\psi}_2) \subseteq [-1, 1]$. Then ψ is called a *classical shearlet*. For more details see [11].

The classical shearlet system forms a tight frame as stated in the following theorem:

THEOREM 4.10 (P. 20 IN [8]) Let $\psi \in L^2(\mathbb{R}^2)$ be a classical shearlet as in Definition 4.9, then SH_ψ is a tight frame for $L^2(\mathbb{R}^2)$.

Even though the classical shearlet system forms a tight frame, some functions are not well represented by it as shown by the example below:

The support in frequency domain of the classical shearlet is illustrated for three level of shearings in Figure 4.3. Consider a function in $L^2(\mathbb{R}^2)$ with support along the ξ_1 -direction. Then the function is only represented by shearlets for $s \rightarrow \infty$. These shearlets would become infinitely narrow. Hence, we motivate to decompose the shearlet system into parts in regards to the frequency directions as follows.

DEFINITION 4.11 For $\phi, \psi, \tilde{\psi} \in L^2(\mathbb{R}^2)$ the *cone adapted shearlet system* $\text{SH}_{\phi, \psi, \tilde{\psi}}$ is defined by

$$\text{SH}_{\phi, \psi, \tilde{\psi}} = \Phi_\phi \cup \Psi_\psi \cup \tilde{\Psi}_{\tilde{\psi}},$$




Figure 4.3: The support of shearlets in frequency domain for different shearings.

where

$$\Phi_\phi = \{\phi_t = \phi(\cdot - t) : t \in \mathbb{R}^2\},$$

$$\Psi_\psi = \{\psi_{a,s,t} = a^{-\frac{3}{4}} \psi(A_a^{-1} S_s^{-1}(\cdot - t)) : a \in (0, 1], |s| \leq 1 + a^{1/2} t \in \mathbb{R}^2\},$$

$$\tilde{\Psi}_{\tilde{\psi}} = \{\tilde{\psi}_{a,s,t} = a^{-\frac{3}{4}} \tilde{\psi}(\tilde{A}_a^{-1} S_s^{-1}(\cdot - t)) : a \in (0, 1], |s| \leq 1 + a^{1/2} t \in \mathbb{R}^2\},$$

$$\text{and } \tilde{A}_a = \begin{bmatrix} a^{1/2} & 0 \\ 0 & a \end{bmatrix}.$$

For the cone adaptive system the *shearlet scaling function*, ϕ , is chosen to have compact frequency support near the origin, which ensures that the system Φ_ϕ is associated with the low-frequency region $\mathcal{C}_0 = \{(\xi_1, \xi_2) : |\xi_1|, |\xi_2| \leq 1\}$ illustrated in Figure 4.4. By choosing ψ to satisfy the conditions of a classical shearlet in Definition 4.9, the system Ψ_ψ is associated with the vertical cones $\mathcal{C}_1 \cup \mathcal{C}_3 = \{(\xi_1, \xi_2) : |\xi_2/\xi_1| \leq 1, |\xi_1| > 1\}$. The shearlet $\tilde{\psi}$ can be chosen likewise with the roles of ξ_1 and ξ_2 reversed, i.e. $\hat{\psi}(\xi_1, \xi_2) = \hat{\psi}(\xi_2, \xi_1)$. Then the system $\tilde{\Psi}_{\tilde{\psi}}$ is associated with the horizontal cones $\mathcal{C}_2 \cup \mathcal{C}_4 = \{(\xi_1, \xi_2) : |\xi_2/\xi_1| > 1, |\xi_2| > 1\}$.

One can construct a tight frame using the cone adaptive shearlet system as shown on page 28 in [8]. Since the corresponding shearlets are compact in frequency domain, they are not compactly supported in the spatial domain. This is an undesired property if one wants to decompose spatially compact




Figure 4.4: Illustration of the cone separation scheme in frequency domain.

objects by shearlets. Hence, we are motivated to construct a shearlet system with compactly supported shearlets in spatial domain. Such a shearlet system, based on the cone adaptive system, is constructed numerically in ShearLab [12]. The system does not form a tight frame, but has shown to retain a stable reproducing formula in [8]. Here the following generating functions are used for ψ .

$$\hat{\psi}_1(x) = \begin{cases} 1 & |x| \leq \frac{1}{4}, \\ \cos[\frac{\pi}{2}\nu(\frac{4}{3}|x| - \frac{1}{3})] & \frac{1}{4} \leq |x| \leq 1, \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\hat{\psi}_2(x) = \begin{cases} \sin[\frac{\pi}{2}\nu(\frac{4}{3}|x| - \frac{1}{3})] & \frac{1}{4} \leq |x| \leq 1, \\ \cos[\frac{\pi}{2}\nu(\frac{1}{3}|x| - \frac{1}{3})] & 1 \leq |x| \leq 4, \\ 0 & \text{otherwise} \end{cases}$$

where $\nu \geq 0$ is a $C^2(\mathbb{R})$ function such that $\nu(x) + \nu(1-x) = 1$ for $0 \leq x \leq 1$. The scaling function, $\hat{\phi}$, is then defined dependent on $\hat{\psi}_1$ and $\hat{\psi}_2$. Some examples of the shearlets are illustrated in Figure 4.12 in Section 4.4. Further details of the implementation for this shearlet system is out of scope for the thesis and we refer to [13].

4.2 The variational formulation

Main author: Nicolai André Brogaard Riis.

Co-author: Jacob Frøsig.

In this section we formulate a frame-based variational formulation. The idea behind the formulation is to include a regularisation term that enforces sparsity of the object in whichever frame is chosen, *in addition to* an existing data fidelity term. This form of variational formulation is general and can be used for most linear inverse problems. For this thesis, however, we focus on the ROI tomography problem. That is, the method is based on computing the regularised solution for the object, $\mathbf{x} \in \mathbb{R}^n$, to the discrete ROI problem in Model 3.2:

$$\mathbf{b}_\Omega^\delta = A_\Omega \mathbf{x} + \mathbf{e}. \quad (4.2)$$

Given some discrete frame system, $\Phi = \{\varphi_\mu\}_{\mu=1}^M$, with frame elements $\varphi_\mu \in \mathbb{R}^n$ and corresponding analysis operator, \mathcal{T}_Φ , the approximate solution can be found by the variational formulation:

$$\mathbf{x}^*|_{\alpha, w, \Phi, p} = \operatorname{argmin}_{\mathbf{x} \in F_M} \left\{ \|A_\Omega \mathbf{x} - \mathbf{b}_\Omega^\delta\|_2^2 + \alpha \|\mathcal{T}_\Phi \mathbf{x}\|_{p, w}^p \right\}. \quad (4.3)$$

Here $\|\cdot\|_{p, w}^p$ is a weighted ℓ_p -norm with strictly positive weights $w_\mu > 0$ on each element given by

$$\|\mathcal{T}_\Phi \mathbf{x}\|_{p, w}^p = \sum_{\mu=1}^M w_\mu |\langle \mathbf{x}, \varphi_\mu \rangle|^p.$$

For $1 \leq p < 2$ the regularisation term enforces sparsity (see e.g. [14]). The minimiser of (4.3) is unique for $p = 1$ if A_Ω is injective [15]. It was shown in [15], [16] that the continuous ROI Radon transform, \mathcal{R}_Ω , is injective assuming the given data is continuous and the object, f , is piecewise constant (PC) and compactly supported in \mathbb{R}^2 . In the discrete case A_Ω might not be injective.

For the above variational formulation it is assumed that \mathbf{x} has a finite decomposition (of M elements) with respect to the basis or frame explained by the assumption on the space:

$$F_M := \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} = \sum_{\mu=1}^M \langle \mathbf{x}, \varphi_\mu \rangle \varphi_\mu \right\}.$$

REMARK. The assumption that \mathbf{x} is finitely decomposed by the frame makes it such that we only need finitely many coefficients to store \mathbf{x} in its entirety. It

follows from an underlying assumption that the object we consider has a sparse representation in our chosen frame.

When Φ forms a tight frame, we can express (4.3) in terms of its frame coefficients, $c_\mu = \langle \mathbf{x}, \varphi_\mu \rangle$, as a vector $\mathbf{c} \in \mathbb{R}^M$ using the synthesis operator, \mathcal{T}_Φ^* . The formulation is as follows:

$$\mathbf{c}^*|_{\alpha, w, \Phi, p} = \operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^M} \left\{ \|A_\Omega \mathcal{T}_\Phi^* \mathbf{c} - \mathbf{b}_\Omega^\delta\|_2^2 + \alpha \sum_{\mu=1}^M w_\mu |c_\mu|^p \right\}. \quad (4.4)$$

The regularised solution to the frame-based variational formulation (4.3) is then simply found by applying the synthesis operator to the regularised coefficients, i.e., $\mathbf{x}^*|_{\alpha, w, \Phi, p} = \mathcal{T}_\Phi^* \mathbf{c}^*|_{\alpha, w, \Phi, p}$.

4.2.1 Choice of parameters

There are generally four classes of parameters that need to be chosen when solving (4.3) using frame-based regularisation: The regularisation parameter, the frame system, the weights and the norm. For the purposes of this thesis we set the norm parameter $p = 1$.

Frame system

The choice of frame system is crucial to the quality of the reconstructed solution. The goal is to choose a frame that represents the object with few coefficients, while not representing noise and artefacts that would otherwise be generated in the reconstruction. The choice of frame is hence dependent on the problem at hand, typically defined by the measurement geometry and type of object.

Regularisation parameter

The regularisation parameter, denoted α , determines how sparsely the reconstructed image is represented in a specific frame system. If the regularisation parameter is chosen too small, the solution can be dominated by noise in the data. If it is chosen too large, then the object is no longer well represented by the few number of frame coefficients and details of the ground truth will be lost. A choice in-between these is thus preferable.

Scale dependent weights

Using scale dependent weights is a way to penalise frame elements differently depending on their scale. In this case, for frame elements at scales $2^0, 2^1, \dots, 2^J$, we set the scale weight by

$$w_j = 2^{-j+1}, \quad j = 1, \dots, J. \quad (4.5)$$

This can be understood as a heuristic: The frame elements that are twice as large must be penalised half as much, effectively penalising high frequency information. This weight scheme has proven an effective noise reduction technique for curvelets in [17] and this form of level dependent weights also show up in wavelet regularisation using Besov norms (see e.g. [18]).

Location based weights: Region-of-Interest

In [15] the idea is to choose a weight sequence for wavelet regularisation such that, for those basis functions that are only supported in the ROI, Ω , the weight 1 is assigned, and for those only supported only outside some outer weight, $w_{out} > 1$ is assigned. In the case that the wavelet is supported both inside and outside Ω , we assign weight depending on the shared support of the wavelet, interpolating between 1 and w_{out} .

We can generalise this idea for a given frame system, $\Phi = \{\varphi_\mu\}_{\mu=1}^M$, formally writing it in the following way: Let $w_{out} > 1$, we define the weight for each frame element, $\varphi_\mu \in \mathbb{R}^n$, as

$$w_\mu = \gamma(V_\Omega(\varphi_\mu)).$$

Here $V_\Omega(\varphi_\mu)$ is the proportional support of $\varphi_\mu \in \mathbb{R}^n$ inside the ROI, Ω , given by

$$V_\Omega(\varphi_\mu) = \frac{\|\chi_\Omega \varphi_\mu\|_2}{\|\varphi_\mu\|_2}. \quad (4.6)$$

Here the mask, χ_Ω , sets all elements in Ω to 1 and 0 otherwise. Finally γ is a linear interpolation function, $\gamma(t) = t + (1 - t)w_{out} \in [1, w_{out}]$. The weight scheme is most easily illustrated by considering the mask, χ_Ω , as shown in Figure 4.5a. Each frame is then assigned a weight depending on its location and support on this mask. An example of the mask applied on a diagonal 2D Haar wavelet is shown in Figure 4.6. The proportional support inside the ROI is then calculated by (4.6).

Location based weights: Information measure

An extension of the above idea is to let the mask vary depending on the normalised density of rays going through a particular pixel. This is calculated by summing the length of all rays' trajectories through the pixel and dividing by the maximum. Using this type of mask, we assume that the information at a particular point decreases as fewer rays go through it.

In the case of ROI measurements, we set frame elements that are supported only inside the ROI to weight 1, and those supported only outside the ROI have weights that decrease as they go further away from the ROI. This can be described by changing the mask in (4.6). The mask now depends on the density of the rays, denoted \mathcal{C}_Ω , as illustrated in Figure 4.5b. Similarly to before, the weight is then calculated using (4.6) with the new mask. An illustration of the mask's influence on a diagonal 2D Haar wavelet is shown in Figure 4.6.

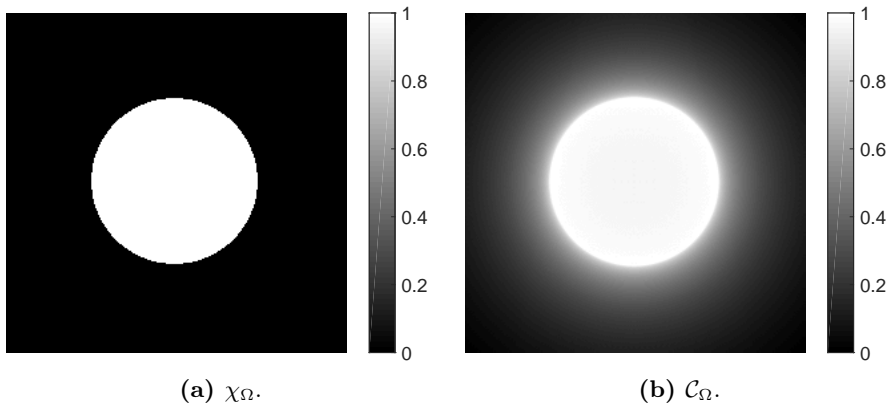


Figure 4.5: Illustration of the masks used to calculate the proportion $V_\Omega(\varphi_\mu)$ in (4.6). This is then used to decide the weight for each frame element using linear interpolation between 1 and some outer weight, w_{out} .

4.2.2 Fast Iterated Soft-Thresholding Algorithm (FISTA)

We now describe an algorithm for solving (4.4) for $p = 1$. Essentially the problem becomes the optimisation of the ℓ^1 -penalised Tikhonov functional for a given measurement \mathbf{b}_Ω^δ :

$$\mathbf{c}^*|_{\alpha, w, \Phi} = \underset{\mathbf{c} \in \mathbb{R}^M}{\operatorname{argmin}} \left\{ \frac{1}{2} \|A_\Omega \mathcal{T}_\Phi^* \mathbf{c} - \mathbf{b}_\Omega^\delta\|_2^2 + \alpha \|W \mathbf{c}\|_1 \right\}. \quad (4.7)$$




Figure 4.6: Illustrating the effect of applying the masks from Figure 4.5 on a diagonal 2D Haar wavelet. The proportion, $V_\Omega(\phi_D)$, is then calculated by taking the norm of the images with and without the mask from (4.6).

Here $W : \mathbb{R}^M \rightarrow \mathbb{R}^M$ is an operator such that $Wc = [w_1 c_1, \dots, w_\mu c_\mu]$.

Define for convenience:

$$\begin{aligned}\mathcal{K} &= A_\Omega \mathcal{T}_\Phi^*, \\ \mathcal{K}^* &= \mathcal{T}_\Phi A_\Omega^T.\end{aligned}$$

The method of *iterated soft-thresholding* (ISTA), (also known as *Method of Surrogate Functionals* [15]) for solving (4.7), is determined by the fixed point equation: (see e.g. [19],[14] for more details)

$$\mathbf{c}^*|_{\alpha,w,\Phi} = S_{s\alpha w}(\mathbf{c}^*|_{\alpha,w,\Phi} - s\mathcal{K}^*(\mathcal{K}\mathbf{c}^*|_{\alpha,w,\Phi} - \mathbf{b}_\Omega^\delta)), \quad (4.8)$$

where $S_\kappa(x) = \text{sgn}(x) \max(|x| - \kappa, 0)$ and $s > 0$ is a step length. It can be shown, as in [19], that iteratively updating from (4.8) converges to a minimizer of (4.7), provided that the step length satisfies

$$0 < s < 2/\|\mathcal{K}\|_2^2. \quad (4.9)$$

The FISTA algorithm was recently developed as a faster converging method of the ISTA algorithm in the paper [20]. It is based on the same fixed-point equation, but uses a linear combination of the previous points for each update. The paper shows that the ISTA method has a worst-case convergence rate of the objective function of $O(1/k)$ where k is the iteration number. The paper similarly show that the new FISTA method has a worst-case convergence rate of the objective function of $O(1/k^2)$. The FISTA algorithm is shown in Algorithm 1.

Algorithm 1: Fast Iterated Soft-Thresholding Algorithm (FISTA)

Input: Noisy data \mathbf{b}_Ω^δ , ROI Radon Transform A_Ω , A_Ω^T and analysis, \mathcal{T}_Φ , and synthesis, \mathcal{T}_Φ^* , operators. $\mathcal{K} = A_\Omega \mathcal{T}_\Phi^*$, $\mathcal{K}^* = \mathcal{T}_\Phi A_\Omega^T$.

Output: Reconstruction, \mathbf{x} , that is the approximate solution to (4.2).

Assume $\mathbf{x} \in F_M$ for some $M \in \mathbb{N}$.

Choose s by (4.9) and some weights $w = \{w_\mu\}_{\mu=1,\dots,M}$

Choose regularisation parameter α .

Initiate start guess $\mathbf{c}^{[0]}$. Let $\mathbf{y}^{[1]} = \mathbf{c}^{[0]}$ and $t_1 = 1$, $k = 1$.

```
while  $k < k_{max}$  do
   $\mathbf{c}^{[k]} \leftarrow S_{s\alpha w}(\mathbf{y}^{[k]} - s\mathcal{K}^*(\mathcal{K}\mathbf{y}^{[k]} - \mathbf{b}_\Omega^\delta))$ 
   $t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2}$ 
   $\mathbf{y}^{[k+1]} = \mathbf{c}^{[k]} + \frac{t_k - 1}{t_{k+1}}(\mathbf{c}^{[k]} - \mathbf{c}^{[k-1]})$ 
   $k = k + 1$ 
```

end

```
return  $\mathbf{x} = \mathcal{T}_\Phi^* \mathbf{c}^{[k_{max}]}$ 
```

REMARK. In a special course at DTU, the ISTA and FISTA algorithms were compared with an ADMM algorithm (see Appendix A.1 for derivations of ADMM update steps for (4.3)). The results showed that, while giving the same result in the end, the ADMM method was too sensitive to the parameter choice of the inner iterative solver. In addition, the FISTA method showed a much faster convergence rate numerically compared to ISTA as expected. Hence going forward the FISTA algorithm was used as the main tool for solving (4.3). We do note, that if a more stable solver for the inner update is found, the ADMM method might converge faster.

4.3 Weighted wavelet sparsity penalty

Main author: Nicolai André Brogaard Riis.

Co-author: Jacob Frøsig.

In this section we propose a weighted Haar wavelet sparsity penalty for the variational formulation in (4.3). This was first studied in [15] with weights depending on the relative location of each wavelet. We add to this by allowing the weights to depend on scale as well. Hence, the weights on the wavelet coefficients can be used to modify the penalty term depending on both the scale and relative location of each wavelet.

Recall, that the frame-based methods solve the ROI problem in Model 3.2:

$$\mathbf{b}_\Omega^\delta = A_\Omega \mathbf{x} + \mathbf{e},$$

using the variational formulation

$$\mathbf{c}^*|_{\alpha, w, \mathcal{H}} = \operatorname{argmin}_{\mathbf{c} \in \mathbb{R}^M} \left\{ \|A_\Omega \mathcal{T}_{\mathcal{H}}^* \mathbf{c} - \mathbf{b}_\Omega^\delta\|_2^2 + \alpha \sum_{\mu=1}^M w_\mu |c_\mu| \right\}. \quad (4.10)$$

Here $\mathcal{T}_{\mathcal{H}}^*$ and $\mathcal{T}_{\mathcal{H}}$ are the Haar wavelet synthesis and analysis operators, respectively. The regularised solution is, as before, simply found by applying the synthesis operator to the regularised coefficients, i.e., $\mathbf{x}|_{\alpha, w, \mathcal{H}} = \mathcal{T}_{\mathcal{H}}^* \mathbf{c}^*|_{\alpha, w, \mathcal{H}}$.

We assume in this method that our object, \mathbf{x} , is finitely decomposed with respect to the discrete 2D Haar wavelet basis denoted \mathcal{H} with Haar wavelets denoted $\phi_{j,t} \in \mathbb{R}^n$. Here $\phi_{j,t}$ describes every wavelet from (4.1) stored as vectors where j and t are the scaling and translation parameters, respectively. Using the notation from earlier, we thus write $\mathbf{x} \in H_J$, where

$$H_J := \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} = \sum_{K \in \{H, V, D, S\}} \sum_{j=1}^{J(K)} \sum_{t=1}^{T(j, K)} \langle \mathbf{x}, \phi_{j,t; K} \rangle \phi_{j,t; K} \right\},$$

where $T(j, K)$ is the number of translations at scale j for the wavelet type K and $J(K)$ is the number of scales for the wavelet type K . For a given image size the number of translations T is fixed for each scale. These parameters are chosen such that the number of wavelets are fixed by

$$M = \sum_{\{k \in H, V, D, S\}} \sum_{j=1}^{J(K)} T(j, K) = n. \quad (4.11)$$

4.3.1 Numerical experiments

The following numerical experiments are meant to give an illustration of how the weighted wavelet regularisation method performs on various test problems with different choices of parameters. Unfortunately there seem to be no "golden rule" for how these parameters must be chosen since it depends heavily on the problem. The experiments are only meant to illustrate the points that we felt were most essential to explain the behaviour of the method. Many factors, such as the placement of ROI are not considered in this section. One thing that should not influence the final image, however, is how the optimisation problem is solved. We use the FISTA algorithm as described in Algorithm 1 to solve (4.10). The ADMM method as described in Algorithm 2 in Appendix A.1 is then used to validate that the images are the same to some numerical accuracy.

For the wavelet synthesis and analysis operators, we use the Matlab implementation of the discrete 2D Haar wavelet system using the functions `wavedec2.m` and `waverec2.m`. The scaling coefficients, $\langle \mathbf{x}, \phi_S \rangle$, are not included in the penalty term by setting their respective weights to 0. The wavelet parameters are fixed using the largest number of scales possible for a given image size unless stated otherwise. This is calculated by the Matlab function `wmaxlev.m`. This number of scales yield the same number of wavelet coefficients as elements in \mathbf{x} as indicated in (4.11).

The ROI system matrix, A_Ω , is generated using the parameters from Table 3.1 as described in Section 3.2 unless stated otherwise. The data is thus as shown in Figure 3.3b from Section 3.2.

The optimal regularisation parameter is calculated by comparing the reconstructed image to the ground truth in these synthetic experiments. The comparison is done using ROI relative error, RE_Ω , and mutual information, MI_Ω , as in Section 3.2 and by visual inspection. The idea here is to showcase the effect of varying the other parameters using a near optimal choice of regularisation parameter. However, depending on what the reconstructed image of the object is intended for, one measure of image quality might be preferable over another. Visual inspection takes precedence over the measures. The visual inspection is focused mainly on singularities.

Wavelet phantom

We first consider the weighted wavelet regularisation method on an object which is represented well by the wavelet basis. To increase the difficulty of the problem, the detector size in Table 3.1 is decreased to 20 cm making the ROI smaller. The wavelet phantom is shown in Figure 4.7a and the ROI is indicated in blue. This phantom can be represented by few wavelet coefficients since it directly consists of discrete 2D Haar wavelets. The best reconstruction, using the standard methods described in Section 2.2.4 on the ROI Model 2.6, was found by the Landweber method after 200 iterations and is shown in Figure 4.7b. We try the wavelet regularisation method using 2 and 8 scales respectively. The best reconstructions were found for regularisation parameters 10 and 1 for these scales and are shown in Figures 4.7c and 4.7d. There is a clear difference between the wavelet method and the best reconstruction using the standard methods. This is not surprising since the phantom is made of 2D wavelets. This showcases that choosing the right frame for the object can give a significant increase in image quality. At the same time, we see how artefacts outside the ROI are generated on the wavelet reconstruction using only 2 scales, whereas the reconstruction using 8 scales do not show these artefacts to the same degree. We believe this is because the homogeneous region outside the phantom is well represented by a few large wavelets.

Adding scale weights

We consider the scale weights in (4.5) and how they change quality of the reconstruction for the Shepp–Logan phantom. Recall our methodology is to first choose the wavelet system and weight parameters and then find the optimal regularisation parameter. In Figure 4.8 we see that adding scale weights makes the ring artefact less pronounced in the reconstructions for both noisy and noise free data. This suggests that the ring artefact is less significant than the singularities from the ground truth in the data fitting, and hence it is removed when penalising smaller wavelets.

Regularisation parameter

The effect of varying the regularisation parameter is illustrated in Figure 4.9. If the parameter is too low, then noise dominates the solution. If the parameter is too high, then the object is represented by too few wavelets thus losing structural details from the ground truth. If chosen just right, then the solution represents the ground truth well. Note that scale weights are used in all the reconstructions.

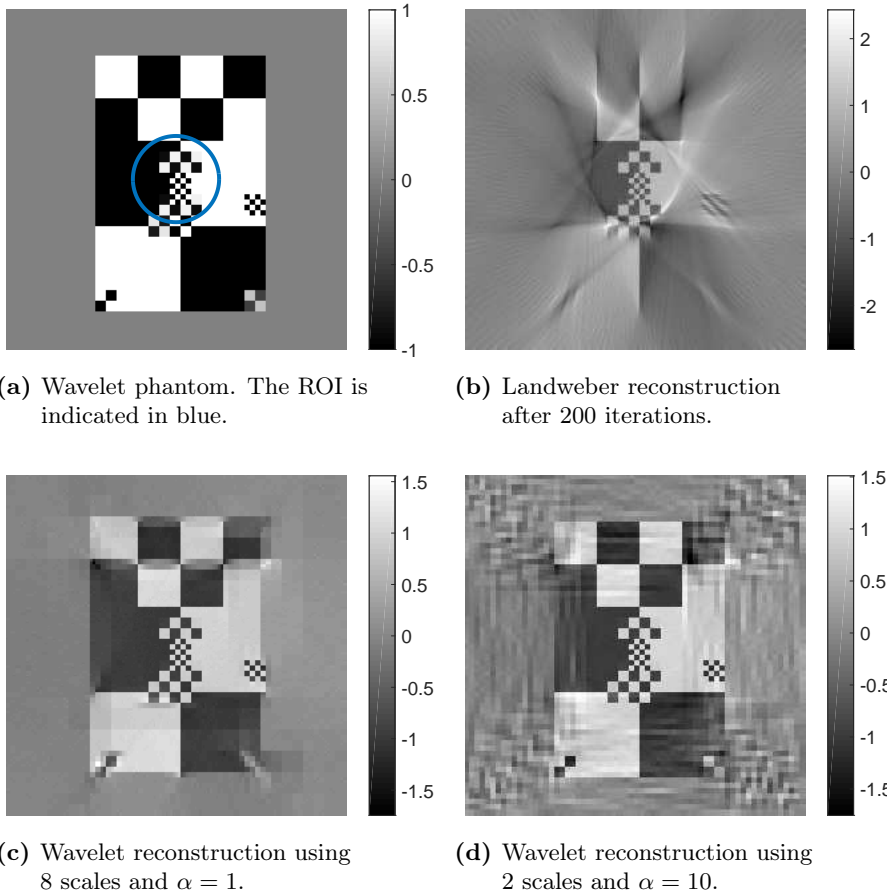


Figure 4.7: Comparing wavelet reconstructions with the Landweber method on a 256×256 phantom which is represented well by wavelets. The ROI data is generated with fan-beam projections that fully illuminate the region indicated by blue in the original phantom. The data consists of 180 projections all around the object with 256 detector pixels in each and has added 2% relative noise.

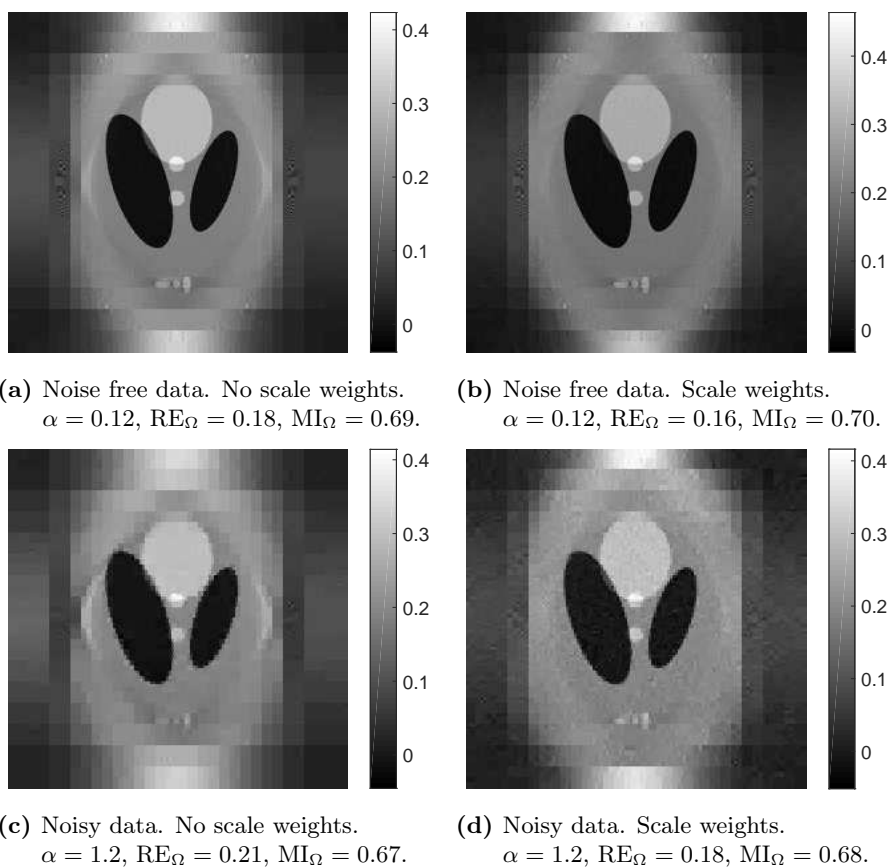


Figure 4.8: Comparing the effect of scale weights on the wavelet coefficients on the Shepp–Logan phantom from the data in Figure 3.3b in Section 3.2 and on a noise free version of the sinogram. The scale weights promote larger wavelets over smaller ones, removing small details that do not fit the data as well.




Figure 4.9: Wavelet based reconstructions on the ROI data in Figure 3.3b from Section 3.2 for varying regularisation parameter. The method uses scale dependent weights and has 8 scales.

Adding location based weights

We finally consider the effect of adding location based weights using the schemes shown in Section 4.2. In Figure 4.11 we see the ROI relative error, RE_{Ω} , and mutual information, MI_{Ω} , measures for varying regularisation parameter and outer weight w_{out} on the data in Figure 3.3b. No improvement is shown by adding location based weights for this particular set-up, as long as the regularisation parameter is chosen near the optimum. The results hold for both the ROI and information based weighting schemes. In addition, this was also tested for other phantoms using the same system matrix generated from Table 3.1 and for smaller increases in outer weight. For all tests there was seen no improvement in image quality by location based weights, even with and without scale dependent weights. We note that some problems might benefit from location based weights, but we have found no improvements in neither error measures nor by visual inspection. In Figure 4.10 we see the reconstructions given the optimal regularisation parameter chosen by the error measures. Note that details are removed outside the ROI by increasing the outer weight. A lower choice of regularisation parameter, in turn, increases the frequency inside the ROI which, in this case, is characterised by noise. We note that for noise free data, with many details inside the ROI the locations based weights might be beneficial.




Figure 4.10: Change in image quality by adding location based weights an addition to the existing scale dependent weights for the weighted wavelet regularisation method on the data in Figure 3.3b. We see that there is no obvious improvement to be found by adding location based weights, although more high frequent elements are added inside the ROI when using a lower regularisation parameter and higher weight.




Figure 4.11: Error measure using weighted wavelet regularisation for location based weights, w_{out} , while varying the regularisation parameter on the data in Figure 3.3b. It is observed that if the regularisation parameter is chosen near the optimum, then location based weights do not effect the quality of the solution.

4.4 Weighted shearlet sparsity penalty

Main author: Jacob Frøsig.

Co-author: Nicolai André Brogaard Riis.

In this section, we consider a different version of the optimisation problem in (4.7) with a frame suited for objects with singularities on smooth curves such as the pipe by FORCE. By having a suited frame, structures that are not well represented by it, such as artefacts and noise, are removed by the sparsity penalisation. It has been shown in [11] that objects with singularities along smooth curves are well represented by the shearlets described in Subsection 4.1.3. Hence, we are motivated to use a weighted shearlet sparsity penalty.

The problem at hand is to fit our data, \mathbf{b}_Ω^δ , to Model 2.6:

$$\mathbf{b}_\Omega^\delta = A_\Omega \mathbf{x} + \mathbf{e},$$

by the variational formulation

$$\mathbf{c}^*|_{\alpha, w, \text{SH}} = \underset{\mathbf{c} \in \mathbb{R}^M}{\operatorname{argmin}} \left\{ \|A_\Omega \mathcal{T}_{\text{SH}}^* \mathbf{c} - \mathbf{b}_\Omega^\delta\|_2^2 + \alpha \sum_{\mu=1}^M w_\mu |c_\mu| \right\}.$$

Then, we get the reconstructed object as

$$\mathbf{x}^*|_{\alpha,w,\text{SH}} = \mathcal{T}_{\text{SH}}^* \mathbf{c}^*|_{\alpha,w,\text{SH}}.$$

For the analysis, \mathcal{T}_{SH} , and synthesis, $\mathcal{T}_{\text{SH}}^*$, operators, we use the discrete cone adaptive shearlet system described in Section 4.1.3. Here, the discrete shearlets, $\psi_{j,s,t;K} \in \mathbb{R}^n$, have compact spatial support with scaling, shearing and translations parameters j , s and t respectively specified for each cone set denoted by $K \in \{\Phi, \Psi, \tilde{\Psi}\}$. Hence, we imply the assumption that our object, \mathbf{x} , is finitely composed by shearlets, i.e., $\mathbf{x} \in \mathcal{S}_{J,S}$ where

$$\mathcal{S}_{J,S} = \left\{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} = \sum_{K \in \{\Phi, \Psi, \tilde{\Psi}\}} \sum_{j=1}^{J(K)} \sum_{s=1}^{S(j,K)} \sum_{t=1}^T \langle \mathbf{x}, \psi_{j,s,t;K} \rangle \psi_{j,s,t;K} \right\}.$$

For a given image size, the number of translations, T , is fixed, i.e., $T = n$ and the number of scales and shears for Φ is fixed to 1. Throughout the number of shearings, $S(j, K)$, depends on the scale, j , and cone set, K . This assumption gets more realistic as we increase the number of shearlets. The total number of shearlets is given by

$$M = n \sum_{K \in \{\Phi, \Psi, \tilde{\Psi}\}} \sum_{j=1}^{J(K)} S(j, K).$$




Figure 4.12: Illustration of different shearlets in spatial and frequency domain.

4.4.1 Numerical experiments

This subsection is dedicated to illustrate the impact of important parameter settings for the weighted shearlet sparsity penalty on the ROI problem given in Model 3.2. Throughout, we use the implementation in Shearlab documented in [13] for the analysis and synthesis operators. Here, a stable version of the cone adapted discrete shearlet system SH, described in Definition 4.11, is generated

with spatial compact support. This system does not form a tight frame, which is de-prioritised to achieve the spatial compactness.

The shearlets thus have a directional component and are compactly supported in spatial domain as illustrated in Figure 4.12 for shearlets in the first and third cones. Note that larger scales in Fourier domain correspond to smaller scales in spatial domain.

The system matrix, A_Ω , is generated using the parameters from Table 3.1 as described in Section 3.2 unless stated otherwise. The data is thus as shown in Figure 3.3b in Section 3.2.

To perform the optimisation we use the FISTA method described in Algorithm 1. We use the step-size $s_k = 1/B$, where B is the estimated squared two norm of AT_{SH}^* computed by using the power method.

Similar to the weighted wavelet sparsity penalty, the regularisation parameter is found using the ROI relative error, RE_Ω , mutual information, MI_Ω , as described in Section 3.2 and by visual inspection.

Scale dependent weights and shearings

In the discrete version of the shearlet system from Shearlab, one chooses the scale levels by specifying the number of scales, J . For each scale we specify a *shear level*, i.e., a parameter defining the number of shearings for a given scale. For a specific scale, j , and shear level, s , the generating shearlet is sheared 2^s times for both Ψ and $\tilde{\Psi}$. This gives us $S(s, \Psi) = S(s, \tilde{\Psi}) = 2 \cdot 2^s + 1$ number of shearings. The shear level is a trade-off between noise/artefact reduction on one hand and computational complexity on the other. The implementation stores coefficients for each translation in all pixels of the image, i.e. $T = n$. Hence, by increasing the number of shearings we increase the number of coefficients significantly. That is, going from s to $s + 1$ shearings for some scale, we get $2(2 \cdot 2^s)n$ additional coefficients.

To illustrate the impact of choosing the right scales Figure 4.13 shows reconstructions from noise free sinograms using shear levels denoted as a vector, $s_{\text{small}} = [3, 3]$ and $s_{\text{large}} = [3, 3, 3, 3, 3]$, i.e., s_{small} has 2 scales and a shear level of 3 for each scale and s_{large} has 5 scales and a shear levels of 3 for each scale. Here, we use scale dependent weights to motivate large structures well-knowing that the shearlets do not scale exactly by 2^j .

In the reconstructions using s_{small} , illustrated in Figure 4.13a, we see that high

frequent shearlet-structured artefacts occur outside the region-of-interest, where no true structure of the ground truth is present. Hence, we are motivated to use larger scales to avoid these artefacts. In the reconstruction using s_{large} in Figure 4.13b, these high frequency artefacts are replaced by larger shearlets allowing us to see the 3 small ellipses from the ground truth. To decrease the computational complexity, we can reduce the number of shearings for larger shearlets. This is illustrated in Figure 4.13c, maintaining the same image quality. Finally, to confirm that scale weights are necessary, we remove them and show the reconstruction in Figure 4.13d. The ROI artefact is then visible again. The ROI ring artefact also remains when varying the number of shearings and scales with no scale weights.

Regularisation parameter


The importance of the regularisation parameter is illustrated in Figure 4.14, where a too small, a proper and a too large α is used. Here, the small regularisation parameter, $\alpha = 10^{-4}$, results in noise corrupted reconstruction and $\alpha = 10$ results in missing details from the ground truth. We observe that by setting the regularisation parameter to $\alpha = 0.05$, we can retrieve more details yet avoid total corruption by noise and the ring artefact.

Location based weights

Figure 4.15 illustrates a reconstruction using the information based weighing scheme with $w_{\text{out}} = 10$. Here the method is shown both with and without scale dependent weights.


For the one using scale dependent weights, we observe that high frequent structures are avoided outside the ROI, but the structure inside the ROI is kept high frequent. The small ellipse structures of the ground truth is still present but blurred. For the one without scale dependent weights, we observe the left and right parts of the ring artefact are present. As expected, most attenuation is inside the ROI since it is not penalised as much.

For the ROI data, the scale dependent weights seem the most important since they remove the ring artefact. However, for other types of measurement geometries the location dependent weights might show improved results.



(a) $s_{\text{small}} = [3, 3]$, $\alpha = 1$,
with scale weights.
 $\text{RE}_{\Omega} = 0.30$, $\text{MI}_{\Omega} = 0.64$.

(b) $s_{\text{large}} = [3, 3, 3, 3, 3]$, $\alpha = 0.02$,
with scale weights.
 $\text{RE}_{\Omega} = 0.17$, $\text{MI}_{\Omega} = 0.64$.



(c) $SL = [1, 1, 1, 3, 3]$, $\alpha = 0.05$,
with scale weights.
 $\text{RE}_{\Omega} = 0.16$, $\text{MI}_{\Omega} = 0.63$.

(d) $SL = [1, 1, 1, 3, 3]$, $\alpha = 0.05$,
without scale weights.
 $\text{RE}_{\Omega} = 0.38$, $\text{MI}_{\Omega} = 0.65$.

Figure 4.13: Shearlet reconstructions on a noise free version of the ROI sino-gram in Figure 3.3b in Section 3.2 using different shear levels and scale weights. We see that having large shearlets with scale weights is necessary to avoid the ring artefact. In addition, larger shearlets need not be sheared as much.




Figure 4.14: Shearlet reconstruction on ROI-data in Figure 3.3b from Section 3.2 using shear levels $[1, 1, 1, 3, 3]$ for different regularisation parameters.




Figure 4.15: Reconstruction on ROI-data in Figure 3.3b from Section 3.2 using information location based weights and shear levels $[1, 1, 1, 3, 3]$ with and without scale dependent weights. We see there is no obvious improvement to be found by adding location based weights. In addition, location based weights cannot be used to replace scale dependent weights for removing the ring artefact.

4.5 Reflections

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

4.5.1 Choice of frame

The weighted frame-based sparsity penalty is able to incorporate prior information of the object by choosing the right frame and weights for a given measurement geometry. For the ROI problem, we see that the most influential weighting scheme is the scale dependent weights, which avoid the ring artefacts.

To illustrate the importance of choosing the right frame for a given object, consider Figure 4.16. In Figures 4.16a and 4.16b the method using shearlet regularisation performs significantly worse than the wavelet regularisation method, since the phantom is not well-represented by shearlets. When comparing Figures 4.16c and 4.16d, we see that the method using shearlets represents details inside the region of interest better compared to the wavelet method, in which the reconstruction gets pixelated. Note, this is not reflected in the relative error and mutual information measures.

4.5.2 Artefact removal

An important feature of the frame based methods is the removal of the ROI ring artefact. The reconstructions of a modified Shepp–Logan phantom with added singularities on the ROI-boundary from noise-free ROI-data is shown in Figure 4.17. Here we observe that the frame-based methods do not remove the singularities from the ground truth on the ROI boundary. This is an important feature, since we do not want to remove structures from the ground truth that are present in the data. This is shown to hold experimentally for different objects and reasonable choices of parameters. Hence, we are led to believe this a general property of the weighted frame-based methods.

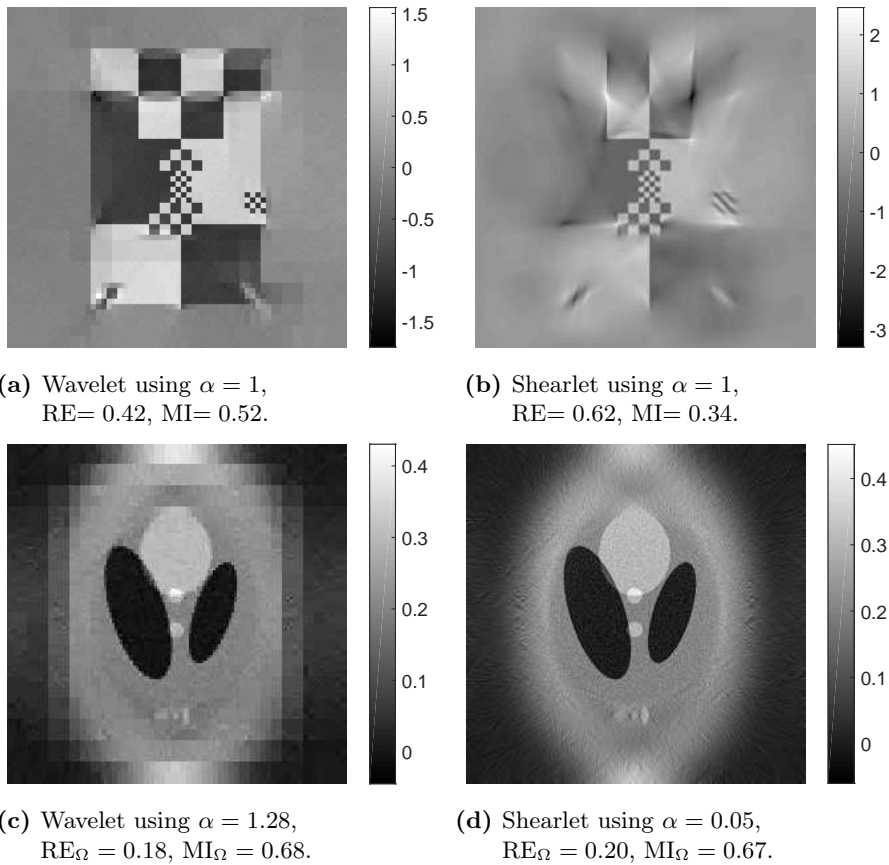





Figure 4.16: Comparing the best reconstructions using weighted wavelet and shearlet regularisation on the wavelet and Shepp-Logan phantoms. Shear levels are $[1, 1, 1, 3, 3]$ and the wavelets have 8 scales. Scale dependent weights are used in all reconstructions. This illustrates the importance of choosing the right frame for a given object.




(a) Shepp–Logan phantom with added singularities at the right and top boundary of the ROI.




(b) Filtered back projection using Hamming filter.



(c) Landweber using 100 iterations.



(d) Wavelet using $\alpha = 1$.



(e) Shearlet using $\alpha = 0.01$.

Figure 4.17: Reconstructions on Shepp–Logan phantom with added singularities. The frame-based methods with scale dependent weights avoid the added artefact at the ROI boundary, while still reconstructing the singularities from the ground truth on the boundary.

Computed tomography on deep sea oil pipes

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

In this chapter, we apply the insight and algorithms, discussed in previous chapters, to the prototype X-ray inspection device for deep sea oil pipes by FORCE. The aim is to reconstruct a 2D slice of the pipe from ROI measurement data. We firstly introduce the physical test set-up and describe how the data is collected. The data is then compared to simulations from a forward model on a synthetic phantom, to determine if the model is a realistic approximation to the measurement set-up. In the latter part, we finally test our reconstruction methods on the real data and discuss how to improve the results.

The real measurement data has been graciously provided by FORCE for use in this project.

5.1 Prototype set-up

The prototype set-up is meant to be a good approximation of how the data is acquired on the finalised device. Therefore, it is important to have a good physical model of the pipe and inspection device. In addition, the prototype should include any exterior factors that can impact the quality of the measurement data.




Figure 5.1: 3D rendering the prototype measurement set-up. The pipe is illuminated by X-rays from a betatron source and those X-rays are detected using a line detector. Note that the pipe is only partially illuminated since the beam is so narrow. This set-up is meant give an approximation to the actual measurement set-up illustrated in Figure 1.1. This illustration has been graciously provided by FORCE.

To accomplish this goal, the set-up uses the same source and detector as is intended for the finalised device. That is, a line detector with 8 modules totalling 512 detector pixels and a betatron source outputting X-rays with intensities up to 6MeV. The object is a section of an actual oil pipe that has been extracted from the seabed. This gives a good prototype of how the data are acquired in the final device. The test set-up is illustrated on a 3D rendering in Figure 5.1.

The source outputs X-rays in a cone-beam, which is equivalent to the fan-beam projections illustrated in Figure 2.3 for the line detector measurements. The measurements are taken by rotating the pipe and not the source and detector in the prototype set-up. To perform this rotation, a steel axis is inserted near the centre of the pipe. This axis will not be present in the actual oil pipes, but is necessary for taking reliable measurements around the pipe. In addition to the steel rotation axis, the pipe has 4 layers as illustrated on the cross-section of the setup in Figure 5.2. The individual pipes, from inner-most to outer-most, consist of the following materials: Steel, PE-foam, PU-rubber and reinforced concrete.




Figure 5.2: Illustration of a cross-section for the prototype set-up. The measurement specifications are in millimetres. The layers of the pipe, from inner-most to outer-most, consist of the following material: Steel, steel, PE-foam, PU-rubber and reinforced concrete.

The X-rays are filtered by wolfram before and after entering the pipe to reduce beam hardening. With this information in mind, we are able to give a rough estimate of the dampening coefficients for each type of material in the pipe. This will be important when comparing the forward model to the actual data, since a synthetic phantom of the pipe is then necessary. This is explained further in Section 5.3.

5.2 Gathering & preprocessing of data

The data is gathered by the line detector with 512 detector pixels sampling a photon count every 25 ms and storing an average of these samples for each angular degree. This gives us a dataset of size 512×360 .

The measured data contains dead pixels located at one end of the detector. Instead of adding prior to the data, we simply remove these pixels. However, because the pipe is rotating and not the detector and source, it is crucial to align the centre of rotation with the centre pixel in the detector, because of the assumed ROI model. To adhere to this, we also remove the same number of pixels in the other end of the detector. This reduced detector size is indicated by the red lines in Figure 5.2. The final dataset has size 491×360 .




Figure 5.3: Measurements of I_0 and I for a specific source position. The zero, I_0 , is measured by removing the pipe, and I is measured with the pipe.

The mathematical model of projections in (2.3) uses a proportional measure, $-\ln(I/I_0)$, rather than photon count, where I_0 and I are the intensities of the beam before and after it goes through the object. To approximate I and I_0 , we sample with and without the object respectively. These samples are shown for one source position in Figure 5.3. Note that the intensity of the X-ray beam is most powerful at its centre. In addition, we observe drops in intensity at transitions between detector modules.

By modelling from the projection data, $-\ln(I/I_0)$, we approximate the attenuation of the object alone, excluding other factors such as the drops in intensity between detector modules. The collection of these projections is then our discrete sinogram, \mathbf{b}_Ω^δ , as shown in Figure 5.5b. We note the attenuation is highest for detector pixels where the X-rays have gone long distances through steel, indicated by the four bright horizontal lines. In addition, we see vertical streaks that seem to "exit" the sinogram and hence must represent features that are outside the ROI and hence not measured from all source positions.

5.3 Comparing forward simulation with obtained data

We generate a forward model as described in Model 2.6 with the physical parameters of the prototype set-up illustrated in Figure 5.2. The parameters necessary for generating the forward model are shown in Table 5.1. Note that the domain size is chosen to include the entire pipe, although no other information of the pipe is included in the forward model.

Table 5.1: The physical parameters used for the forward model to simulate the prototype set-up. The detector length is adjusted to take dead pixels into account.

Grid size (N)	512
Number of source locations	360
Number of detector pixels	491
Domain size	55 [cm] × 55 [cm]
Source to centre distance	59 [cm]
Source to detector distance	100 [cm]
Detector length	41.1 · (491/512) [cm]

To verify that our forward model is a good approximation to the measurement set-up, we need a simulation of the real object to be transformed by our forward model. A synthetic pipe phantom of size 512 × 512 is shown in Figure 5.4 with ROI indicated in blue. The attenuation coefficients are chosen as follows: For specific materials a *mass* attenuation coefficient, κ , is used to describe the attenuation coefficient as

$$x [1/\text{cm}] = \kappa [\text{cm}^2/\text{g}] \cdot \rho [\text{g}/\text{cm}^3],$$

where ρ is the density of the material.

The mass attenuation coefficients depend on the intensity of the beam. The beam's intensity is assumed to be 2MeV after leaving the wolfram filter and the corresponding mass attenuation coefficients are acquired by the NIST database [21]. The simulated data is observable in Figure 5.5a.

We observe similar structure in the sinograms. Noting that, the ROI is just outside the steel pipe in Figure 3.2, we are to believe that the high attenuations in the upper and lower part of the sinograms arise from the out-most steel pipe. Similarly, the high attenuation closer to the vertical centre of the sinograms is believed to arise from the rotation axis. Note the small intensity attenuation difference between the simulated and the real sinogram. This can be explained

by some errors in the assumption on the intensity coefficients of the object and measurement errors. As a conclusion, Model 3.2 seems to mimic the overarching characteristics of the prototype set-up well and we are motivated to use the forward model in our reconstruction methods.




Figure 5.4: Synthetic pipe phantom with 512×512 pixels. The attenuation coefficients are chosen to fit a 2MeV intensity beam going through the pipe described in Figure 5.2.




Figure 5.5: Comparing the real ROI measurement data with the simulated data generated from the forward model described by Table 5.1 on the object in Figure 5.4 on a fine grid. The simulated data has added 2% relative noise.


5.4 Reconstructions

We now compare the filtered back projection, Landweber, weighted wavelet and shearlet reconstructions on the real measurement data. The range of iteration numbers and regularisation parameters are narrowed down by reconstructing on the simulated data, for which the ground truth is known. Then, to find the optimal regularisation parameters on the real measurement data, the reconstructed images are inspected visually. The filter for filtered back projection is chosen as the Hamming filter. The wavelet and shearlet methods both use scale dependent weights justified by the analysis in Sections 4.3.1 and 4.4.1. The shear levels are chosen as $[1, 1, 1, 3, 3]$ and the wavelets have 8 scales.

The results are shown in Figure 5.6. The ring artefact is very clear on the filtered back projection method and completely dominates the solution. In the Landweber and frame-based methods, the intensity values of the pipe inside the ROI fit the assumed attenuation coefficients quite well. It is clear, however, that the details outside are very obscured. It is hard to tell which of the methods provide the best reconstructions. Some of the defects in the concrete are more clearly distinguished in the wavelet and shearlet reconstructions compared to the Landweber method.


In general, the pipe problem from FORCE is shown not to be well suited for ROI measurements. This is justified by our synthetic and real experiments. The ROI boundary lies close to the second steel pipe, which means we cannot trust the reconstruction of this boundary in the Landweber method. In addition, the outer layers of the pipe are completely lost in the reconstructions. This can be explained by microlocal analysis from Corollary 2.24 in Section 3.3.2, since no rays pass through orthogonal to the direction of the singularities. Hence, we cannot expect to reconstruct these singularities, since they are not visible in the data.

This leads us to propose another measurement geometry that captures these singularities.



(a) Filtered back projection with Hamming filter.

(b) Landweber after 80 iterations.



(c) Wavelet regularisation method with $\alpha = 0.21$ and scale dependent weights.

(d) Shearlet regularisation method with $\alpha = 0.03$ and scale dependent weights.

Figure 5.6: Reconstructions from the real data in Figure 5.5b. The shear levels are $[1, 1, 1, 3, 3]$ and the wavelets have 8 scales. All the methods are unable to reconstruct most features outside the ROI. The filtered back projection method shows a significant artefact at the ROI.

5.5 Exterior tomography

In exterior tomography the data is measured outside some central region, keeping the centre of rotation the same as in ROI tomography. This is in contrast to ROI, where the data is measured inside this central region. It can be understood by shifting the detector and source some distance orthogonal to the line going through the source position and centre of rotation as illustrated in Figure 5.7. The goal here is to capture those singularities that are invisible in the ROI data, namely the boundaries of the pipes lying outside the ROI, along with any defects that lie along these boundaries.

For the prototype set-up, this is done by lifting the pipe relative to the source and detector as illustrated in Figure 5.7. Since the pipe is rotating, this is equivalent to shifting the detector and source in the final device.

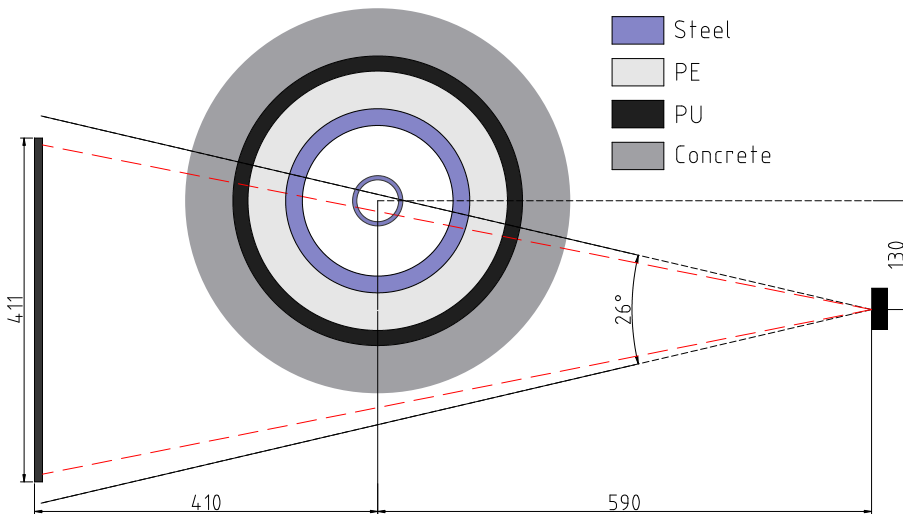


Figure 5.7: Illustration of a cross-section of the exterior tomography measurement set-up. The measurements specifications are in millimetres. Note the centre of rotation is still at the centre of the pipe, but the relative position of the detector and source is shifted.

We thus have the following model for exterior tomography:

MODEL 5.1 (DISCRETE EXTERIOR TOMOGRAPHY WITH NOISE)

We model the attenuation of X-rays measured by the scheme illustrated in Figure 5.7 as follows:

$$\mathbf{b}_E^\delta = A_E \mathbf{x} + \mathbf{e}, \quad \text{for } \mathbf{b}_E^\delta \in \mathbb{R}^m, \mathbf{x} \in \mathbb{R}^n \text{ and } A_E \in \mathbb{R}^{m \times n}.$$

Here $(b_E^\delta)_i$ is a measured exterior projection from (2.4) with added Gaussian distributed white noise, $\mathbf{e} \in \mathbb{R}^m$ and A_E contains rows corresponding to the measured X-rays in exterior tomography.

We modify the linear detector version of `fanbeamtomo.m` to allow shifting the relative position of the source and detector. This is done relatively easily in the matrix version of the implementation. The system matrix is hence generated using the physical parameters for the set-up as shown in Table 5.2.

Table 5.2: The physical parameters used for the forward model to simulate the exterior tomography prototype set-up. Here the detector length is adjusted for dead pixels.

Grid size (N)	512
Number of source locations	360
Number of detector pixels	507
Domain size	55 [cm] × 55 [cm]
Source to centre distance	59 [cm]
Source to detector distance	100 [cm]
Detector length	$41.1 \cdot (507/512)$ [cm]
Source & detector shift	-13 [cm]

The measurements and simulations are then carried out using exactly the same parameters as the ROI set-up, except for the added shift in source and detector. Note that the number of detector pixels is 507, since it is no longer necessary to remove pixels to centre the pipe in the data. The resulting sinograms for both the simulated data and real measurement data are shown in Figure 5.8. The model fits quite well to the real measurement data. Note how the outer layers of the pipe, seen as the top-most boundaries in the sinogram, are uneven in the real measurement data.




Figure 5.8: Comparing the real exterior measurement data with the simulated data generated from the forward model described by Table 5.2. The simulated data has no added noise and is generated from the 512×512 object in Figure 5.4.





Figure 5.9: Exterior mask for information based weights for the measurement set-up in Figure 5.7.


The reconstructions using filtered back projection and Landweber methods are shown in Figure 5.10. The regularisation parameter is chosen by visual inspection of a range reconstructions. The methods show ring artefacts where the sinogram is cut-off, just as in ROI tomography. Note for exterior tomography, we have a cut-off in two places. The outer ring artefact is more clear for the Landweber method in Figure 5.10b, whereas the inner ring artefact is more clear in the filtered back projection method in Figure 5.10a, although both are present in the reconstructions. We expect the frame-based methods with scale weights can avoid these artefacts, if they are less significant in the data fitting.

The reconstructions for the weighted wavelet and shearlet methods using scale weights are shown in Figures 5.11a and 5.12a, respectively. Note the artefacts both inside the inner steel pipe and outside the entire pipe. The wavelet method does not reconstruct the artefact outside the pipe, since it is not well represented by the frame. However, the artefact inside the pipe is well-represented, and so it remains in the reconstruction. We expect the information based weights can help remove this artefact.

Finally, we try the information based weights using the generated mask shown in Figure 5.9. The exterior measurement geometry may benefit from this type of weighting scheme. The reconstructions are shown in Figures 5.11b and 5.12b. For the wavelet method the added location based weights show a slight improvement by visual inspection. The shearlet method show no obvious improvement by the added informations weights.



(a) Filtered back projection using Hamming filter.



(b) Landweber after 450 iterations.

Figure 5.10: Reconstructions using standard methods from the real data in Figure 5.8b. The exterior measurements provide better data for reconstruction of singularities, however, the standard methods still show added ring artefacts.



(a) $\alpha = 0.1$.(b) $\alpha = 0.1, w_{out} = 5$.

Figure 5.11: Wavelet reconstructions with scale dependent weights from the real data in Figure 5.8b. Compared to the standard methods the the ring artefact outside the pipe is less pronounced. The centre ring artefact is represented by a few wavelets and thus remains. Finally, adding information based weights show a slight improvement in the image quality.



(a) $\alpha = 0.01$.(b) $\alpha = 0.01$, $w_{out} = 10$.

Figure 5.12: Shearlet reconstructions with scale dependent weights from the real data in Figure 5.8b. Similar to the wavelet method the ring artefact is less pronounced even though it is quite well-represented by the frame. In addition, new shearlet structured artefacts show outside the pipe. Finally, adding location based weights show no obvious improvement in image quality.

5.6 Reflections

We have shown, that the problem proposed by FORCE is better suited for exterior tomography, since more singularities are captured in the data. Even though the reconstructions on real data using exterior tomography and frame-based methods yield good results, we now elude to the fact that some singularities are still not visible for this particular measurement set-up.

5.6.1 Combining measurement data

Consider the synthetic pipe with added singularities shown in Figure 5.13. From the reflections in Chapter 3, we note that some of these singularities will be invisible using ROI measurements, in addition to the outer layers of the pipe. This is because there are no rays that pass orthogonal to the singularities direction. Similarly for exterior measurements, some of these singularities, particularly the ones that "point" towards the centre of rotation, are invisible in the measurements. We note that the set of singularities that are visible by combining ROI and exterior data are all possible singularities.

Hence, we are motivated to combine these sets of data. This is done numerically, by creating a combined system matrix and sinogram as follows

$$A = \begin{bmatrix} A_{\Omega} \\ A_E \end{bmatrix}, \mathbf{b} = \begin{bmatrix} \mathbf{b}_{\Omega} \\ \mathbf{b}_E \end{bmatrix}.$$

We are then able to use the frame-based reconstruction methods to solve the combined problem:

$$A\mathbf{x} = \mathbf{b}.$$

To illustrate the benefit of combining the data, we showcase reconstructions where the data is generated *with inverse crime* and has no noise. That is, we use the same forward model for generating the data and reconstruction. In Figure 5.15 we see the results. Note, how some singularities are not reconstructed using the ROI data in Figure 5.14a and exterior data in Figure 5.14b. However, when combining the data, all singularities are visible as shown in Figure 5.15c. We compare with the Landweber method in Figure 5.15d, which still reconstructs the outer exterior ring artefact.





Figure 5.13: Pipe phantom with added singularities. It is sufficient to consider horizontal singularities since the pipe is rotationally symmetric.



(a) ROI sinogram from the system matrix generated from Table 3.1. (b) Exterior sinogram from the system matrix generated from Table 5.2.

Figure 5.14: The sinogram generated from the forward models using the phantom shown in Figure 5.13. Notice how some singularities are visible in the data, while others are not for different measurement geometries.

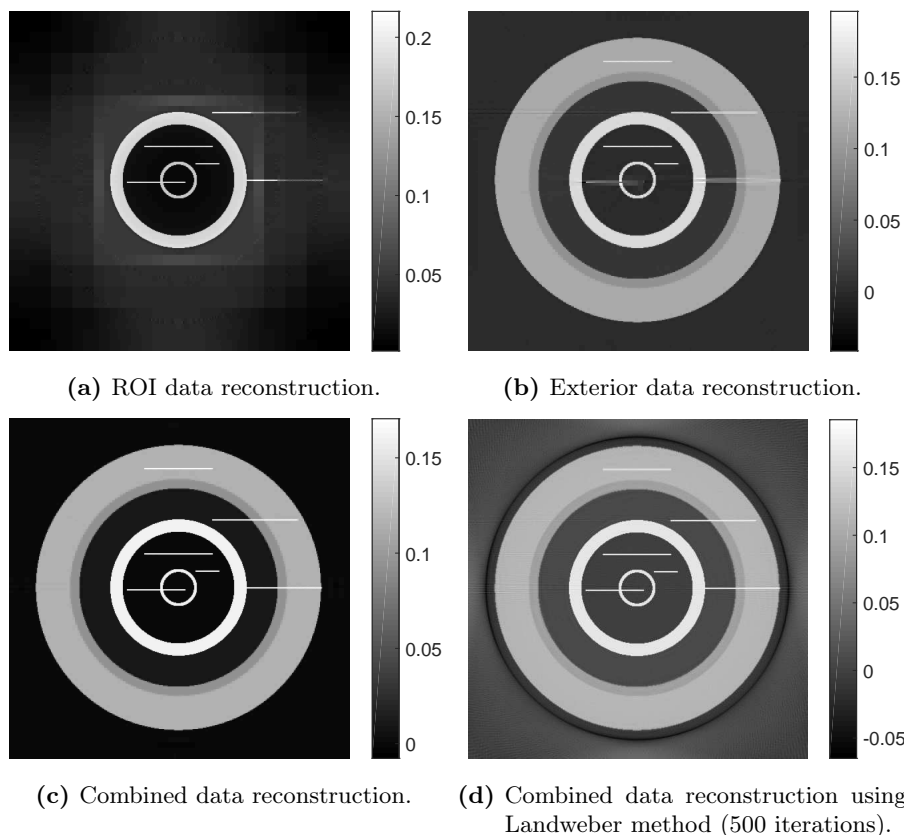


Figure 5.15: Comparing the reconstructions for ROI, exterior and combined data sets shown from Figure 5.14. The system matrices for ROI and exterior measurement geometries are generated from Tables 3.1 and 5.2, respectively. The weighted frame-based method is wavelet regularisation with $\alpha = 0.2$ and scale dependent weights. By combining the data from both exterior and ROI measurements we are able to capture and reconstruct all singularities. In addition, the frame-based method does not generate additional artefacts in the reconstruction.

5.6.2 Increased detector size for exterior measurements

Another method to capture all singularities is to increase the detector size, such that the exterior measurements has rays going through the centre of rotation. This increase in detector size is not particularly large as illustrated in Figure 5.16. Hence, we increase the detector size to 46.2 [cm].

The generated sinogram is shown in Figure 5.17a. Notice the two new singular points, which are not present in the sinogram generated from a smaller detector shown in Figure 5.14b. These singularities in the sinogram correspond to the two singularities from the ground truth that point towards to centre of rotation. We reconstruct using the weighted wavelet sparsity penalty and Landweber method as shown in Figures 5.17b and 5.17c, respectively. Indeed, all singularities from the ground truth are reconstructed in the image. The wavelet method avoids adding addition singularities, whereas the Landweber method still creates the exterior ring artefact, in addition to other small artefacts.






Figure 5.16: Illustration of a cross-section of the exterior tomography prototype set-up with a larger detector. The measurements specifications are in millimetres. Note the centre of rotation is still at the centre of the pipe, but relative position of the detector and source has been shifted. With the increased detector size we expect to capture all singularities of the object.



(a) Sinogram from simulated exterior measurements with a larger detector size.



(b) Weighted wavelet reconstruction. $\alpha = 0.2$. (c) Landweber reconstruction with 500 iterations.

Figure 5.17: Comparing the reconstructions for the extended detector size with exterior measurement geometry. As expected all singularities are visible in the data and hence can be reconstructed by the methods. The Landweber method still creates additional artefact in the pipe as evident from the negative valued ring outside the pipe.

CHAPTER 6

Conclusion & future work

Main authors: Nicolai André Brogaard Riis & Jacob Frøsig.

We now give a final conclusion for the project in a broader context, combining the insights gained throughout the chapters.

We start by evaluating if the mathematical model of X-ray tomography is a good approximation to reality. The models are shown to maintain the key properties of how X-rays propagate through objects, even through the hierarchy of approximations. This claim is mainly supported by how well the synthetic data fit the real measurements. This is cemented by the predictions made in microlocal analysis of the singularities visibility in the continuous model.

We have shown the limitations of region-of-interest tomography, in the sense that some singularities outside the region-of-interest are not visible in the data and hence can not be reconstructed. In addition, artefacts are shown in standard reconstruction methods on ROI measurements at the boundary between fully and partially illuminated regions.

A variational formulation with a weighted frame-based sparsity penalty is shown to provide good reconstructions, for objects that are sparsely represented by the frame. In addition, the added singularities seen for standard methods are shown to be less prominent in reconstructions by weighted frame-based methods. This is attributed to the fact that, the singularities are not as significant in the data fitting and are hence removed by the weighted sparsity penalty.

Finally, the ROI measurement geometry is shown to be an ineffective method of

capturing singularities of oil pipes as in the problem proposed by FORCE. An exterior tomography measurement geometry is proposed and is shown to better capture the singularities. Finally, by either combining the ROI and exterior data, or increasing the detector size for the exterior measurement set-up, it is experimentally shown that all singularities are captured, given enough data. The weighted frame-based method is shown to give the best reconstructions for the exterior measurement geometry compared to the considered standard methods.

Future work

Several interesting studies presented themselves while working on this thesis, some of which could not be included, because of time limitations, including the following.

The included continuous model for ROI-tomography uses an altered version of the Radon transform as the forward operator. It would be interesting to study this transform in more detail, with the aim of revealing general properties of, say, its adjoint.

Two different frames, i.e., shearlets and wavelets was studied in the thesis. It would be interesting to elaborate this study to other frames suited for different kinds of objects. This could emerge in a catalogue describing which frame system to use for specific types of objects.

A study of the stability in the reconstruction quality when varying the model, regularisation and frame parameters could be beneficial. This could be used to avoid parameter tuning on new problems and quantify the uncertainty of the reconstructed images. To this end, an interesting topic could be to apply *uncertainty quantification* in computed tomography.

As discussed in Section 3.3, we observed a ring artefact when using standard iterative methods for the ROI tomography. A thorough examination of why these artefacts occur might present useful insight into how they can be avoided.

One of the issues of using the frame-based methods, is the computational complexity of the frame decomposition. Hence, a faster implementation is highly beneficial to obtain fast image reconstructions, thus decreasing time spent tuning parameters. This could be done by implementing a fast GPU version of the synthesis and analysis operators, or by parallelising the FISTA algorithm.

A.1 Alternating Direction Method of Multipliers (ADMM)

Another method of solving the optimization problem is the Alternating Direction Method of Multipliers (ADMM). This is derived by rewriting the variational formulation in (4.7) into a constrained optimization problem:

$$\begin{aligned} \mathbf{c}^*_{\alpha,w} &= \underset{\mathbf{c} \in \ell_2}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathcal{K}\mathbf{c} - \mathbf{b}^\delta\|^2 + \alpha \|\mathbf{d}\|_1 \right\}, \\ &\text{s.t. } W\mathbf{c} = \mathbf{d}. \end{aligned}$$

The augmented Lagrangian for the above problem is

$$L_\rho(\mathbf{c}, \mathbf{d}, \lambda) = \frac{1}{2} \|\mathcal{K}\mathbf{c} - \mathbf{b}^\delta\|_2^2 + \alpha \|\mathbf{d}\|_1 + \lambda^T (\mathbf{d} - W\mathbf{c}) + \frac{\rho}{2} \|\mathbf{d} - W\mathbf{c}\|_2^2.$$

The ADMM iterations are then given by (see e.g. [22])

$$\begin{aligned} \mathbf{c}^{[k+1]} &= \underset{\mathbf{c} \in \mathbb{R}^K}{\operatorname{argmin}} L_\rho(\mathbf{c}, \mathbf{d}^{[k]}, \lambda^{[k]}), \\ \mathbf{d}^{[k+1]} &= \underset{\mathbf{d} \in \mathbb{R}^K}{\operatorname{argmin}} L_\rho(\mathbf{c}^{[k+1]}, \mathbf{d}, \lambda^{[k]}), \\ \lambda^{[k+1]} &= \lambda^{[k]} + \rho(\mathbf{d}^{[k+1]} - W\mathbf{c}^{[k+1]}). \end{aligned}$$

Taking derivatives equal to 0 in c and d for the ADMM iterations we find the following update steps:

$$\begin{aligned}\mathbf{c}^{[k+1]} &= (\mathcal{K}^* \mathcal{K} + \rho W^2)^{-1} (\mathcal{K}^* \mathbf{b}^\delta + \rho W \mathbf{d}^{[k]} - W \lambda^{[k]}) \\ \mathbf{d}^{[k+1]} &= S_{\alpha/\rho}(W \mathbf{c}^{[k+1]} + \lambda^{[k]}/\rho) \\ \lambda^{[k+1]} &= \lambda^{[k]} + \rho(\mathbf{d}^{[k+1]} - W \mathbf{c}^{[k+1]})\end{aligned}$$

where $S_\kappa(x) = \text{sgn}(x) \max(|x| - \kappa, 0)$. The step-size parameter, ρ , can be updated using the rule from [22] as follows:

$$\rho^{k+1} := \begin{cases} \tau^{\text{incr}} \rho^k & \text{if } \|\mathbf{r}^k\|_2 > \mu \|\mathbf{s}^k\|_2 \\ \rho^k / \tau^{\text{decr}} & \text{if } \|\mathbf{s}^k\|_2 > \mu \|\mathbf{r}^k\|_2 \\ \rho^k & \text{otherwise,} \end{cases} \quad (\text{A.1})$$

where $\mu > 1$, $\tau^{\text{incr}} > 1$ and $\tau^{\text{decr}} > 1$ (typically $\mu = 10$ and $\tau^{\text{decr}} = \tau^{\text{incr}} = 2$). Here $\mathbf{r}^k = \mathbf{d}^k - W \mathbf{c}^k$ is the primal residual and $\mathbf{s}^k = \rho W(\mathbf{d}^k - \mathbf{d}^{k-1})$, the dual residual. This leads us to the ADMM algorithm shown in Algorithm 2.

Note that the first update in the while loop, for the ADMM algorithm, requires a solve on a linear system of equations. In the case of frame-based variational regularisation, it requires an iterative solver, such as a conjugate gradient method.

Algorithm 2: Alternating Direction Method of Multipliers (ADMM).

Input: Noisy data \mathbf{b}_Ω^δ , ROI Radon Transform A_Ω , A_Ω^T and analysis, \mathcal{T}_Φ , and synthesis, \mathcal{T}_Φ^* , operators. $\mathcal{K} = A_\Omega \mathcal{T}_\Phi^*$, $\mathcal{K}^* = \mathcal{T}_\Phi A_\Omega^T$.

Output: Reconstruction, \mathbf{x} , that is the approximate solution to (4.2).

Assume $\mathbf{x} \in F_M$ for some $M \in \mathbb{N}$.

Choose some initial $\rho > 0$ and some weights $w = \{w_\mu\}_{\mu=1}^M$.

Choose regularisation parameter α .

Initiate Start guess $\mathbf{c}^{[0]}, \mathbf{d}^{[0]}, \lambda^{[0]}$.

while $k < k_{max}$ **do**

$$\left| \begin{aligned}\mathbf{c}^{[k+1]} &= (\mathcal{K}^* \mathcal{K} + \rho W^2)^{-1} (\mathcal{K}^* \mathbf{b}_\Omega^\delta + \rho W \mathbf{d}^{[k]} - W \lambda^{[k]}) \\ \mathbf{d}^{[k+1]} &= S_{\alpha/\rho}(W \mathbf{c}^{[k+1]} + \lambda^{[k]}/\rho) \\ \lambda^{[k+1]} &= \lambda^{[k]} + \rho(\mathbf{d}^{[k+1]} - W \mathbf{c}^{[k+1]}) \\ \text{Update } \rho &\text{ by (A.1)} \\ k &= k + 1\end{aligned}\right.$$

end

return $\mathbf{x} = \mathcal{T}_\Phi^* \mathbf{c}^{[k_{max}]}$

A.2 Additional theory

DEFINITION A.1 The Schwartz space $\mathcal{S}(\mathbb{R}^n)$ is defined as the set of all functions $f \in C^\infty(\mathbb{R}^n)$ for which

$$\|f\|_{\alpha,\beta} := \sup_{x \in \mathbb{R}^n} |x^\beta D^\alpha f(x)| < \infty$$

for all multi-indices $\alpha, \beta \in \mathbb{N}_0^n$, where $D^\alpha = \partial^\alpha / \partial x^\alpha$.

DEFINITION A.2 (FOURIER TRANSFORM) The Fourier transform of a function $f \in L^1(\mathbb{R}^n)$ is defined by the operator $\mathcal{F} : L^1(\mathbb{R}^n) \rightarrow C_0(\mathbb{R}^n)$ of a function

$$\mathcal{F}f(\xi) := \hat{f}(\xi) = (2\pi)^{-n/2} \int_{\mathbb{R}^n} f(x) e^{-i\langle x, \xi \rangle} dx.$$

That the Fourier transform maps into all continuous functions vanishing at infinity is shown by the Riemann-Lebesgue Lemma. Hence it does not necessary hold that $\hat{f} \in L^1(\mathbb{R}^n)$. We can now define the Inverse Fourier transform

DEFINITION A.3 (INVERSE FOURIER TRANSFORM) The inverse Fourier transform of a function $g \in L^1(\mathbb{R}^n)$ is defined by

$$\mathcal{F}^{-1}g(x) := (2\pi)^{-n/2} \int_{\mathbb{R}^n} g(\xi) e^{i\langle \xi, x \rangle} d\xi.$$

Hence, if $\hat{f} \in L^1(\mathbb{R}^n)$ then the equality $\mathcal{F}^{-1}\hat{f}(x) = f(x)$ holds. In particular it can be shown that if $f \in \mathcal{S}(\mathbb{R}^n)$ then $\hat{f} \in \mathcal{S}(\mathbb{R}^n) \subset L^1(\mathbb{R}^n)$.

One of the most central results in tomography is the Fourier Slice Theorem. It relates the Radon transform of a function to the function itself by Fourier transforms. In the following we use the 1D Fourier transform of $\mathcal{R}f(\theta, s)$ along s and denote it by $\widehat{\mathcal{R}f}(\theta, \sigma) = \mathcal{F}_s \mathcal{R}f(\theta, \sigma)$. Whenever an angle, θ , is fixed we denote the Radon transform for a fixed angle by $\mathcal{R}_\theta f(s)$.

THEOREM A.4 (FOURIER SLICE THEOREM) For $f \in \mathcal{S}(\mathbb{R}^n)$ and $\theta \in S^{n-1}$ we have

$$\widehat{\mathcal{R}_\theta f}(\sigma) = (2\pi)^{(n-1)/2} \hat{f}(\sigma\theta), \quad \sigma \in \mathbb{R}.$$

PROOF. Applying the definition of the one-dimensional Fourier transform of

$\mathcal{R}_\theta f$ over s and then the Radon transform we get

$$\begin{aligned}\widehat{\mathcal{R}_\theta f}(\sigma) &= (2\pi)^{-1/2} \int_{\mathbb{R}} \mathcal{R}_\theta f(s) e^{-is\sigma} ds \\ &= (2\pi)^{-1/2} \int_{\mathbb{R}} \int_{\theta^\perp} f(\theta s + y) dy e^{-is\sigma} ds \\ &= (2\pi)^{-1/2} \int_{\mathbb{R}} \int_{\theta^\perp} e^{-is\sigma} f(\theta s + y) dy ds.\end{aligned}$$

Now, let $\varphi : \theta^\perp \times \mathbb{R} \rightarrow \mathbb{R}^n$, be defined by $\varphi(y, s) = s\theta + y$, and let $u = (y, s)$. Note that $s = x \cdot \theta = \theta \cdot x$ over θ^\perp and that $\theta^\perp \times \mathbb{R}$ is an open set. Then by change of variables we see that

$$\begin{aligned}\widehat{\mathcal{R}_\theta f}(\sigma) &= (2\pi)^{-1/2} \int_{\mathbb{R}} \int_{\theta^\perp} e^{-is\sigma} f(\varphi(u)) du \\ &= (2\pi)^{-1/2} \frac{1}{|\det D\varphi(u)|} \int_{\mathbb{R}^n} e^{-i\sigma\theta \cdot x} f(x) dx.\end{aligned}$$

Noting that $|\det D\varphi(u)| = 1$, we get the desired result

$$\begin{aligned}\widehat{\mathcal{R}_\theta f}(\sigma) &= (2\pi)^{-1/2} (2\pi)^{n/2} \left((2\pi)^{-n/2} \int_{\mathbb{R}^n} e^{-i\sigma\theta \cdot x} f(x) dx \right) \\ &= (2\pi)^{(n-1)/2} \hat{f}(\sigma\theta).\end{aligned}$$

□

Using the Fourier Slice Theorem, one can derive an inversion formula for the Radon transform.

THEOREM A.5 (INVERSE RADON TRANSFORM) *For $f \in \mathcal{S}(\mathbb{R}^n)$ and the Radon transform $\mathcal{R}f(\theta, s)$, $\theta \in S^{n-1}$, $s \in \mathbb{R}$ we have that f is given by*

$$f(x) = \frac{1}{2} (2\pi)^{-n+1/2} \int_{S^{n-1}} \int_{-\infty}^{\infty} \mathcal{F}_s \mathcal{R}f(\theta, \sigma) e^{i\sigma\langle x, \theta \rangle} |\sigma|^{n-1} d\sigma d\theta.$$

PROOF. Assume $f \in \mathcal{S}(\mathbb{R}^n)$. Then by applying the inverse Fourier transform on $\hat{f} \in \mathcal{S}(\mathbb{R}^n)$, we can write

$$f(x) = (\mathcal{F}^{-1} \hat{f})(x) = (2\pi)^{-n/2} \int_{\mathbb{R}^n} \hat{f}(\xi) e^{i\langle x, \xi \rangle} d\xi. \quad (\text{A.2})$$

Let $U = S^{n-1} \times (0, \infty)$ and $V = \mathbb{R}^n \setminus \{0\}$. Note, U and V are open sets and since the singleton $\{0\}$ has measure zero in \mathbb{R}^n the integral over U and \mathbb{R}^n are equal. Furthermore, let

$$\varphi : U \rightarrow V, \quad (\sigma, \theta) \mapsto \sigma\theta = \xi$$

for $\theta \in S^{n-1}$ and $\sigma \in (0, \infty)$. Then we can write any $\xi \in V$ as $\xi = \varphi(\sigma, \theta) = \sigma\theta$. Noting that $|\det[(D\varphi)(\sigma, \theta)]| = |\sigma|^{n-1}$ we have by change of variables on (A.2) the following.

$$f(x) = (2\pi)^{-n/2} \int_{S^{n-1}} \int_0^\infty \hat{f}(\sigma\theta) e^{i\langle x, \sigma\theta \rangle} |\sigma|^{n-1} d\sigma d\theta.$$

The Fourier Slice Theorem A.4 and $\mathcal{R}f$ being even then yields

$$\begin{aligned} f(x) &= 2\pi^{-n+1/2} \int_{S^{n-1}} \int_0^\infty \mathcal{F}_s \mathcal{R}f(\theta, \sigma) e^{i\langle x, \sigma\theta \rangle} |\sigma|^{n-1} d\sigma d\theta \\ &= \frac{1}{2} (2\pi)^{-n+1/2} \int_{S^{n-1}} \int_{-\infty}^\infty \mathcal{F}_s \mathcal{R}f(\theta, \sigma) e^{i\langle x, \sigma\theta \rangle} |\sigma|^{n-1} d\sigma d\theta \end{aligned}$$

□

Bibliography

- [1] Frank Natterer and Frank Wubbeling. *Mathematical Methods in Image Reconstruction*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001.
- [2] Per Christian Hansen and Maria Saxild-Hansen. AIR Tools - A MATLAB package of algebraic iterative reconstruction methods. *Journal of Computational and Applied Mathematics*, 236(8):2167–2178, 2012.
- [3] ASTRA Toolbox. <http://www.astra-toolbox.com>. Accessed: 2017-2-14.
- [4] J. Hadamard. *Lectures on Cauchy's problem in linear partial differential equations*. Dover Publications, 1952.
- [5] Per Christian Hansen. *Discrete Inverse Problems: Insight and Algorithms*. Society for Industrial and Applied Mathematics SIAM, 2010.
- [6] Venkateswaran P. Krishnan and Eric Todd Quinto. Microlocal analysis in tomography. In Otmar Scherzer, editor, *Handbook of Mathematical Methods in Imaging*, pages 847–902. Springer-Verlag New York, 2 edition, 2015.
- [7] Lars Hörmander. *The Analysis of Linear Partial Differential Operators I: Distribution Theory and Fourier Analysis*. Grundlehren der mathematischen Wissenschaften 256. Springer Berlin Heidelberg, 2 edition, 1983.
- [8] Gitta Kutyniok and Demetrio Labate. *Shearlets: multiscale analysis for multivariate data*. Applied and numerical harmonic analysis. Birkhäuser, 2012.
- [9] Ole Christensen. *Functions, Spaces, and Expansions*. Birkhäuser Verlag GmbH, 2010.

- [10] E. W. C. and Ingrid Daubechies. Ten lectures on wavelets. *Mathematics of Computation*, 61(204):941, 1993.
- [11] KH Guo, D Labate, WQ Lim, G Weiss, and E Wilson. Wavelets with composite dilations and their mra properties. *Applied and Computational Harmonic Analysis*, 20(2):202–236, 2006.
- [12] Shearlab. <https://www.shearlab.org>. Accessed: 2017-2-16.
- [13] Gitta Kutyniok, Wang-Q Lim, and Rafael Reisenhofer. Shearlab 3d: Faithful digital shearlet transforms based on compactly supported shearlets. *Acm Transactions on Mathematical Software*, 42(1):5–47, 2016.
- [14] I Daubechies, M Defrise, and C De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.
- [15] Esther Klann, Eric Todd Quinto, and Ronny Ramlau. Wavelet methods for a weighted sparsity penalty for region of interest tomography. *Inverse Problems*, 31(2):025001, 2015.
- [16] Hengyong Yu and Ge Wang. Compressed sensing based interior tomography. *Physics in Medicine and Biology*, 54(9):2791–2805, 2009.
- [17] Jürgen Friel. *Reconstructions in limited angle x-ray tomography: Characterization of classical reconstructions and adapted curvelet sparse regularization*. Technische Universität München, PhD thesis, 2013.
- [18] Jennifer L. Mueller and Samuli Siltanen. *Linear and nonlinear inverse problems with practical applications*. SIAM,, 2012.
- [19] D. A. Lorenz. On the role of sparsity in inverse problems. *Journal of Inverse and Ill-posed Problems*, 17(1):61–68, 2009.
- [20] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *Siam Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [21] NIST X-ray attenuation coefficient tables. <https://www.nist.gov/pml/x-ray-mass-attenuation-coefficients>. Accessed: 2017-2-12.
- [22] Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press,, 2004.