



Tomographic Imaging of Flow Processes

Rasmussen, Peter Winkel

Publication date:
2022

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Rasmussen, P. W. (2022). *Tomographic Imaging of Flow Processes*. Technical University of Denmark.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

PhD Thesis
Doctor of Philosophy

 **DTU Compute**
Department of Applied Mathematics and Computer Science

Tomographic Imaging of Flow Processes

Peter Winkel Rasmussen

Kongens Lyngby 2022



DTU Compute
Department of Applied Mathematics and Computer Science
Technical University of Denmark

Richard Petersens Plads
Building 324
DK-2800 Kongens Lyngby, Denmark
Phone +45 4525 3031
compute@compute.dtu.dk
www.compute.dtu.dk

*“I was a victim of a series of accidents,
as are we all.”*

– Kurt Vonnegut, *The Sirens of Titan*

Summary

X-ray computed tomography (CT) is a powerful metrology that allows for non-destructive imaging of the internal structure of objects. Improvements in X-ray sources, X-ray detectors and increasing computation power have paved the way for dynamic tomography, which enables a new range of experiments.

Dynamic experiments are ubiquitous in research fields such as carbon storage, geothermal energy, multiphase flow, rock deformation and research related to the oil and gas industry. The dynamic nature of the experiments constrains the imaging process in that the imaging has to be sufficiently fast to capture it. This constraint can drastically compromise the image quality of tomographic reconstructions if either exposure time or the number of projections is reduced to accommodate the pace of the dynamic experiment.

In this thesis, we have developed novel reconstruction algorithms for dynamic experiments to provide high-quality reconstructions despite substandard data. We designed the algorithms for experiments that allow for a high-quality scan of the experiment before any dynamics take place. All algorithms we developed are based on the Simultaneous Iterative Reconstruction Technique, and the best performing of these algorithms had two additions, which improved it. The first addition is we initialise a reconstruction for a given time step with a reconstruction of the previous time step. The first time step is initialised with a high-quality reconstruction of the static system. The second addition is that each pixel in the dynamic reconstructions is constrained such they can only take a single value or a range of values. This constraint is derived from a segmentation of the high-quality reconstruction. We tested the algorithms with simulated data of differing quality to gauge their performance under different conditions. We showed that this algorithm performs far better than conventional methods and allows for a substantial reduction in imaging time.

We also examined multiple stopping rules for reconstruction algorithms that estimate the optimal point to terminate a reconstruction. Terminating a reconstruction before or after the optimal point will either result in a blurry or noisy reconstruction. All tested methods provided accurate estimates of the optimal stopping point for data with a moderate or small signal to noise ratio (SNR). Two of the methods do not require any knowledge of the noise level in the data, which makes them very practical.

Finally, we performed a dynamic CT experiment to study the scaling formation process. We injected a carbon steel flow cell with water supersaturated with BaSO_4

for over 150 hours while imaging it with a CT scanner. We used a modified version of one of the reconstructions algorithms we developed to reconstruct the dynamic data to improve the SNR of the reconstructions. Quantitative analysis of the reconstruction allowed us to gain new insights into the scaling formation process. From this analysis, we show that scaling formation has three distinct growth phases.

This thesis and the contributions within it shows the applicability of dynamic tomography for a wide range of fields and how it provides novel insight within these fields.

Resumé

Røntgen computer-tomografi (CT) er en kraftfuld målemetode, der gør det muligt at tage billeder af objekters indre struktur uden at skade dem. Forbedringer af røntgen kilder, røntgen detektorer og forøgelse af regnekraft har gjort det muligt at lave dynamisk tomografi, hvilket muliggør en række nye eksperimenter.

Dynamiske eksperimenter er udbredt i forskning af kulstoflagring, geotermisk energi, flerfasestrømning, klippedeformation og forskning relateret til olie- og gasindustrien. Dynamiske eksperimenter sætter krav til varigheden af et CT-skann i den forstand, at den skal være tilstrækkeligt kort til at optage det dynamiske fænomen. For at mindske skanningstiden kan det være nødvendigt at reducere antallet af projektioner eller eksponeringstiden per røntgenbillede, hvilket kan reducere billedkvaliteten af en CT-skanning.

I denne afhandling har vi udviklet nye rekonstruktionsalgoritmer for dynamiske eksperimenter, der kan bevare billedkvalitet på trods af mangelfulde data. Vi har designet algoritmerne til eksperimenter, hvor det er muligt at fortage et højkvalitetsskan, før det dynamiske eksperiment begynder. Disse algoritmer er baseret på *Simultaneous Iterative Reconstruction Technique* og den bedste af vores algoritmer har to tilføjelser, som forbedrer den. Den første af vores tilføjelser er, at vi benytter rekonstruktionen af højkvalitetsskannet til at initialisere det første tidsskridt i den dynamiske rekonstruktion. De næste tidsskridt bliver initialiseret med rekonstruktionen af forrige tidsskridt. Den anden tilføjelse er, at pixels i den dynamiske rekonstruktion er begrænset således, at de kun kan antage bestemte værdier. Begrænsningen er baseret på en segmentering af højkvalitetsrekonstruktionen. Vi testede algoritmerne ved at rekonstruere kunstige data af forskellig kvalitet. Disse test viste at vores algoritmer fungerer langt bedre end traditionelle algoritmer, hvilket tillader en betydelig reduktion i skanningstid.

Vi testede også forskellige metoder til at estimere det optimale punkt til at stoppe en rekonstruktionsalgoritme. Hvis en rekonstruktionsalgoritme bliver afsluttet for tidligt eller sent vil det hhv. resultere i en uskarp eller en støj fyldt rekonstruktion. Alle metoder vi testede, gav akkurate skøn af det optimale stoppunkt for data med et moderat eller højt støj indhold. To af metoder kræver ingen information om støjniveauet i data, hvilket gør dem nemme at anvende i praksis.

Til sidst udførte vi et dynamisk CT-eksperiment for at studere scale-afsætning. Scale er en fællesbetegnelse for salte og mineraler, som aflejrer sig på rør. Vi injicerede en strømningscelle i kulstofjern med vand, der var overmættet med BaSO_4 i over 150 timer imens vi fortog røntgen CT. Vi brugte en modificeret udgave af en af

de algoritmer vi udviklede for at forøge signal støj forholdet af rekonstruktionerne. Kvantitativ analyse af rekonstruktionerne gjorde det muligt for os at få ny indsigt i, hvordan scale-afsætning foregår. Ud fra vores analyse observerede vi at scale-afsætning gror i tre forskellige faser.

Denne afhandling og bidragene i den viser at dynamisk tomografi er brugbart i mange forskellige forskningsområder og at det giver ny viden inden for disse områder.

Preface

This PhD thesis was prepared in the Section for Visual Computing at the Department of Applied Mathematics and Computer Science (DTU Compute) at the Technical University of Denmark in fulfillment of the requirements for acquiring a PhD degree in computer science. The project was a part of the μ -Cracks: Flow & Deformation project and was funded by DTU Offshore – Danish Offshore Technology Centre.

During the project I developed tomographic reconstruction algorithms in order to improve the image quality obtained in dynamic computed tomography experiments at DTU Offshore. The project was initially focused on multiphase flow in porous media, but the scope of the project changed as it progressed. Three research papers have been prepared over the course of this project. Two of these papers have been peer reviewed while the third is still in preparation.

The project was supervised by Associate Professor Anders Nymark Christensen and co-supervised by Senior Researcher Henning Osholm Sørensen and Professor MSO Anders Bjorholm Dahl. The research was carried out almost exclusively at the Section for Visual Computing as unforeseen circumstances prevented me from collaborating with external researchers. I did, however, collaborate with Professor, dr.techn. Per Christian Hansen from the Scientific Computing Section at DTU Compute during the spring of 2021.

Kongens Lyngby, February 28, 2022

Peter Winkel Rasmussen

Acknowledgements

I would like to thank my three supervisors Anders Nymark Christensen, Henning Os-holm Sørensen and Anders Bjorholm Dahl. It is not necessarily a given that a PhD student has a single engaged supervisor, but I have been lucky enough to have three of them. Hands-on supervision is a luxury not many are afforded, and while it can result in confusion and panic, it can just as easily provide confidence and excitement. Your guidance has been greatly appreciated.

Furthermore, I would like to thank all my wonderful colleagues at the Section for Visual Computing. Both Christmas workshops, summer schools and “tour des of-fices” have been great fun with you. A special thanks to Morten Hannemose for indulging my \LaTeX inquiries and to Niels Jeppesen for assisting me with graph cuts.

Til Elers' Kollegium, afslutningen af min ph.d. markerer også afslutningen på min tid som Alumne på Elers' Kollegium. Jeg havde aldrig troet, at jeg skulle bo et så fan-tastisk sted med så mange søde mennesker. Jeg vil gerne takke alle de nuværende og tidligere alåmner, som jeg har haft fornøjelsen at stifte bekendtskab med i de sidste år.

Finally, a big thanks go out to my sister and her husband, who took the time to proofread this thesis, my parents and “drene”. One could not ask for a kinder or more supportive family or better friends.

Publication List

Published

Contribution A: P. W. Rasmussen et al. (December 2021). “Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction.” In: *Scientific Reports* 11.1, page 12501. ISSN: 20452322. DOI: [10.1038/s41598-021-91776-1](https://doi.org/10.1038/s41598-021-91776-1)

Accepted

Contribution B: P. C. Hansen, J. S. Jørgensen, and P. W. Rasmussen (2021). “Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography.” In: *The 21st International Conference on Computational Science and its Applications: Application of Numerical Analysis to Imaging Science*

In Preparation

Contribution C: I. A. Løge et al. (2022). “Crystal formation from high resolution 4D X-ray μ CT: Initiation, expansion, and densification.” In preparation

Contents

Summary	ii
Resumé	iv
Preface	vii
Acknowledgements	viii
Publication List	x
Published	x
Accepted	x
In Preparation	x
Contents	xi
1 Introduction	2
1.1 Oil Production in the North Sea	2
1.2 This Research Project	4
1.3 X-ray Imaging of Dynamic Processes	5
2 X-ray Radiation	12
2.1 Generation of X-ray Radiation	13
2.2 The Attenuation of Light	16
3 X-ray Computed Tomography	20
3.1 The Radon Transform	20
3.2 Inverse Radon Transform & Fourier Reconstruction	22
3.3 Analytical Reconstruction	25
3.4 Algebraic Reconstruction	27
3.5 Recent Developments in Reconstruction Algorithms	32
4 Practical Aspects of CT Experiments	36
4.1 X-ray Tomography Instruments	36
4.2 Detection of X-ray Radiation	37
4.3 From Raw Data to Sinograms	40

4.4	Noise	41
4.5	Ring Artefacts	42
4.6	Cone Beam Geometry	44
4.7	Misalignment of Geometry and Artefacts	45
4.8	Auto-focus	49
4.9	Consistency of Projection Data	50
4.10	Implementation Details	53
5	Contribution A: Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction	56
6	Contribution B: Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography	72
7	Contribution C: Crystal formation from high resolution 4D X-ray μCT: Initiation, expansion, and densification	84
8	Conclusion	104
A	CT Scanners at DTU Offshore	106
B	Sample Movement Corrections for Contribution C	108
	Bibliography	110

CHAPTER 1

Introduction

1.1 Oil Production in the North Sea

Denmark started recovering oil and gas from the North Sea in 1972 when the first oil field in the Danish part of the North Sea started production. What followed was a rapid exploration of the area in the late 1970s and early '80s, along with the construction of offshore platforms to exploit the newfound resources. A map of the Danish North Sea can be seen in [Figure 1.1](#) where the producing fields are marked. In total 19 fields that produce either oil or gas have been developed since the exploration of the North Sea started in the 1960s (Danish Energy Agency, 2021a; TotalEnergies, 2021).




Figure 1.1: Danish oil and gas fields in the North Sea. The figure is obtained from the Danish Energy Agency (2021a) with permission.

The development of the fields has been of significant importance for the Danish economy (Nordsøfonden, 2021). Oil and gas are valuable natural resources as they provide energy, heat and electricity. Additionally, they are also processed to create the chemical feedstock used in many commonplace products such as most plastics, artificial fertilisers and pharmaceuticals (International Energy Agency, 2018). In the mid-90s the production of oil surpassed consumption making Denmark a net exporter of oil until the mid-10s which can be seen in Figure 1.2. This decline since 2005 is simply a result of ageing oil fields and increased downtime during production due to maintenance (Danish Energy Agency, 2017).

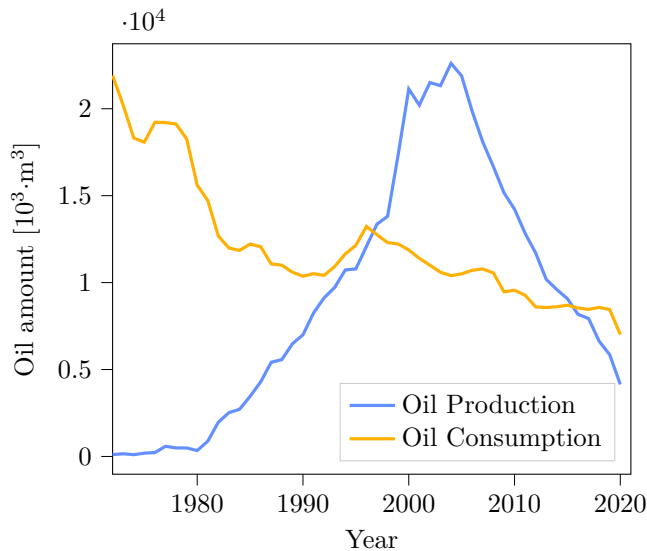


Figure 1.2: Oil production and consumption in Denmark since 1972. Data are acquired from Danish Energy Agency (2021b) and BP (2020).

The decline in oil production prompted the Danish Parliament and a consortium of companies in the oil industry to create the Danish Hydrocarbon Research and Technology Centre (DHRTC) in 2014, which has the overall goal of supporting the Danish oil industry through research. Initially, the centre was focused on enhancing recovery of oil and gas, reducing costs of operation and reducing the environmental impact of production.

However, the political landscape has changed since its inception due to the increasingly overwhelming evidence for the harmful effects of anthropogenic global warming (Masson-Delmotte et al., 2021). In 2020 it was decided that further exploration of North Sea should be cancelled and in 2021 it was decided by The Danish Parliament that extractions of hydrocarbons from the North Sea is finished by the year 2050 (The Danish Parliament, 2021).

These changes has caused a shift in the reasearch activities of the DHRTC, which

now seeks to support the transition from fossil fuels to sustainable energy sources thereby limiting the impact of Danish offshore activities as much as possible (The Danish Hydrocarbon Research and Technology Centre, 2022). The research scope of the centre have broadened and now include research into the abandonment of oil and gas fields, CO₂ storage, eliminating byproducts from oil production and improve materials used for oil production. The DHRTC has now also been renamed to DTU Offshore – Danish Offshore Technology Centre (DTU Offshore) to reflect this change in research.

1.2 This Research Project

This PhD project was initially started to support research into enhanced oil recovery from the Lower Cretaceous reservoirs by developing the tools needed for in situ measurements of samples from these reservoirs. The composition of Lower Cretaceous chalk poses unique challenges to the production of oil, and imaging samples of this particular kind of chalk is necessary to improve the understanding of it.

My role was specifically to create algorithms for tomographic reconstruction suited for dynamic imaging of multiphase flow in porous media. However, the scope of my project broadened, and the focus shifted to more general applications of dynamic X-ray tomography for the types of research that DTU Offshore performs. There are three contributions included in this thesis which are:

1. **Contribution A:** Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction.
2. **Contribution B:** Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography.
3. **Contribution C:** Crystal formation from high resolution 4D X-ray μ CT: Initiation, expansion and densification.

The first contribution presents tomographic reconstruction algorithms suited for in situ core flooding experiments. The algorithm maintains image quality despite high noise and few projections. The robustness of the algorithm means that imaging time can be reduced, which improves temporal resolution. [Section 1.3](#) and [Chapters 2](#) and [5](#) are meant to provide the of the background for this contribution.

The second contribution is also related to tomographic reconstruction algorithms. In this contribution, we examine different heuristics for when to stop reconstruction algorithms. Stopping too early results in a blurry reconstruction, while stopping too late yields a noisy one, which makes it desirable to estimate this ideal stopping point. [Chapter 5](#) provides the background for this contribution.

The third and final contribution presents a study of scale formation by using dynamic X-ray tomography. A flow cell in carbon steel was continuously injected with water supersaturated with BaSO₄ for over 150 hours. Imaging the cell during the

injection and using a modified reconstruction method from [Contribution A](#) made it possible to perform a detailed study of the scaling process. [Chapter 4](#) covers the practical aspects of going from raw data to the final reconstruction. These steps are not described in [Contribution C](#) as they are not relevant for the research topic; however, it is still necessary to account for them to obtain a high-quality reconstruction.

1.3 X-ray Imaging of Dynamic Processes

X-ray imaging has seen increasingly more use in petrophysics in the last ten years as it provides a reliable way to image an opaque sample in 3D. These gradual improvements in spatial and temporal resolution partially account for this growth but increasing computer power has been even more important as it has allowed much larger samples to be imaged (Wildenschild and A. P. Sheppard, 2013). A detailed explanation of X-ray radiation and the theory behind tomography is given in [Chapters 2](#) and [3](#). At this point, it is sufficient to know that X-ray computed tomography (CT) is an imaging technique based on X-ray radiation that makes non-destructive imaging of objects possible in 3D. There is a plethora of uses for tomography in geosciences, pore-scale imaging and porous media of which many are described in the reviews Wildenschild and A. P. Sheppard (2013), Blunt et al. (2013), Cnudde and M. Boone (2013), and Bultreys, De Boever, and Cnudde (2016).

1.3.1 Enhanced Oil Recovery and Core Flooding

As previously mentioned, research in enhanced oil and gas recovery was one of the early main goals of DTU Offshore. Oil recovery is, broadly speaking, separated into three different phases, which are the primary, secondary and tertiary phases.

During primary recovery, oil flows freely to the oil well due to the high pressure in the reservoir. The pressure of the oil field drops continuously during recovery, and it will eventually be too low for primary recovery. Reservoir pressure can be sustained artificially by injecting either fluid or gas. Secondary recovery refers to maintaining reservoir pressure in this way. Seawater is commonly used in the North Sea during secondary recovery due to its availability and ability to recover oil from chalk (Zhang, Tweheyo, and Austad, 2007).

While secondary recovery seeks to sustain pressure in the reservoir, tertiary recovery attempts to modify the conditions in the reservoir to increase the recovery. The tertiary phase is, therefore, also known as Enhanced Oil Recovery (EOR). There are multiple types of EOR, each of them using different approaches to improving recovery (Manrique et al., 2010). Many of these methods are researched at DTU Offshore, such as injections with carbonated water, foams, smart water, which is water with a modified ionic composition, etc. (Nielsen, Nesterov, and Shapiro, 2016; Mohammadkhani, Shahverdi, and Esfahany, 2018; Seyyedi, Sohrabi, et al., 2018; Eftekhari and Farajzadeh, 2017; Seyyedi and Sohrabi, 2017).

A typical experiment when researching EOR is the core flooding experiment. A core flooding experiment seeks to measure how much of a given quantity of oil is produced from a sample of sedimentary rock known as a core plug. The core plug is kept under pressures and temperatures similar to the conditions found in reservoirs to mimic real-world conditions, and water is injected as it would be during secondary or tertiary recovery. The quantity of oil in a core plug can be measured precisely before injection and is usually referred to as original oil in place (OOIP). The OOIP makes it possible to calculate the recovery factor i.e. how much oil is recovered during the injection. The recovery factor (RF) is often given as a function of the pore volumes injected (PVI) in the sample.

An example of the results of a core flooding experiment, as reported in Seyyedi, Tagliaferri, et al. (2018), can be seen in Figure 1.3. In this experiment, the effect of the salinity and temperature of the flooding fluid was tested. The horizontal axis displays the number of PVI, and the vertical axis shows the RF as a percentage of the OOIP. During the first 5 PVI, seawater (SW) is injected, and the RF is steadily climbing until it reaches an equilibrium. Secondary recovery is initiated by injecting low salinity seawater (LSSW), which increases the RF. The final injection consists of heated seawater, and we see that this does not increase the RF for this experiment.

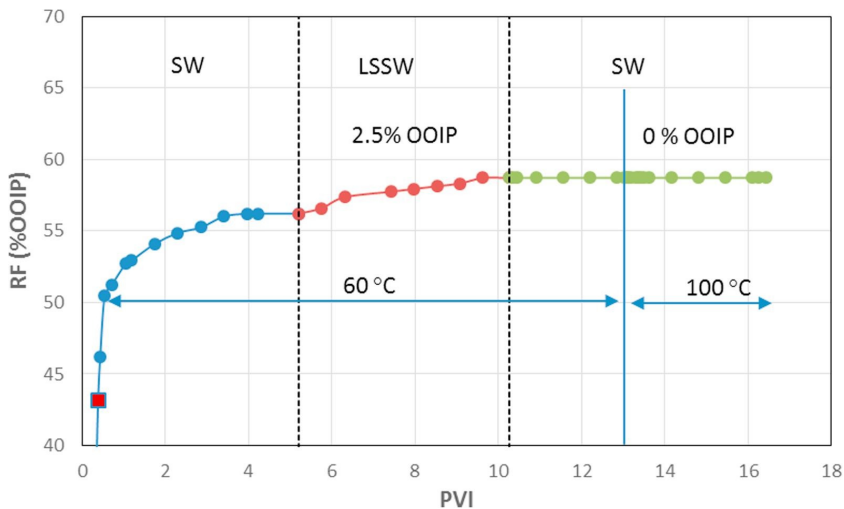


Figure 1.3: The recovery factor (RF) expressed as a percentage of the original oil in place (OOIP) as a function of pore volumes injected (PVI) during a core flooding experiment. SW denotes that seawater is injected while LSSW denotes low salinity seawater is injected. The figure obtained from Seyyedi, Tagliaferri, et al. (2018) with permission.

Core flooding experiments such as the previously mentioned one provide a clear answer to whether an EOR technique increases the RF, but they do not necessarily provide an answer to the underlying mechanism behind the techniques. Instead, other modalities have to be used, which is the case in Mohammadkhani, Shahverdi,

and Esfahany (2018) where the ionic composition of the injection fluid was examined before and after injection to understand the mechanism behind low salinity water flooding.


Another way of examining the underlying methods of smart water would be performing X-ray computed tomography (CT) on the sample. X-ray CT can be used to image the sample before and after injection, and this can illuminate how the oil is displaced. However, it would be far more useful if the sample was imaged during flooding i.e. in situ. Dynamic tomography is used in a plethora of research besides EOR, such as the physics of fluid flow, carbon storage and deformation of porous media. The following section will detail different applications of tomography in the oil and gas industry and related research fields.

1.3.2 Core Flooding with In Situ Measurements


Core flooding experiments are kept under ambient pressure and temperature, as mentioned in the previous section. Real-world conditions allow for accurate studies of fluid and or gas interactions with the sample. To replicate the conditions of the North Sea samples have to be heated to approximately 60 °C and kept at pressures hundreds of times larger than atmospheric pressures. Ensuring that these conditions are maintained while also having the capability of fluid flow results in fairly complex experimental setups. This can be difficult to integrate with a CT scanner and it is, therefore, often necessary to design a flow cell specifically for in situ measurements.

An example of such a flow cell can be seen in Figure 1.4. This cell was designed during my project specifically with in situ measurements in mind. The cell is designed to hold 30 MPa of pressure and has been tested at 23 MPa and 90 °C. The cell has the possibility of applying axial pressure on the samples which makes it possible to study deformation of porous media. We have only been able to perform preliminary in situ experiments at room temperature and ambient pressures but the cell seem to perform as expected. An example of a piece of chalk imaged in the cell can be seen in Figure 1.4b. Core flooding is not only used for EOR experiments but whenever it is necessary to replicate real-world conditions. In recent years core flooding experiments that explore the potential of different kinds of rock sediments to store CO₂ has been performed with CT capabilities (Y. Yang, Y. Li, et al., 2020; Y. Yang, Bruns, et al., 2018; Y. Yang, Hakim, et al., 2018; Menke, Bijeljic, Andrew, et al., 2015; Menke, Andrew, et al., 2016; Menke, Bijeljic, and Blunt, 2017; Menke, Reynolds, et al., 2018; Scanziani et al., 2020; Øren et al., 2019).

Y. Yang, Y. Li, et al. (2020) performed in situ measurement of a core flooding experiment where carbonates were injected with carbonated brine to examine their storage potential. The injected brine is acidic due to the CO₂ dissolved in it, which causes it to dissolve the carbonate when injected. The authors examined the dissolution pattern for different pore structures, which is needed for accurate carbon storage capacity predictions. The imaging was not dynamic but was instead done at discrete points during the experiment, a trait that is common for many experiments that use



(a) A vertical cross section of a schematic of the tri-axial flow cell developed during my project. The three inlets at the bottom are used to create the confining and axial pressure and fluid injection. The grainy dark grey indicates stainless steel, while the grainy light grey indicates steel.



(b) A vertical cross section of a piece of chalk imaged while in a tri-axial flow cell.

Figure 1.4: A schematic of a flow cell designed for in situ measurements along with a CT scan of a piece of chalk imaged in this flow cell.

laboratory X-ray sources. Genuine dynamic experiments require much faster imaging which is only available at synchrotron facilities. An example of a similar experiment performed at the Diamond Lightsource synchrotron is found in Menke, Andrew, et al. (2016), which also tested reactive flow. The fast imaging made it possible to scan the tested samples between 51 and 94 times over a 2 hour period, which is approximately 100-200 times faster than conventional imaging.

The fast imaging of synchrotrons, along with improvements in spatial resolutions, have made it possible to resolve the pore structure of many types of sedimentary rock. Pore-scale resolution has made it possible to examine the physics of flow at the fundamental level. In Berg et al. (2013), Singh, Scholl, et al. (2017), Singh, Menke, et al. (2017), Gao et al. (2020), Spurin et al. (2019), and Armstrong et al. (2016), the authors use synchrotrons to study Haines jumps and snap-off. Haines jumps occur during drainage where a fluid such as oil is injected into a porous medium. As the oil is injected, it will occasionally suddenly jump forward at the interface of the oil. The pressure inside the medium will drop briefly during a Haines jump, and the fluid will redistribute afterwards (Sun and Santamarina, 2019).

Snap-off occurs when a wetting fluid replaces a non-wetting fluid in a medium. An example of this is when water is injected into an oil-filled porous medium. Water will displace the oil gradually in a piston-like way; however, it will occasionally rapidly move forward and, in the process, it will surround some of the oil, which is now trapped in the porous medium.

Both Haines jumps and snap-off are very fast events and can therefore only be imaged at synchrotrons. The phenomena are, however, important to understand. Efficient recovery of oil requires that snap-off is minimised while carbon storage seeks to maximise snap-off.

1.3.3 Deformation

Tri-axial cells, such as the one shown in Figure 1.4, can also be used to study the compaction and faulting of porous media (Renard, Cordonnier, Dysthe, et al., 2016; Renard, Cordonnier, Kobchenko, et al., 2017; Renard, Weiss, et al., 2018; Renard, McBeck, et al., 2019; Chen et al., 2020). Compaction bands in porous media are regions where material have been compressed, which lowers porosity and fluid permeability (Sternlof et al., 2006). This is relevant for both managing aquifers and oil reservoirs as the presence of compaction bands drastically inhibit fluid flow. In Chen et al. (2020) carbonates were imaged in a tri-axial cell at increasing axial pressures. The authors were able to identify the nucleation and dynamics of compaction bands by linking the reconstructions from the CT scans to the stress-strain curves.

Tri-axial cells are used to study faults in rocks sample, which is one of the main approaches to understand larger faults such as earthquakes (Ben-Zion, 2008). Renard, Cordonnier, Dysthe, et al. (2016) describes a tri-axial cell developed for X-ray imaging that can replicate conditions 10 km underground. The cell was used in Renard, Cordonnier, Kobchenko, et al. (2017), Renard, Weiss, et al. (2018), and Renard,

McBeck, et al. (2019) where samples were imaged at The European Synchrotron. The fast imaging made it possible to image microfractures as they arise, join and finally result in macroscopic faults. Experiments such as these provide insights into the conditions present before faults occur and can improve the ability to predict events such as earthquakes.

1.3.4 Scaling and Corrosion

In situ X-ray CT has also been used to study scaling and corrosion processes, where scaling refers to the build-up of salts and minerals on pipes. (Connolly et al., 2006; Barker et al., 2018; Godinho and P. Withers, 2018; Oliveira et al., 2019; Haaksman et al., 2017). Crude oil and natural gas often contain corrosive chemicals that damage the pipes they are extracted via and salts that deposit them, which inhibits fluid flow causing economic loss (J. Li et al., 2017). Understanding the scaling process is key to preventing downtime of oil wells (J. Li et al., 2017). Connolly et al. (2006) demonstrates that synchrotron in situ X-ray CT makes it possible to study localised corrosion behaviour in an aluminium aerospace alloy. The study examines whether friction stir welding enhances the susceptibility to local corrosion and how localised corrosion develops. This is not easily measured using traditional methods, which is why X-ray CT was used.

Scaling and corrosion is a concern in many fields and affects heat exchangers, geothermal plants and water treatment, and tomography is just as useful within these fields as the previously mentioned ones (Touir et al., 2009; Mundhenk et al., 2013; Sassi and Mujtaba, 2011). In Godinho and P. Withers (2018), the authors sought to study how precipitation occurs in porous media. The precipitation process occurs naturally in geological formations, but it can also happen in geothermal energy recovery. The build-up of precipitates can inhibit fluid transport, which would be problematic for a geothermal plant.

The authors injected a piece of quartz with a solution that was supersaturated with calcite to study the precipitation process. The precipitation and crystallisation of the mineral were imaged at discrete points in the experiment with a CT scanner. Based on their findings, they show that the overall growth rate of precipitates cannot be estimated from the bulk fluid composition as it is modified by fluid transport. It is, therefore, necessary to image the sample as the precipitation occurs to see how the fluid transport changes.

1.3.5 Our Contribution to this Field

The previous sections show that X-ray CT is a valuable technique that has a high level of applicability in a wide range of fields. Many of the examples given use either synchrotrons for fast imaging of dynamic processes or laboratory CT scanners for static or time-lapse imaging. Based on the literature, we saw a need to enable laboratory CT scanners to perform dynamic in situ experiments.

We achieved this by developing new algorithms to reconstruct data from CT scanners, which is presented in [Contribution A](#), and demonstrated the practical value of these algorithms in [Contribution C](#), where they were used to reconstruct data from a dynamic in situ experiment where the scaling formation process of BaSO_4 was examined. By using dynamic X-ray CT, we were able to study the scale morphology, rate of growth and deposition of material during the scaling formation process.

CHAPTER 2

X-ray Radiation

Wilhelm Röntgen discovered X-ray radiation, quite by accident, on November 8th, 1895, where he quickly noticed their penetrating nature (American Physical Society, 2001). Röntgen started investigated with the newfound ray to understand its properties and submitted his first work by the end of the year (Röntgen, 1896). In this period he famously imaged the hand of his wife where her wedding ring noticeable attenuated by the radiation, see [Figure 2.1](#).



Figure 2.1: An X-ray image of the hand of Anna Bertha Ludwig imaged by Wilhelm Röntgen (Röntgen, 1895) (CC BY-NC 4.0).

The penetrating nature of X-ray radiation is related to its short wavelength, between 0.1 and 10 nanometers or conversely, its high energies, approximately between 1 keV and 100 keV. The ability of X-ray radiation to penetrate matter is also highly dependent on the material it interacts with. Denser materials or materials with a higher atomic number will attenuate more than lighter, less dense elements. This can readily be seen in [Figure 2.1](#), where the tissue of the hand is transparent while the bones and

wedding band are more opaque.

2.1 Generation of X-ray Radiation

All generation of X-ray radiation is done via acceleration of electrons (Holbrow et al., 2009b). From classical electromagnetism, it is known that acceleration or deceleration of a charged particle creates an electric dipole which makes the charged particle emit electromagnetic radiation (Bilderback, Elleaume, and Weckert, 2005). X-ray radiation used for tomography is generated with an X-ray tube or from a synchrotron light source. The principles behind these are given in the following sections.

2.1.1 X-ray tube

The most common way of producing X-ray photons is with an X-ray tube. An X-ray tube consists of a cathode and an anode cased in a vacuum tube. A drawing of an X-ray tube can be seen in Figure 2.2, where the cathode and anode are displayed in bright red.

The cathode usually consists of a tungsten filament which provides the electrons through thermionic emission used to create the X-ray radiation. The filament is heated with an electric current such that electrons in the metal overcome their binding energy. The binding energy E_v that holds the free electrons in the metal in place has two main components.

The first component consists of electrostatic forces that prevent the free electrons from escaping as they move towards the surface. A net positive charge is created below electrons that are sufficiently energetic to move past the outer ion layer of the metal. The combination of a negatively charged electron and surplus of positive charge below it creates a dipole layer which further slows down the electron trying to escape (Buzug, 2008).

The second component of the binding energy is caused by the mirror charge left behind when an electron is removed from the metal. The mirror charge creates an electric field between it and the escaping electron, which requires additional energy to overcome (Buzug, 2008). Overcoming E_v requires that the filament is heated to approximately 2400 K, which is why tungsten is commonly used as it has a high melting point. The emission current density of the filament is a function of the temperature of the filament, and it can be described by the Richardson-Dushman equation that states

$$j_e = C_{RD} T^2 \exp\left(-\frac{\varphi}{kT}\right), \text{ with } C_{RD} = \frac{4\pi m_e k^2 e}{h^3}, \quad (2.1.1)$$

where j_e is the emission current density, C_{RD} is the Richardson-Dushman constant, T is the filament temperature, φ is the work function of the metal, k is the Boltzmann constant, m_e is the mass of the electron, e is the charge of the electron and h is




Figure 2.2: A drawing of a water cooled X-ray tube (Coolth, 2010). C and A indicate the cathode and anode, W_{in} and W_{out} are the inlet and outlet for the water cooling and U_h and U_a are voltage across the coil and between the cathode and anode. The green X indicate the generated radiation.

Planck's constant. A cloud of electrons forms around the filament as it heats, which start to accelerate toward the positively charged anode. Figure 2.2 shows this, where electrons move from the coiled cathode toward the anode. The size of the electron beam spot i.e. the point of impact on the anode can be controlled with an electrode. A small beam spot results in a tighter focus, which increases the photon flux at the expense of having a smaller field of view. The actual generation of X-ray radiation is a result of electrons interacting with the atoms of the anode. Electrons are decelerated by the atomic nucleus and orbital electrons in the anode. The decelerating electrons form electric dipoles that radiate electromagnetic waves. The filament electrons can interact with the nucleus and orbital electrons in multiple ways, which determines the energy of the electromagnetic waves that are radiated (Buzug, 2008).

The efficiency of X-ray tubes is extremely poor. Only 1% of the energy of the electrons is converted into X-ray radiation while the remaining energy is lost to heat (Hemberg, 2004). The heat generated by the electron beam represents a major problem as it can easily melt or damage the anode. It is, therefore, necessary to cool the anode, which can be done in multiple ways. The schematic of the tube in Figure 2.2 has active water cooling. The anode can also be designed in a way that allows for a more intense electron beam. For example, the anode can consist of a rotating piece of metal or, in more advanced anodes, consist of a stream of liquid metal. Both of these designs make it possible to increase the emission current density of the cathode without damaging the anode.




Figure 2.3: The radiation spectrum of a tungsten anode with different tube voltages, which are shown in the legend. The horizontal axis displays the photon energy, and the vertical axis displays the normalised count. The spectrum is based on the results in J. Boone and Seibert (1997).

An example of the output spectrum of an X-ray tube at different voltages can be seen in [Figure 2.3](#). The smooth part is attributed to bremsstrahlung that has several underlying mechanisms (Ebel, 1999). The photoelectric effect is responsible for the two characteristic peaks in the spectrum. It occurs when an electron is absorbed by the anode and a photon is emitted instead. This photon will have an energy equivalent to the orbital energy of the electron. We also see that the normalised count is zero for energies above the tube voltage as an electron cannot produce a photon with more energy than itself.

2.1.2 Synchrotron Light Sources

Synchrotrons are the second commonly used method of obtaining X-ray photons for tomography. These machines are highly specialised and only exist at large research facilities. The interest in these facilities have gradually increased, and today there are more than 50 synchrotrons worldwide (Bilderback, Elleaume, and Weckert, 2005). However, beam time is still valuable at these sites and obtaining it is not a given.

A synchrotron is a circular particle accelerator that accelerates charged particles with radio waves emitted in radiofrequency cavities and controls their movement using magnets. The magnets bend the electron beam such that the electrons remain in orbit. Bending the charged particles require that they are accelerated towards the centre of the storage ring. This acceleration causes the charged particles to emit synchrotron

radiation. The power emitted by a charged particle that performs a uniform circular motion in a constant magnetic field \mathbf{B} is given by the relativistic Larmor formula (Tsang, 1998)

$$\mathcal{W} = \frac{q^4}{6\pi m^2} v^2 B^2 \gamma^2, \text{ with } \gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}}, \quad (2.1.2)$$

where q is the particle charge, m is the particle mass, v is the speed of the particle, B is the size of the magnetic field and c is the speed of light. We see from this formula the power emitted is proportional to $B^2 \gamma^2$ and inversely proportional to m^2 . The dependence of the mass of the particle means that electrons are always used in synchrotron light sources as they are approximately 2000 times lighter than the proton.

Relativistic effects force the radiation into a narrow cone in front of the electron beam, which makes the brilliance much larger in synchrotrons than in X-ray tubes (Tsang, 1998). This effect also means that the beam is highly collimated, which increases the spatial resolution of measurements. Additionally, the light emitted is polarised, a characteristic that enables spectroscopy and diffraction experiments. The energy emitted has a broad spectrum, but it is possible to tune it by using monochromators (Bilderback, Elleaume, and Weckert, 2005).

Modern synchrotrons do not solely rely on bending magnets to produce X-ray radiation, instead, they use more advanced magnets such as undulators to produce it. Undulators have a sinusoidal magnetic field rather than the uniform magnetic field used by bending magnets. The size and spatial period of the magnetic field of undulators make it possible to control the wavelength of X-ray radiation produced by relativistic electrons in an undulator (Bilderback, Elleaume, and Weckert, 2005).

The monochromaticity, polarisation and brilliance of synchrotron light allows for experiments which is simply not possible with laboratory sources.

2.2 The Attenuation of Light

The attenuation of light is a complicated phenomenon that depends on multiple physical processes. The processes are dependent on the energy of the light and the material in question. The material dependence can readily be seen in Figure 2.1. Fortunately, most of the processes are only relevant at small or large energies, which makes it possible to make a simple approximation of the attenuation of light. However, before approximation is described, a brief explanation will be given of the processes that attenuate light.

2.2.1 The Interaction of Light With Matter

Light can be attenuated via either scattering or absorption. The dominant source of attenuation depends on the energy of the incident X-ray photons. Figure 2.4

shows the attenuation coefficient of light for iron. The total attenuation coefficient is shown in orange from 1 keV to 1 MeV where the filled blue area between 5 keV and 140 keV denotes the energies associated with X-ray radiation. Below this area, we have “tender” X-ray radiation and UV light and above it, we have gamma rays, although there is no clear cut-off between X-ray radiation and gamma rays. The following sections will briefly describe the different contributions to the attenuation coefficient.

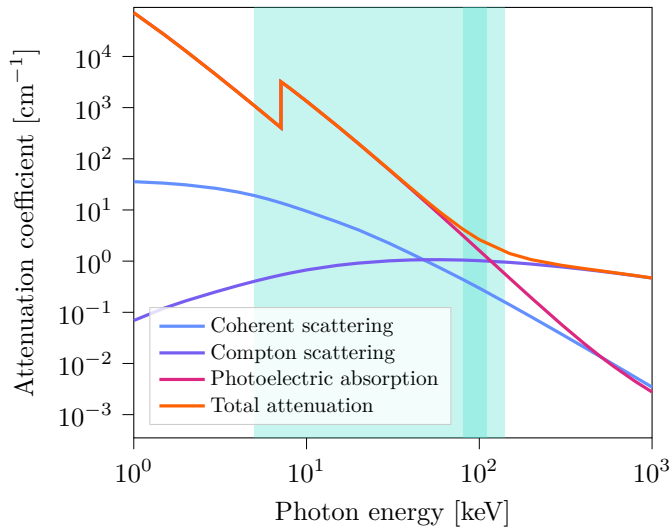


Figure 2.4: An example of the attenuation coefficient of light in iron between 1 keV and 1 MeV. The total attenuation coefficient is shown in orange and its contributions are shown as the remaining graphs defined in the legend. The light blue band between 5 keV and 140 keV displays the energies attributed to X-ray radiation while the dark blue band between 80 keV and 110 keV display the energies relevant for this project. The attenuation coefficients are obtained with the NIST XCOM: Photon Cross Sections Database (Berger et al., 2022).

2.2.1.1 Coherent Scattering

Coherent scattering refers to scattering processes where the path of the incoming photon is changed when it scatters, while the energy remains unchanged (elastic scattering). There are two contributions to this process, Rayleigh and Thompson scattering. Rayleigh scattering occurs when light scatters a bound particle, while Thompson scattering happens when light scatters off a single free or quasi-free charged particle. The incident light wave induces a dipole in the charged particle it hits during the scattering. This dipole then reemits the light without altering its wavelength. Both of these processes predominantly happen at low energies. Rayleigh scattering occurs in the limit where the wavelength of the incident light is much larger than

the size of the particle it scatters off, and Thompson scattering occurs when the energy of the light is much smaller than the mass-energy of the particle it scatters off (Pierrehumbert, 2015). The sum of these effects is shown as the blue line in Figure 2.4, where we see it decreases steadily as the photon energy increases.

2.2.1.2 Compton Scattering

Compton scattering is the high energy limit of Thompson scattering where the incident photon has sufficient energy to make the electron scatters off recoil. This recoil reduces the energy of the photon (inelastic scattering), however, unlike photoelectric absorption (which is described below), the photon continues to travel through the medium, and it might have sufficient energy to interact again (Holbrow et al., 2009b). Besides being a function of the photon energy, Compton scattering depends on the electron density of the material, which means most biological tissue have an equivalent level of Compton scattering.

2.2.1.3 Photoelectric Absorption

Photoelectric absorption (also known as the the photoelectric effect) occurs when a photon removes an orbital electron bound to an atom while being absorbed in the process. This cannot occur if the energy of the incident photon is smaller than the binding energy of the electron. The binding energy is determined by which atomic orbital the electron inhabits, where electrons further away from the nucleus are bound less tightly. The binding energy of each orbital is quantised, which means that a small increase in photon energy can result in a large increase in attenuation (Holbrow et al., 2009a). The quantisation can readily be seen in Figure 2.4 where the photoelectric absorption is indicated in magenta. There is a noticeable increase in attenuation at approximately 7 keV, which corresponds to the binding energy of an inner orbital electron in iron.

2.2.2 Lambert-Beer's Law

It is apparent from Figure 2.4 that the attenuation of light is a complex process with multiple components. However, Compton scattering and the photoelectric effect are the main contributions to the attenuation coefficient for the energies we are concerned with, in this project, which means the attenuation coefficient μ can be modelled as independent of the photon energy. The intensity of the light source is a function of the distance it has traversed $I(s)$. After passing through a homogeneous object of width Δs with an attenuation coefficient μ the intensity will be given by

$$I(s + \Delta s) = I(s) - \mu \Delta s I(s). \quad (2.2.1)$$

Rearranging the equation and going to the infinitesimal limit of s we obtain:

$$\lim_{\Delta s \rightarrow 0} \frac{I(s + \Delta s) - I(s)}{\Delta s} = \frac{dI}{ds} = -\mu I(s). \quad (2.2.2)$$

We separate the variables and integrate to solve the equation

$$\int \frac{dI}{I(s)} = - \int \mu ds \quad (2.2.3)$$

$$-\log(I(s)) = -\mu s + C \quad (2.2.4)$$

$$I(s) = e^{-\mu s + C} \quad (2.2.5)$$

If we solve this for the special case where $I(0) = I_0$ we get

$$I(s) = I_0 e^{-\mu s}. \quad (2.2.6)$$

This equation is known as the Lambert-Beer law and it provides a simple way to recover the attenuation of a homogenous object. However, in most cases, objects will be composed of multiple materials all of which with a unique attenuation value. In this case Lambert-Beer's law will read

$$I(s) = I_0 \exp \left[- \int_0^s \mu(x) dx \right]. \quad (2.2.7)$$

Noticeably, we see that we cannot recover $\mu(s)$ by measuring $I(s)$ as we cannot tell how the light was attenuated along the path from 0 to s since this information is lost in the line integral. It is possible to recover the lost information with a full set of line integrals which means multiple measurements are required. The details of this will be explained in Section 3.1.

CHAPTER 3

X-ray Computed Tomography

The following chapter provides the mathematical background of tomography. The background includes a description of the Radon transform, analytical reconstruction algorithms and algebraic reconstruction algorithms. The chapter is concluded with a description of recently developed reconstruction algorithms and how [Contribution A](#) relates to these.

3.1 The Radon Transform

In 1917 Johann Radon published the paper “*On the Determination of Functions From Their Integral Values Along Certain Manifolds*” in which he, unknowingly, formed the basis of tomography (Radon, 1986). In this paper, he examined a function $f(\mathbf{x}) = f(x, y) \in \mathbb{R}^2$ integrated along each straight line in the space of straight lines $\mathbf{L} \subset \mathbb{R}^2$, which results in the function $F(\mathbf{L})$. He discovered that $f(\mathbf{x})$ can be recovered from $F(\mathbf{L})$. Immediately, this seems reminiscent of what we learned about Lambert-Beer’s Law in [Section 2.2.2](#), where the attenuation coefficients of an object can be described as $\mu(\mathbf{x})$ and the passing of an X-ray through an object can be viewed as a line integral. The function $f(x, y) \in \mathbb{R}^2$ is subject to three regularity conditions:

1. $f(x, y)$ is continuous.
2. The integral

$$\iint \frac{|f(x, y)|}{\sqrt{x^2 + y^2}} dx dy$$

must converge.

3. For an arbitrary point on the plane $\mathbf{p} = (x', y')$ we have

$$\lim_{r \rightarrow \infty} \frac{1}{2\pi} \int_0^{2\pi} f(x' + r \cos \phi, y' + r \sin \phi) d\phi = 0$$

with this, the Radon transform of $f(x, y)$ is defined as

$$F(\mathbf{L}) = \mathcal{R}f(\mathbf{L}) = \int_{\mathbf{L}} f(x, y) d\ell, \quad (3.1.1)$$

where $d\ell$ is an infinitesimal piece the straight lines in \mathbf{L} . It is more convenient to represent \mathbf{L} with its distance from the origin s and the angle θ between the x -axis and the normal vector of \mathbf{L} , which is the rotation between the x -axis the unit normal vector to \mathbf{L} . The coordinates of these lines will be given by

$$\mathbf{L}_{\theta, s} = \{(x, y) \mid x \cos \theta + y \sin \theta = s\} \quad (3.1.2)$$

with this the Radon transform is given by (Feeman, 2015c)

$$[\mathcal{R}f](\theta, s) = \int_{\mathbf{L}_{\theta, s}} f(s \cos \theta - \ell \sin \theta, s \sin \theta + \ell \cos \theta) d\ell \quad (3.1.3)$$




Figure 3.1: A depiction of the Radon transform of a disk for a specific angle θ .

An example of the Radon transform of a disk that is offset from the origin is shown in **Figure 3.1** for a specific angle θ . Line integrals are performed along the blue lines, which are perpendicular to the axis labelled s . We see that the Radon transform equals zero for most values of s as the value of $f(\mathbf{x})$ is equal to zero for most of the plane. Between the edge of the disk and the centre of it, we see an increase in the value of $\mathcal{R}f(s)$ since the line integrals pass over more of the disk. Note that the Radon transform is almost symmetric around the axis of rotation, which means $[\mathcal{R}f](\theta, s) = [\mathcal{R}f](\theta + \pi, -s)$. We, therefore, get the full Radon transform from




Figure 3.2: The Radon transform of an off-center disk in Figure 3.1 for $\theta \in [0, \pi[$. The horizontal axis indicates the angle θ of the straight line $L_{\theta,s}$, while the vertical axis shows the position on the line s . The value of the Radon transform corresponds to the colour of the image, with black being equal to zero and white being equal to the maximal value in the sinogram.

$\theta \in [0, \pi[$. A Radon transform of the disk shown in Figure 3.1 for all values of θ results in what is colloquially known as a sinogram, the name stemming from the fact the Radon transform of off-centre points are sinusoids. The sinogram of the disk is shown in Figure 3.2, where the sinusoidal behaviour can be seen.

3.2 Inverse Radon Transform & Fourier Reconstruction

While Radon did provide a method of inversion in his paper, it will not be presented in the following. Instead, a method that utilises the Fourier transform will be explained as this method is a prerequisite to the commonly used filtered backprojection. Fourier reconstruction is based on the Fourier slice theorem, also known as the projection-slice theorem. The basic steps of this reconstruction method are:

1. Project a function $f(x, y)$ to a line using the Radon transform $[\mathcal{R}f](\theta, s)$.

2. Perform a one-dimensional Fourier transforms on Radon transformed function $\mathcal{F}_1(\mathcal{R}f)$.
3. These one-dimensional Fourier transforms can be collected to form a two-dimensional function $F(k_x, k_y)$.
4. The inverse two-dimensional Fourier transform is used on this function to recover $f(x, y)$.

The steps outlined above are shown in [Figure 3.3](#). A sketch of the proof for the Fourier slice theorem is shown in the following ([Feeman, 2015d](#)). We start with the two-dimensional Fourier transform of our object $f(x, y)$

$$F(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-2\pi i(xk_x + yk_y)} dx dy \quad (3.2.1)$$

and the Radon transform of it

$$p_{\theta}(s) = [\mathcal{R}f]_{\theta}(s) = \int_{\mathbf{L}_{\theta, s}} f(x, y) d\ell \quad (3.2.2)$$

We apply the one-dimensional Fourier transform on $[\mathcal{R}f]_{\theta}(s)$ which yields

$$P(k_s) = \int_{-\infty}^{\infty} p_{\theta}(s) e^{-2\pi i s \psi} ds = \int_{-\infty}^{\infty} \int_{\mathbf{L}_{\theta, s}} f(x, y) e^{-2\pi i s \psi} d\ell ds \quad (3.2.3)$$

We now insert the definition of s from [\(3.1.2\)](#) and get

$$P(k_s) = \int_{-\infty}^{\infty} \int_{\mathbf{L}_{\theta, s}} f(x, y) e^{-2\pi i \psi(x \cos \theta + y \sin \theta)} d\ell ds \quad (3.2.4)$$

We can substitute the variable we integrate with respect to, to obtain

$$P(k_s) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-2\pi i \psi(x \cos \theta + y \sin \theta)} dx dy \quad (3.2.5)$$

This expression is very reminiscent of the two-dimensional Fourier transform shown in [\(3.2.1\)](#). In fact they are equal in the case where

$$P(k_s) = F(\psi \cos \theta, \psi \sin \theta) \quad (3.2.6)$$

This means that a straight line through the two-dimensional Fourier transforms rotated at an angle θ is equivalent to the one-dimensional Fourier transform of the Radon transform of $f(x, y)$ at an angle θ .

Fourier reconstruction provides a method to reconstruct projection data obtained during X-ray tomography; however, there are some practical problems with using this method. In practice, it is not possible to sample the two-dimensional Fourier

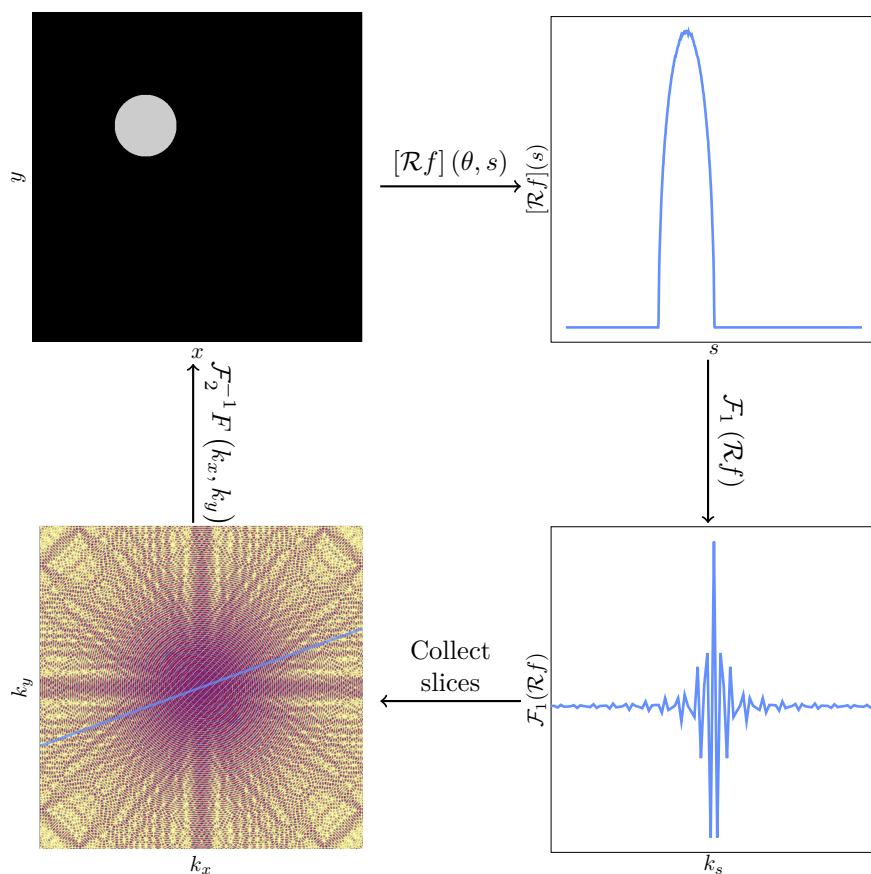


Figure 3.3: A schematic of the Fourier Reconstruction method. It starts with the same disk shown in Figure 3.1 in the top left corner, which is Radon transformed to the line in the top right corner. A one-dimensional Fourier transform is applied to this line to produce the spectrum seen below. The one-dimensional Fourier transforms is performed for all values of θ . Each Fourier transform corresponds to a line in the two-dimensional Fourier transform of $f(x, y)$ shown in the bottom left figure. The inverse two-dimensional Fourier transform can be applied to recover $f(x, y)$.

domain fully as X-ray detectors have a finite number of detector elements, and step size between angles will also be finite. Additionally, the Radon transform is done radially, which means the centre of the Fourier space is sampled more densely. The inverse two-dimensional Fourier transform requires that sample points are equidistant and it is, therefore, necessary to regrid the sample points. Regridding points in the Fourier space is difficult to do accurately; however, it was achieved by the `gridrec` algorithm (Marone and Stampanoni, 2012).

3.3 Analytical Reconstruction

The previous section detailed the Radon transform and provided a method to invert this transform by using Fourier reconstruction. The following section will describe the most commonly used reconstruction algorithm, which is the filtered backprojection (FBP) algorithm. However, before that, an explanation of the backprojection operation will be provided.

3.3.1 Backprojection

A naive way to approach the reconstruction of $f(x, y)$ from the projection function $p_\theta(s)$ would be to simply backproject the profiles obtained from the Radon transform back along the path the integral is performed along. The backprojection operation is mathematically equal to

$$g(x, y) = \int_0^\pi p_\theta(x \cos \theta + y \sin \theta) d\theta = \mathcal{B}[p_\theta](x, y) \quad (3.3.1)$$

However, this operation will result in a function where all points in the plane xy are positive since $g(x, y)$ cannot receive negative values from $p_\theta(x \cos \theta + y \sin \theta)$ as it is either 0 or positive. The backprojection operation when used on the sinogram in Figure 3.2 is shown in Figure 3.4. N_θ shows number of angles between 0 and π that are backprojected.




Figure 3.4: Example of the backprojection operation when applied to the sinogram of a disk. N_θ denotes the amount of equally spaced values of θ in the interval $[0, \pi]$.

Initially, the backprojection does not resemble a disk, but as more and more angles are

added we start to recover something akin to the original disk. However, it seems blurry compared to the original. The blurring stems from the fact that the backprojection operation smears $p_\theta(s)$ over the plane (Feeman, 2015a), or alternatively it could be interpreted as smoothing the function $f(x, y)$ with the kernel $\frac{1}{\pi\sqrt{x^2+y^2}}$ (Feeman, 2015b). We will see in the following section how this problem can be solved by filtering $p_\theta(s)$ before backprojecting it.

3.3.2 Filtered Backprojection

To derive the filtered backprojection we turn to the inverse two-dimensional Fourier transform

$$f(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(k_x, k_y) e^{2\pi i(k_x x + k_y y)} dk_x dk_y \quad (3.3.2)$$

which is transformed into polar coordinates

$$f(x, y) = \int_0^{2\pi} \int_0^{\infty} F(\psi \cos \theta, \psi \sin \theta) e^{2\pi i\psi(x \cos \theta + y \sin \theta)} \psi d\psi d\theta \quad (3.3.3)$$

When the coordinate system changes from Cartesian to polar, the surface elements $dk_x dk_y$ used in the integral change from squares to wedges. The shape of these wedges depends on the distance to the origin, which is why obtain the surface element $\psi d\psi d\theta$. The symmetry of the Fourier transform can be used to reformulate the limits of the integral (Feeman, 2015d). We now get

$$f(x, y) = \int_0^\pi \int_{-\infty}^{\infty} F(\psi \cos \theta, \psi \sin \theta) e^{2\pi i\psi(x \cos \theta + y \sin \theta)} |\psi| d\psi d\theta \quad (3.3.4)$$

We can use the Fourier slice theorem now, which gives us

$$\begin{aligned} f(x, y) &= \int_0^\pi \int_{-\infty}^{\infty} P_\theta(\psi) e^{2\pi i\psi s} |\psi| d\psi d\theta \\ f(x, y) &= \int_0^\pi h_\theta(s) d\theta \text{ where } h_\theta = \int_{-\infty}^{\infty} P_\theta(\psi) e^{2\pi i\psi s} |\psi| d\psi \end{aligned} \quad (3.3.5)$$

We recognise the outer integral as the backprojection operator from (3.3.1). Inside the integral, we have h_θ , which is the inverse Fourier transform of $P_\theta(s)$ filtered with $|\psi|$. The filter $|\psi|$ is called the ramp filter and acts as a high-pass filter on the Fourier transform of $p_\theta(s)$.

$$f(x, y) = \mathcal{BF}_1^{-1} |\psi| \mathcal{F}_1 [p_\theta(s)] \quad (3.3.6)$$

Figure 3.5 shows the same example as in Figure 3.4 this time reconstructed with the filtered backprojection rather than the backprojection algorithm.




Figure 3.5: Example of the filtered backprojection operation when applied to the sinogram of a disk. N_θ denotes the amount of equally spaced values of θ in the interval $[0, \pi]$.

Unlike the backprojection operator, this method correctly reconstructs the sharp edges of the disk. However, the number of projections N_θ drastically affects the appearance of the reconstruction. The disk becomes visible at relatively few projections, but there is a very noticeable pattern (streak artefacts) in the area surrounding the disk, which is a result of the few projections used for the reconstruction.

Another problem with the filtered backprojection algorithm is that the filtering with ψ amplifies noise in the projection data as noise is usually high frequency in the Fourier space. It is, therefore, necessary to use other filters than the ramp filter (Feeman, 2015b). An example of an alternative to the ramp filter is the *Ram-lak* filter, where frequencies above a threshold is set to zero.

3.4 Algebraic Reconstruction

This section will describe another approach used for tomographic reconstruction, namely, the algebraic reconstruction algorithms. These methods approach the problem of reconstruction as a linear system of equations that have to be solved. Because of this, linear algebra is the natural way to formulate the problem. This also means all operations are explicitly discretised compared to the continuous formulation we saw in the previous section.

3.4.1 Tomography as an Inverse Problem

The problem of recovering an imaged object \mathbf{x} based on a set of projections \mathbf{b} is viewed as solving the system of equations

$$\mathbf{b} = \mathbf{A}\mathbf{x} \quad (3.4.1)$$

where $\mathbf{b} \in \mathbb{R}^m$ is the projections or Radon transforms of $\mathbf{x} \in \mathbb{R}^n$, which is the discretised version of $f(x, y)$. The Radon transform \mathcal{R} is now represented with $\mathbf{A} \in \mathbb{R}^{m \times n}$ and is referred to as the system matrix or forward projection operator. An element a_{ij} represents how much an element x_j in \mathbf{x} contribute to an element b_i in \mathbf{b} . Based on Figure 3.6 we see that many if not most of the elements of \mathbf{A} are equal to

zero as most values in \mathbf{x} are empty and do not contribute to \mathbf{b} which means that \mathbf{A} is usually very sparse. However, increases in the number of pixels of modern X-ray detectors and the number of projections that are gathered means that both n and m are large, which means that \mathbf{A} is a very large matrix. The size of \mathbf{A} means that the entire matrix cannot be stored in the memory of a computer (Aarle, Palenstijn, Cant, et al., 2016).




Figure 3.6: The discretised version of the previously shown disk.

3.4.2 Algebraic Reconstruction Algorithms

There is a wide range of algebraic reconstruction algorithms and a subset of these will be described in the following sections.

3.4.2.1 Kaczmarz's Method

Kaczmarz's method was first put forth by Stefan Kaczmarz in 1937 but it was rediscovered by Gordon, Bender & Herman in 1970 in the context of tomography, where it was named Algebraic Reconstruction Technique (ART) (Kaczmarz, 1937; Gordon, Bender, and Herman, 1970). ART is historically significant in that it was the very first reconstruction algorithm used for a CT scan (Ambrose, 1973). This method consists of solving one of the equations in (3.4.1) per iterations of the algorithm by projecting the current estimate of \mathbf{x} onto it. A single equation of (3.4.1) is given by

$$b_i = \mathbf{a}_i \mathbf{x} \quad (3.4.2)$$

where b_i is the i th value in the projection data and \mathbf{a}_i is the i th row in the projection data. Projecting a point \mathbf{z} onto the affine hyperplane defined by (3.4.2) is done using

$$\mathcal{P}(\mathbf{z}) = \mathbf{z} + \frac{b_i - \mathbf{a}_i \cdot \mathbf{z}}{\|\mathbf{a}_i\|_2^2} \mathbf{a}_i. \quad (3.4.3)$$

The solution to (3.4.1) i.e $\mathbf{A}^{-1}\mathbf{b}$ is where all hyperplanes intersect. This point can eventually be reached (given that it exists) by continuously projecting \mathbf{z} onto a new hyperplane. With this we arrive at Kaczmarz's Method shown in Algorithm 1.

Algorithm 1: Pseudo-code of Kaczmarz's method.


Input: $x_{\text{initial}}, N_{\text{itr}}$

```

1  $x_0 \leftarrow x_{\text{initial}}$ 
2 for  $k$  from 1 to  $N_{\text{itr}}$  do
3    $i \leftarrow k \bmod m$ 
4    $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \frac{b_i - \mathbf{a}_i \cdot \mathbf{x}^{(k)}}{\|\mathbf{a}_i\|_2^2} \mathbf{a}_i$ 
5 end
6 return  $x$ 

```

We see that only one row is used at a time to update the solution, and that N_{itr} controls how many rows in \mathbf{b} that are used in the reconstruction. The modulo operation in line 2 ensures that we return to the first row after the m th row is reached. Kaczmarz's Method has the benefit that it converges quickly and the convergence speed can be improved by updating with rows randomly instead of going through them sequentially (Herman and Meyer, 1993). The second term in the projection operation is often multiplied with a relaxation parameter $\lambda^{(k)}$, which can also aid in faster convergence.



(a) $N_{\text{itr}} = 0.25m$ (b) $N_{\text{itr}} = 0.5m$ (c) $N_{\text{itr}} = 4m$ (d) $N_{\text{itr}} = 8m$ (e) $N_{\text{itr}} = 32m$

Figure 3.7: Example of Kaczmarz's method when applied to the sinogram of a disk. m in the figure text refers to the total number of rows in the system matrix \mathbf{A} .

In Figure 3.7 we see an example of the convergence of Kaczmarz's method when attempting to reconstruct the disk. Noticeably, we see that the outline of the disk is not visible in Figure 3.7a and Figure 3.7b. We have not iterated through all acquisition angles for these reconstructions, which means we have no information about

this region in the reconstruction. The outline of the disk is complete in Figure 3.7c where we have iterated through all the rows of the system matrix 4 times, however, the algorithm has not converged yet at this point. It is still not quite converged after 32 iterations through the rows of the system matrix in Figure 3.7e which shows just how slow the algebraic reconstruction methods are compared to FBP where only a single backprojection operation is required.

Kaczmarz's method is rarely used in practice for tomographic reconstruction as it is not robust against noise. Adding noise to the projection data effectively means that the system of equations in (3.4.1) is inconsistent. The hyperplanes of these equations do not intersect, which means the system of equations does not have a solution. This lack of robustness can be addressed by using all the projection data simultaneously, which will be discussed in the next section.

3.4.2.2 Simultaneous Algorithms

Besides the "row action" methods such as Kaczmarz's method, there exists a range of reconstruction algorithms that use all the rows of the system matrix \mathbf{A} to update the solution to the inverse problem. The simplest of these methods is Cimmino's method, and it is very similar to Kaczmarz's method except that the update step uses an average of $\mathcal{P}(\mathbf{z})$ for $i = 1$ to $i = m$ (Cimmino, 1938).

$$\mathbf{x}^{(k+1)} = \frac{1}{m} \sum_{i=1}^m \mathcal{P}(\mathbf{x}^{(k)}) \quad (3.4.4)$$

Expanding this expression we get

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \frac{1}{m} \sum_{i=1}^m \frac{b_i - \mathbf{a}_i \cdot \mathbf{x}^{(k)}}{\|\mathbf{a}_i\|_2^2} \mathbf{a}_i \quad (3.4.5)$$

The sum can be expressed as a multiplication of vectors and matrices

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \frac{1}{m} \sum_{i=1}^m \frac{\mathbf{a}_i}{\|\mathbf{a}_i\|_2^2} (b_i - \mathbf{a}_i \cdot \mathbf{x}^{(k)}) \quad (3.4.6)$$

$$= \mathbf{x}^{(k)} + \mathbf{A}^T \mathbf{M}^{-1} (\mathbf{b} - \mathbf{A} \mathbf{x}) \quad \text{where } \mathbf{M} = \text{diag} (m \|\mathbf{a}_i\|_2^2) \quad (3.4.7)$$

Cimmino's method is a part of a large variety of simultaneous algorithms which have the general form

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \lambda^{(k)} \mathbf{D}^{-1} \mathbf{A}^T \mathbf{M}^{-1} (\mathbf{b} - \mathbf{A} \mathbf{x}) \quad (3.4.8)$$

which differ in their definitions of \mathbf{D} and \mathbf{M} , both of which are diagonal matrices.

We have used the Simultaneous Iterative Reconstruction Technique (SIRT) throughout this project, which is specific version of the reconstruction algorithms shown in (3.4.8). SIRT is characterised by having \mathbf{D} and \mathbf{M} equal to the column and row

sums of \mathbf{A} i.e. $\mathbf{D} = \text{diag}(\|\mathbf{a}_j\|_1)$ and $\mathbf{M} = \text{diag}(\|\mathbf{a}_i\|_1)$ (Gilbert, 1972; Gregor and Benson, 2008). Note that SIRT refers to any reconstruction algorithm with the form seen in (3.4.8) in some research fields, but in our case, it refers to a specific algorithm.




Figure 3.8: Example of SIRT when applied to the sinogram of a disk.

Figure 3.8 shows how SIRT performs when it is used to reconstruct the disk phantom. The convergence of the algorithm is quick compared to ART, based on a comparison between Figures 3.7e and 3.8d. All projections are used 32 times to produce these reconstructions, which makes them comparable. This also means that the forward and backward projection operator has been used 32 times, which is much more computationally demanding than the FBP reconstruction in Figure 3.5e that is performed with only one backprojection. However, the major benefit of algebraic reconstruction algorithms is that prior knowledge can be incorporated, which can regularise and constrain the solution.

3.4.2.3 Prior Knowledge

Prior knowledge is commonly used with algebraic reconstruction algorithms to obtain a higher quality reconstruction. Note that the word “prior” is not as strictly defined as in Bayesian statistics but refers to any knowledge known in advance that can be implemented into the reconstruction algorithm to improve it.

An example of prior knowledge is the non-negativity constraint. The image obtained in X-ray tomography is a measure of the attenuation coefficients of the sample. These values can never be negative; however, noise in the projection data can sometimes give rise to negative values in the reconstruction. These can be removed by introducing a projection operator, that projects the solution onto a set that only contains positive numbers. Including the projection operator in (3.4.8) gives us

$$\mathbf{x}^{(k+1)} = \mathcal{P}_{\mathcal{C}} \left(\mathbf{x}^{(k)} + \lambda^{(k)} \mathbf{D}^{-1} \mathbf{A}^T \mathbf{M}^{-1} (\mathbf{b} - \mathbf{A}\mathbf{x}) \right) \quad (3.4.9)$$

where $\mathcal{P}_{\mathcal{C}}$ is the projection operator that, in the case of non-negativity, projects the solution onto the set $\mathcal{C} = \mathbb{R}_+^n$ i.e. the set that contains all real positive numbers. Another way to show this is with

$$0 \leq x_i \text{ with } i = 1, 2, \dots, n \quad (3.4.10)$$

An upper limit can be introduced if the maximum attenuation coefficient is known in the sample, which gives us the constraint

$$0 \leq x_i \leq \mu_{i,\max} \text{ with } i = 1, 2, \dots, n \quad (3.4.11)$$

where $\mu_{i,\max}$ is the largest possible attenuation coefficient for that voxel in the reconstruction. This type of constraint was used extensively in [Contribution A](#). If the largest attenuation coefficient in the entire sample is used μ_{\max} rather than the maximum value in each voxel we would have the constraint known as box-constraints in the CT community.

Modern algebraic reconstruction algorithms generally use some form of prior knowledge to improve image quality. The next section will describe the principles behind modern reconstruction algorithms and their use cases.

3.5 Recent Developments in Reconstruction Algorithms

Algebraic reconstruction methods are, as previously mentioned, computationally very demanding and are therefore less frequently used than analytical reconstruction methods (Beister, Kolditz, and Kalender, 2012). However, interest in these methods has risen as computation power has increased several orders of magnitude since the first CT scan. This section highlights some of the more recent developments within the field.

3.5.1 Variational Regularisation

In Lin et al. (2018) multiple methods are combined to create an algorithm that can reconstruct undersampled and noisy projection data. The authors regularise the algorithm with a penalty term that discourages non-smooth solutions. The penalty term consists of the sum of the Huber function applied to the sum of the difference between a voxel and its first order neighbours. The differences of the voxels are weighed with the distance between them. This is fairly similar to total variation regularisation; however, the Huber function is better at preserving edges than the 2-norm. The elements of the weighing matrix \mathbf{M}^{-1} (see (3.4.8)) are given by $m_i = \exp(-b_i)$ where b_i is the value of the i th ray. This means rays are weighed according to the strength of the signal, which reduces noise in the reconstruction. Both Nesterov's momentum and ordered subsets are used to speed up the convergence of the algorithm.

3.5.2 Discretisation of the Reconstruction

Regularising such that reconstructions appear smooth is commonly used and can be very effective if the sample being imaged has clearly defined borders. In the case

of Lin et al. (2018) the algorithm was developed to image multiphase flow in porous media where all phases can be distinguished, which makes the algorithm perform very well.

Tailoring the reconstruction algorithm with the sample constituents in mind is a powerful way to constrain and it can provide nearly noise-free reconstructions. This is done in Batenburg and Sijbers (2011), Zhuge, Palenstijn, and Batenburg (2016), and Six, De Beenhouwer, and Sijbers (2019). The discrete algebraic reconstruction technique (DART) was developed in Batenburg and Sijbers (2011). As the name suggests it makes use of discretisation to drastically lower the number of equations in (3.4.1). The algorithm is initialised with a reconstruction performed with a continuous algebraic reconstruction algorithm. This reconstruction is discretised via thresholding. The borders in the thresholded image are identified and a reconstruction is done where only border pixels are updated. The border reconstruction is merged with the previous reconstruction and a new threshold is performed.

DART was modified in Zhuge, Palenstijn, and Batenburg (2016) where the total variation regularized discrete algebraic reconstruction technique (TVR-DART) is presented. The method was developed to increase the robustness of DART and reduce user input. The reconstruction is regularised similarly to Lin et al. (2018) except that a soft thresholding function is applied to the reconstruction before the difference is calculated. The use of a soft threshold and the Huber norm ensures that the function is differentiable and nearly twice differentiable, which makes optimisation possible. This, along with an automatic estimation of the thresholds and discrete values in the reconstruction makes it much simpler from a user perspective as few parameters require user input.

Another modification of DART is found in Six, De Beenhouwer, and Sijbers (2019) where the authors have created the polychromatic-DART (poly-DART) algorithm. Poly-DART is very similar to the original DART algorithm except that the forward projection operator is replaced such that it accounts for the poly-chromaticity of most X-ray sources. Accounting for the full X-ray spectrum of the source reduces beam hardening artefacts, which increases the robustness of the DART algorithm.

Discretisation has also been used to improve reconstructions of dynamic tomography, which can be seen in Van Eynndhoven, Batenburg, and Sijbers (2014) and Van Eynndhoven, Batenburg, Kazantsev, et al. (2015). The region-based SIRT algorithm with intermediate piecewise constant function estimation (rSIRT-PWC) is presented in Van Eynndhoven, Batenburg, Kazantsev, et al. (2015) and is designed with flow experiments in mind. The algorithm identifies stationary regions based on a segmentation done on the object before fluid flow and fixes these voxels. The remaining pixels are characterised as either fully or partially dynamic and are treated as PWC functions that change value according to time i.e. a dynamic voxel might initially be empty and later be filled with fluid and this change can be characterised by a PWC function. The algorithm provides excellent results even with a very low amount of projections.

Reconstruction algorithms designed with fluid flow in mind is fairly common and Myers, Kingston, et al. (2011) and Myers, Geleta, et al. (2015) are examples of this.

Myers, Kingston, et al. (2011) is based on SIRT but it uses prior information in multiple ways to regularise it. The static projection data is subtracted from the dynamic projection data, a soft thresholding function is used to ensure that changes between timeframes only occur at the fluid front, the static parts of the reconstruction are forced to remain unchanged and finally, a segmentation operation is used to binarise the reconstruction. Myers, Geleta, et al. (2015) is an extension of this method, however, it is expressed as a statistical model where everything is modelled as conditional probabilities.

3.5.3 Temporal Regularisation

Designing reconstruction algorithms specifically for dynamic phenomena has become increasingly common and Kazantsev, Thompson, et al. (2015), Kazantsev, Van Eynhoven, et al. (2015), Kazantsev, Guo, et al. (2016), and Nikitin et al. (2019) are examples of this. In Kazantsev, Thompson, et al. (2015), Kazantsev, Van Eynhoven, et al. (2015), and Kazantsev, Guo, et al. (2016) a spatio-temporal regularisation term is introduced, which substantially increases image quality. Nonlocal means is used instead of the finite difference seen in Zhuge, Palenstijn, and Batenburg (2016) and Lin et al. (2018) as it avoids the piecewise constant appearance of total variation regularisation. The method uses all temporal information available to regularise instead of only using adjacent time frames. Using all temporal could lead to motion artefacts in the dynamic regions of a reconstruction and the authors, therefore, calculate the similarity of regions to avoid smoothing dynamic regions.

Nikitin et al. (2019) also utilises all temporal information to regularise the solution, however, unlike most dynamic reconstruction algorithms they do not assume that the data is quasi-static. Instead, it is assumed that the reconstruction can be described as a linear combination of basis functions. The algorithm is regularised with the 1-norm of the derivative of the reconstruction. This derivative also includes the temporal dimension and a parameter is introduced to control the weight of this term. The method performs much better than the traditional FBP and is especially improved by the spatio-temporal regularisation.

3.5.4 Deep Learning for Computed Tomography

Deep learning has become a major area of research within computer science in the last 10 years and it has recently been introduced to tomography. It has seen use both as reconstruction algorithms (X. Yang et al., 2020; Huang et al., 2020; Würfl et al., 2016) and as a means to improve algorithms obtained with more traditional methods (Zhi et al., 2021; Hendriksen, Pelt, and Batenburg, 2020; Hendriksen, Bühner, et al., 2021; Kang, Min, and Ye, 2017; Bühner et al., 2021). Both X. Yang et al. (2020) and Huang et al. (2020) provide algorithms designed to work for limited angle tomography, where the movement of the rotation stage is restricted. In general, neural networks require many examples during their training, which often consist of comparing good

and bad examples. Obtaining both high- and low-quality reconstructions of the same data, especially if the experiment is dynamic or performed on a synchrotron, where beamtime is scarce. X. Yang et al. (2020) therefore create a self-supervised network that trains on the data it is reconstructing. Huang et al. (2020) take another approach in that they use the U-net neural network, which trains quickly and performs the training with synthetic data.

In Hendriksen, Pelt, and Batenburg (2020) and Hendriksen, Bührer, et al. (2021) the authors developed the neural network Noise2Inverse, which denoises the reconstructions of noisy data. The network is self-supervised, which means it can train on the data it is attempting to denoise, eliminating the need for special training data. The network performs very well at a variety of tasks compared to other denoisers. The performance of these early efforts begs the question is denoising a subpar reconstruction is better than reconstructing with a regularised reconstruction algorithm.

3.5.5 Our Contribution to this Field

This chapter provided the mathematical background for X-ray computed tomography and showed some recent examples of modern tomographic reconstruction algorithms. Many of these algorithms are designed to improve robustness against noise and to prevent streak artefacts despite using a small number of projections and does so by using prior knowledge. The algorithms described in Section 3.5.2 achieves this robustness by using knowledge about the sample that is imaged to improve the reconstruction of it.

We used the same approach in Contribution A. Here we exploited that some dynamic experiments allow for a high-quality scan of the experiment. A segmentation of the high-quality scan is used to constrain the reconstruction. For instance, if we are performing a core flooding experiment, we know that the rock sample should remain unchanged, which means the attenuation values of the rock can be fixed. Myers, Kingston, et al. (2011) is similar in this regard, where only the dynamic part of the projection data is reconstructed and the static part of the reconstruction is fixed. A novel aspect of Contribution A is the use of initialisation to improve image quality. The reconstruction of the high-quality scan makes for an excellent starting point for a reconstruction algorithm (\mathbf{x}^k in (3.4.8)) and subsequent time steps can be initialised with the previous time step. Using box constraints and initialisation alone can significantly improve the quality of a reconstruction and does not require any modification of the reconstruction algorithm, which makes it very simple to implement.

CHAPTER 4

Practical Aspects of CT Experiments

The following section is meant to describe the practical aspects of working with a CT scanner. While Lambert-Beer's Law, the Radon transformation and the analytical and algebraic reconstruction algorithms provide the cornerstones to understanding a CT scanner there are still several experimental aspects of working with a CT scanner that has to be accounted for in practice. These are effects such as normalisation of projections, accounting for noise, dealing with ring artefacts etc. However, this section will start with a brief review of the development of the CT scanner.

4.1 X-ray Tomography Instruments

The first commercial CT scanner was developed for medical use by Sir Godfrey Hounsfield between 1967 and 1971, where it culminated in the first scan of a patient (Richmond, 2004). He was awarded the Nobel prize in medicine in 1979 that he shared with Allan Cormack, who had worked on developing a similar setup.

Hounsfield's first commercial machine, the EMI Mark I, consisted of only two detector elements that were translated and rotated to record two slices simultaneously (Hounsfield, 1973). The scanner proved to be an immediate success, with the first scan providing the location of a brain cyst (Richmond, 2004).

The fact that the detector had to be translated meant that the scanning procedure was slow, with a slice taking approximately five minutes to image. However, the value of the CT scanner as a diagnostic tool was immediately recognised by physicians, and the technology was rapidly improved. The number of detector elements was increased such that an entire slice could be acquired at once. This eliminated the need to translate detector elements which drastically improved scanning time. This was followed by a rapid push towards multi-slice scanners to increase axial coverage, which has several benefits in a clinical setting (Panetta, 2016).

Both clinical and industrial scanners moved to digital detector systems in the middle of the 1990s, replacing the previously used photographic plates. In industrial scanners, the flat-panel detector was adopted that gave rise to cone-beam CT scanners. Cone-beam CT scanners are by far the most common type of scanner when it comes

to research and material science, and the setup will be described in greater detail in Section 4.6.

Clinical scanners have generally moved towards very short acquisition times, with a full-body scan only taking a few seconds. The quick scan time also makes it possible to scan the heart of a patient without motion artefacts, and additionally, it improves the patient experience by eliminating their need to hold their breath.

This is in contrast to industrial scanners that have pushed for increasingly smaller spatial resolutions at the expense of scan time. A typical laboratory CT scanner with a resolution of approximately one μm will require a few hours to complete a scan (Maire and P. J. Withers, 2014).

4.2 Detection of X-ray Radiation

There are multiple ways of constructing an X-ray detector, and the choice of technology depends on the context of the experiment. We have used two CT scanners throughout this project; one with a flat-panel detector and one with a charge-coupled device (CCD) based detector. The first of these CT scanners is housed at DTU Offshore and is referred to as the “nanoCT”, see Appendix A for a description of its components. The second scanner is a ZEISS XRadia 410 Versa located at the 3D Imaging Center of DTU.

The light collected by both of these CT scanners is first converted from X-ray radiation to visible light with a scintillator, which is then read out by either an array of pixels or the CCDs. The pixels of the nanoCT are made of amorphous silicon (a-Si) photodiodes that convert the visible light to a digital signal, which is read out thin-film transistors (TFT) which is common for most flat-panel detectors (Kump et al., 1998). The pixels of CCD detectors are much smaller than those seen in flat-panel detectors, which means they offer a far better spatial resolution; however, the sensitive area is also much smaller than flat-panel detectors (Hoheisel, 2006; Gruner, Tate, and Eikenberry, 2002).

A drawing of a flat-panel detector can be seen in Figure 4.1. A scintillator works by absorbing the X-ray radiation and then re-emitting the light in the visible range that can then be imaged. The re-emitted light is not necessarily emitted in the same direction as the incident X-ray that effectively smears the position of the incoming radiation. This is indicated in Figure 4.1 by the line spread function.

There are two key features that dictate the performance of a scintillator which is the detective quantum efficiency (DQE) and the time shape of its response. The DQE is defined as the squared ratio of the SNR of an actual detector and an ideal detector (Jones, 1959)

$$\text{DQE} = \frac{(S/N)_{\text{real}}^2}{(S/N)_{\text{ideal}}^2}, \quad (4.2.1)$$

where S is the signal, N is the noise. Noise will still be present in an ideal detector due to the finite photon count. A large DQE is, in general, a very desirable feature for




Figure 4.1: A drawing of a typical flat-panel detector. The figure is obtained from Beevil (2017) and has been cropped.


an X-ray detector. However, it cannot come at the expense of a slow time response. The time shape of the luminescence of a scintillator can be modelled as a double exponential decay (Leo, 1994):

$$N(t) = A \exp\left(\frac{-t}{\tau_f}\right) + B \exp\left(\frac{-t}{\tau_s}\right). \quad (4.2.2)$$

$N(t)$ is the number of photons emitted as a function of the time t , τ_f and τ_s are the two decay constants. τ_f is much smaller than τ_s . The primary or fast decay time for CsI (caesium iodide) use is three orders of magnitude faster than its secondary decay time. CsI is a commonly used material for industrial scanners; however, the afterglow is too long to be used in modern clinical scanners. These scanners are often equipped with scintillators that have a fast response time and a small afterglow (Panetta, 2016).

As previously mentioned, a-Si photodiodes digitise the signal that is read out by TFTs or CCDs. This type of detector is an energy integrating detector which means the spectrum of the incident X-ray radiation is lost. This is not the case for photon-counting detectors, which are currently being developed. Photon-counting detectors measure X-ray radiation directly instead of using the conversion step with the scintillator, which makes it possible to measure the X-ray spectrum. This is of great interest in medicine as this increases soft-tissue contrast, improve spatial resolution and allows for a reduction in exposure (Willemink et al., 2018).

An example of two different detectors is shown in Figure 4.2. In Figure 4.2a, we see a bright field (also known as a clear field or flat field) collected from the nanoCT scanner at DTU Offshore. This scanner has a detector with a CsI scintillator, and it uses a-Si photodiodes to detect the signal from the scintillator. One can easily see the structure of the detector due to variations in the efficiency of the detector elements. This detector also has several faulty pixels that require correction before radiographs collected with this scanner can be reconstructed.



(a) Bright field from the nanoCT at DTU Offshore.

(b) Bright field from the ZEISS XRadia 410 Versa at the 3D Imaging Center.

Figure 4.2: Bright fields collected with different scanners at DTU.

In Figure 4.2b, there is an example of a bright field from the ZEISS XRadia 410 Versa at the 3D Imaging Center at DTU. The readout is done using CCDs rather than a-Si photodiodes in this detector, and its response is noticeable more uniform compared to the nanoCT.

X-ray detectors often have a background signal in the absence of an X-ray source. This background signal is called a dark field, and it is collected to ensure proper normalisation of the projection data gathered on the scanner. The dark field of the nanoCT at DTU Offshore is shown in Figure 4.3. Here we also see that this particular detector has a fairly inhomogeneous response. The ZEISS XRadia 410 Versa has a negligible dark field, and it is therefore not recorded during normal operation.

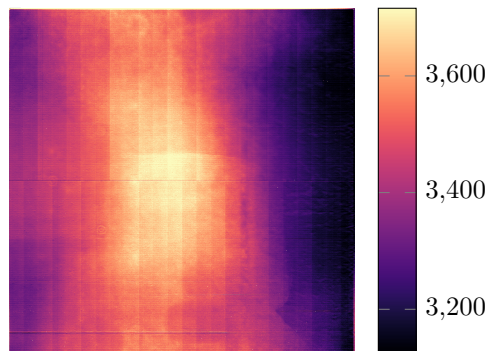


Figure 4.3: Dark field from the nanoCT at DTU Offshore.

4.3 From Raw Data to Sinograms

The projection data gathered from a CT scanner has to be converted from the raw output of the detector to something reminiscent of the Radon transform of the object as this is the basis for tomographic reconstruction. The Radon transform seen in (3.1.3) expressed with cartesian coordinates is given by

$$[\mathcal{R}f](\theta, s) = \int_{L_{\theta, s}} f(x, y) d\ell$$

Lambert-Beer's law seen in (2.2.7) can also be expressed in two dimensions in the same way the Radon transform is expressed

$$\begin{aligned} I(\theta, s) &= I_0 \exp \left[- \int_{L_{\theta, s}} \mu(x, y) d\ell \right] \\ -\log \left(\frac{I(\theta, s)}{I_0} \right) &= \int_{L_{\theta, s}} \mu(x, y) d\ell. \end{aligned} \quad (4.3.1)$$


From this we can identify that the function $f(x, y)$ which describes the object being imaged is equivalent to the 2D function of its attenuation values $\mu(x, y)$ with this we see the following

$$[\mathcal{R}f](\theta, s) = -\log \left(\frac{I(\theta, s)}{I_0} \right). \quad (4.3.2)$$

A CT scanner records $I(\theta, s)$ when operating, which should be divided with I_0 , which is the bright field described in the previous section. We also know that this is not the true initial intensity as the detector might have a non-zero dark field which should be subtracted from both the projection data and the bright field. With this, we arrive at

$$[\mathcal{R}f](\theta, s) = -\log \left(\frac{I(\theta, s) - D}{I_0 - D} \right), \quad (4.3.3)$$

where D is the previously mentioned dark field. The term inside the outer parentheses is the transmission and can take on values between zero and one. An example of the normalisation process is shown in Figure 4.4. Here a piece of chalk imaged at the NanoCT at DTU Offshore can be seen in its corrected and uncorrected state. After the correction, values in the radiograph have to be non-negative as the negative logarithm to a number between zero and one is between zero and infinity. However, noise in the radiographs, bright fields or dark fields can result in negative values in the transmission image. The logarithm of negative numbers is not defined, and faulty pixels, therefore, require correction before taking the negative logarithm.



(a) Raw data from the X-ray detector. (b) Data after equation (4.3.3) has been applied.

Figure 4.4: A raw radiograph and corrected radiograph of a piece of chalk from the Kraka oil field.

4.4 Noise

Noise in the projection data can be detrimental to the quality of a tomographic reconstruction. Noise in the projection data primarily stems from photon noise (also known as shot noise or quantum noise) and electrical noise from the detector (Whiting et al., 2006).

Photon noise is due to the finite number of photons that are detected, which results in Poissonian noise. However, an X-ray tube produces a spectrum of X-ray photons rather than a single wavelength as explained in Section 2.1. Additionally, X-ray detectors in CT scanners do not have a uniform detection efficiency as a function of energy (at least for energy integrating detectors). The combination of these effects means each energy level will give rise to a different Poissonian distribution which means the probability distribution function that best describes the photon noise in a CT scanner is a Compound Poisson distribution (Whiting et al., 2006). While photon noise cannot be described as a pure Poisson distribution, it will still be proportional to the total number of photons that are detected. This means larger samples and samples with large attenuation values will, in general, be noisier. This can be mitigated to some extent by increasing the source voltage. An increase in the source voltage will increase the mean energy of the emitted X-ray radiation, which will improve their ability to penetrate as we saw in Figure 2.4 (Rodríguez-Sánchez et al., 2020). The dependency on sample depth also means that the centre of a sample has more noise than the edge. This is usually not a problem, however, the difference between the noise level at the edge of the sample and the centre can be noticeable during dynamic imaging where the exposure time is usually decreased.

Electrical noise arises right before the electric signal from the TFTs or CCDs is digitised. The analogue signal is susceptible to noise and it usually stems from

the connections between the CCD (for instance), electrical amplifier and the digital to analogue converter (Duan et al., 2013). Electrical noise is usually constant and independent of the settings on the CT scanner. It only becomes a concern when the photon count is small and the signal detected by the X-ray detector is comparable to the electronic noise level.

4.5 Ring Artefacts

Ring artefacts is a type of error that is fairly frequent in tomographic reconstructions. It is caused by defects in either the detector pixels or the scintillator material in front of them. They appear as concentric circles in reconstruction centred around the axis of rotation. The normalisation done in equation (4.3.3) corrects most of the variation in the detector response as seen in Figure 4.4b. Here the structure of the detector is divided out, leaving a uniform background. However, the detector response is nonlinear, which means it is not possible to use (4.3.3) to account for large differences.




Figure 4.5: The central slice of a piece of chalk imaged with the nanoCT at DTU Offshore. Notice the ring artefacts centred in the lower middle part of the figure.

An example of the appearance of ring artefacts can be seen in Figure 4.5 which shows the central cross-sectional slice of a piece of chalk. The centre of rotation is located approximately in the lower middle part of the image, where the concentric rings get increasingly smaller.


Multiple methods have been proposed in the literature to mitigate ring artefacts with fairly different approaches. Most methods either focus on pre-processing the projection data (Boin and Haibel, 2006; Vo, Atwood, and Drakopoulos, 2018) or post-processing the reconstruction (Sijbers and Postnov, 2004; Jha et al., 2014). Both approaches attempt to eliminate the ring artefacts by identifying problematic regions

and smoothing them out. Other methods attempt to account for the detector response either directly or indirectly. In Croton et al. (2019) the detector response is mapped and incorporated into the normalisation of the radiographs. The method requires that the source is easily moved relative to the detector, which is not the case for the nanoCT used for most of this project. The ZEISS XRadia 410 Versa used for [Contribution C](#) attempts to prevent ring artefacts during the acquisition of projection data by translating the position of the sample between each projection. This effectively averages out the response of each pixel preventing ring artefacts.

We have elected to use pre-processing of the projection data to mitigate ring artefacts in the reconstruction throughout this project. First, problematic pixels are identified in histograms of the bright fields and dark fields, which provide a library of pixels that requires correction before reconstruction.

Examples of histograms of the dark and bright fields can be seen in [Figure 4.6](#). Pixels within the 0.01% and 99.999% quantile are part of the blue area and the outliers are shown in yellow. The bins are very thin which makes them appear black. The pixels in the yellow area are the ones that cause ring artefacts. The ZEISS XRadia 410 Versa has a very uniform response which means it is not necessary to pre-process data gathered from this scanner based on [Figure 4.6c](#). This is not the case for the nanoCT which has several outliers that can be seen in both bright and dark fields [Figures 4.6a](#) and [4.6b](#).

We find the position of the pixels in the yellow area with a threshold based on the histogram. After this, we replace the value of problematic pixels with the median value of their neighbours. This has to be done on each projection separately. Using the same value for multiple projections would result in ring artefacts, and it should therefore be avoided.



(a) Histogram of a bright field from DTU Offshore nanoCT. (b) Histogram of a dark field from DTU Offshore nanoCT. (c) Histogram of a bright field from the ZEISS XRadia 410 Versa.

Figure 4.6: Histograms of the bright and dark fields from scanners used in this project.

4.6 Cone Beam Geometry

As previously mentioned, cone-beam geometry is the most commonly used geometry for laboratory CT scanners (Maire and P. J. Withers, 2014), and both the ZEISS XRadia 410 Versa and the nanoCT at DTU Offshore use this geometry.

The geometry consists of an X-ray source that emits a cone of X-ray radiation, hence the name, through the object towards a flat-panel detector. The sample is placed on a rotating gantry while the position of the X-ray source and the detector is fixed. A schematic of a cone-beam geometry can be seen in Figure 4.7.




Figure 4.7: A schematic of the geometric of a cone-beam CT scanner. Here S indicates the source position, D indicates the detector centre, and AoR is an abbreviation for the axis of rotation. A sample would be placed in the middle of this axis and rotated along the angle theta. The vectors \vec{u} and \vec{v} describe the detector plane. The angles ϕ , ψ and θ indicate the possible rotations of the detector plane.

The X-ray source position is defined by the point S , the axis of rotation (AoR), the object position O and the centre of the detector D with the distance between the source and detector $|SD| = C_l$ being referred to as the camera length. The spread of the cone-beam gives rise to geometric magnification, which can be calculated using

$$M = \frac{|SD|}{|SO|}. \quad (4.6.1)$$

The detector can be described as a plane with the vectors \vec{u} and \vec{v} and it can be rotated in three different ways indicated by the angles ϕ , ψ and θ . Note the axis of rotation can also be rotated along ψ and ϕ , and this would be indistinguishable from a rotation of the detector. Similarly, a positive shift along the z -axis of the source position S is indistinguishable from a negative shift of the detector D along the z -axis.

The cone-beam geometry differs from the parallel beam geometry, which is equivalent to the Radon transform described in Section 3.1. Rays emitted from the source in the xy -plane will still behave as if it was a parallel beam geometry, however, obtaining a “complete” set of projection makes it necessary to collect radiographs in the domain $\theta \in [0, 2\pi[$. The FBP algorithm cannot provide a stable solution for rays outside of the central plane, and a constrained algorithm should be used instead (Tuy, 1983).

The de-facto industry standard which approximates the FBP algorithm for a cone-beam geometry is the Feldkamp-Davis-Kress algorithm (Feldkamp, Davis, and Kress, 1984; Cnudde and M. Boone, 2013; Maire and P. J. Withers, 2014). The FDK algorithm is not perfect, and reconstruction artefacts in the form of blurring will appear when moving away from the xy -plane, see Figure 4.7, and this effect gets increasingly pronounced with larger cone angles i.e. a less parallel beam (Maire and P. J. Withers, 2014). Even with these problems, cone-beam geometry has several benefits that make it attractive for most cases. It is far more efficient to use a conical beam since fewer X-ray photons have to be collimated, and the geometric magnification from the conical beam increases the spatial resolution of the CT scanner without having to modify the detector. Additionally, the FDK algorithm is, like the FBP algorithm, computationally very efficient and easy to use.

4.7 Misalignment of Geometry and Artefacts


The geometry shown in Figure 4.7 represents an ideal setup where the line between in the source position S and detector centre D is in the xy -plane and is perpendicular to the detector plane. In practice, such a geometry is rarely realised, especially when non-commercial scanners are used, such as the nanoCT at DTU Offshore. This scanner does not record enough geometrical information to ensure the reconstruction of the projection data is satisfactory. A misaligned CT scanner will cause the elements in the system matrix \mathbf{A} to have incorrect values resulting in faulty forward and backward projection operators. These operators will result in different artefacts depending on which geometry parameter that is misaligned.

An overview of the different misalignment parameters and an approximate description of the artefacts they give rise to can be found in Table 4.1. Examples of the appearance of reconstruction artefacts due to an incorrect geometry can be seen in Figure 4.8 and Figure 4.9. These figures show the artefacts in a chalk sample and artefacts in a carbon steel flow cell where the inside is covered is partially covered with BaSO_4 .


In practice, multiple parameters usually require tuning, so the initial appearance of reconstruction might contain a mix of the artefacts shown. Parameters such as the camera length C_l and the horizontal source position S_u give rise to very similar artefacts, which can make it cumbersome to tune these parameters by hand, and an automatic method is therefore desirable.

Table 4.1: A table of the different misalignment parameters in a cone-beam CT scanner. Figures 4.8 and 4.9 provides examples of the appearance of reconstruction artefacts caused by different geometrical parameters.


Misalignment	Misalignment parameter	Equivalent parameter	Artefact Appearance
Horizontal shift of D	D_u	—	Uniform defocusing.
Vertical shift of D	D_v	$-S_v$	Artefacts without blurring.
Horizontal shift of S	S_u	—	Blurring away from the rotation centre. Doubles edges.
Vertical shift of S	S_v	$-D_v$	Artefacts without blurring.
$ SD $ or camera length C_i	—	—	Blurring that increases with distance to rotation centre. Doubles edges.
Rotation of AoR in ϕ	R_ϕ	D_ϕ	—
Rotation of AoR in ψ	R_ψ	D_ψ	Blurring away from the rotation centre. Doubles edges.
Rotation of detector in ϕ	D_ϕ	R_ϕ	—
Rotation of detector in ψ	D_ψ	R_ψ	Blurring away from the rotation centre. Doubles edges.
Rotation of detector in θ	D_θ	—	Uniform defocusing. Doubles edges. Points become stars.



(a) Detector centre shifted by 5 pixels in the xy -plane. (b) Detector centre shifted by 25 pixels along the AoR. (c) Source centre shifted by 144 pixels in the xy -plane.




(d) Camera length reduced by 10%. (e) Detector rotated by 1.5° in θ . (f) Detector rotated by 1.5° in ϕ .




(g) Detector rotated by 1.5° in ψ . (h) The reconstruction with optimal parameters.


Figure 4.8: Examples of reconstruction artefacts for different geometrical parameters in a piece of chalk. The chalk was imaged at DTU Offshore using the nanoCT.



(a) Detector centre shifted by 5 pixels in the xy -plane. (b) Detector centre shifted by 25 pixels along the AoR. (c) Source centre shifted by 144 pixels in the xy -plane.



(d) Camera length reduced by 10%. (e) Detector rotated by 1.5° in θ . (f) Detector rotated by 1.5° in ϕ .



(g) Detector rotated by 1.5° in ψ . (h) The reconstruction with optimal parameters.

Figure 4.9: Examples of reconstruction artefacts for different geometrical parameters in a carbon steel flow cell covered with BaSO_4 . The cell was imaged at DTU Offshore using the nanoCT.

4.8 Auto-focus

The previous section revealed that an incorrect geometry greatly diminishes the quality of reconstructions. It is, therefore, desirable to know the scanner geometry accurately. Automatically obtaining the scanner geometry has been well researched, and in general, there are two approaches to the problem.

The first approach involves imaging a specially made phantom with known dimensions. Based on either the projection data or reconstruction of it, it is possible to determine the geometry parameters of the scanner (Ferrucci et al., 2018; G. Li et al., 2018; Xu et al., 2017; K. Yang et al., 2006; Jacobson, M. D. Ketcha, et al., 2018; Xiang, Wang, and Cai, 2016; Noo et al., 2000; Jacobson, M. Ketcha, et al., 2017).

The second approach is to maximise the sharpness of the reconstruction. The logic behind this method is that a reconstruction that maximises image sharpness is performed with the correct geometry. There is no objective measure of image sharpness, so proxies of image sharpness are used instead (Kingston, Sakellariou, A. Sheppard, et al., 2010; Kingston, Sakellariou, Varslot, et al., 2011; Meng, Gong, and X. Yang, 2013; Panetta, Belcari, Del Guerra, Bartolomei, et al., 2012; Wicklein et al., 2012; Lesaint et al., 2017; Patel et al., 2008; Herbst et al., 2019; Panetta, Belcari, Del Guerra, and Moehrs, 2008; Muders and Hesser, 2014; Ouadah et al., 2016).

The phantom based methods are ideal when the geometry rarely changes, which could be in a clinical setting where the position of the source and detector is fixed. This is not the case in a research setting where the camera length is adjusted to obtain the desired geometrical magnification for the sample. Having to image a phantom each time the setup of the scanner changes before performing the actual imaging is not feasible in a research environment. We have therefore focused on developing an automatic method to obtain the correct geometrical parameters. Our approach is inspired by Kingston, Sakellariou, A. Sheppard, et al. (2010) and Kingston, Sakellariou, Varslot, et al. (2011), where image sharpness is defined as the norm of the gradient of the image

$$\mathcal{S}(I(x, y)) = \|\nabla I(x, y)\|_2$$


This is the sharpness of a single slice in a reconstruction; however, using the sum of multiple slices provides a more robust measure of the sharpness of the reconstruction. Additionally, some artefacts are more pronounced further away from the detector centre, such as the rotation in ψ . The total sharpness of a reconstruction is given by

$$\mathcal{S}_{\text{tot}}(I(x, y)) = \sum_z \|\nabla I(x, y)\|_2. \quad (4.8.1)$$

I have not examined in detail how many slices is sufficient to ensure a stable measure of sharpness; however, every 10th slice seems to provide a robust measure of sharpness. To find the optimal geometry, we use the minimiser `iminuit` (Dembinski et al., 2022; James and Roos, 1975). The negative sharpness $-\mathcal{S}_{\text{tot}}(I(x, y))$ is used as a cost function and all parameters are included in the minimisation. Each function call

requires a full reconstruction which can be time-consuming even when performed on GPUs, so it is usually desirable to downscale the projection data to speed up computation time.

We have not performed a systematic test of the auto-focus procedure yet, but preliminary results seem promising. An example of the performance of the method is presented in Figure 4.10. In the figure, there are two reconstructions of the same sample with different geometry parameters. The sample consists of a piece of aluminium interspersed with specks of tungsten. The tungsten attenuates X-ray radiation far more than the aluminium, which results in a large contrast between them. The sample was imaged with the nanoCT at DTU Offshore at the beginning of this project to test the performance of the scanner. The camera length was not measured precisely during the acquisition, which results in the artefacts seen in Figure 4.10a.



(a) Reconstruction performed with the initial parameters measured by hand. (b) Reconstruction performed with the parameters obtained with our auto-focus method.

Figure 4.10: Examples of reconstructions of a piece of aluminium with dots of tungsten.

The projection data was initially downsampled by a factor of four, after which the auto-focus code was run to obtain a rough estimate of the scanner geometry. The auto-focus algorithm was then used on the full data set with this rough geometry to get the final geometry. The final geometry results in the reconstruction seen in Figure 4.10b, which is much sharper and without artefacts.

4.9 Consistency of Projection Data

Performing dynamic tomographic experiments usually requires fairly elaborate experimental setups. Complicated experiments potentially have more error sources which

we learned during the project. This was especially apparent in the flow experiments we performed in this project.

In **Contribution C**, we wanted to examine the crystallisation fouling process or scale formation inside pipes. We performed in situ measurements to study the scale buildup by imaging the process with a CT scanner. Our experimental setup consisted of a carbon steel flow cell that was injected with water supersaturated with BaSO_4 . The injection and collection of the water were done by connecting both ends of the flow cell with plastic tubing. The tubing we used was fairly stiff, which resulted in unwanted movement of the sample as it rotated on the stage during acquisition. The movement was especially pronounced for the initial and final projections in a revolution where the tube was stretched the most. It was, therefore, necessary to develop multiple tools to quantify the movement of the sample, which also made it possible to correct this unwanted movement during the reconstruction.

A simple way of getting a qualitative overview of the projection data is “doubling” the sinogram and examining its appearance. Since the acquisition angles for cone-beam tomography are typically spread out equally between zero and 2π , one would expect a smooth transition from the final projection to the initial projection because zero and 2π should correspond to the same position.

An example of a “double”-sinogram can be seen in **Figure 4.11**. There is a sharp transition from the final projection to the first projection seen in the middle of the figure at the blue line. This shift means that the sample has moved relative to the detector during the acquisition. Additionally, it appears that the change is fairly smooth as there are no sharp transitions in the sinogram.





Figure 4.11: An example of a “double”-sinogram. The sinogram of a carbon steel flow cell is repeated such that the final projection is followed by the sinogram again. A line is drawn to show the point where the sinogram is repeated.

The transition region in **Figure 4.11** can be examined further by looking at the differences between the first, next to last and last projections. In **Figure 4.12a** the difference between the last and the second to last projection is displayed. The difference is almost uniform, but there are faint vertical lines from the edge of the flow cell. The angular shift between the last and second to last projection should be identical to the angular shift between the last and the first projection. The difference between these projections should therefore be equivalent to what is seen in **Figure 4.12a**; however,

this is not the case as shown in [Figure 4.12b](#). The difference is far greater, and it appears that the first projection is shifted to the left.



(a) The difference between the last and second to last projection. (b) The difference between the first and the last projection.

Figure 4.12: Projection differences for the carbon steel flow cell examined in [Contribution C](#).

The shift observed in [Figure 4.12b](#) can be quantified by examining the contour of the sinogram. The contour of any sinogram should behave exactly like a sine wave if the sample is moving as expected. By fitting the edge of a sinogram with a sine wave, it is possible to quantify the horizontal movement of the sample during acquisition.

An example of this for the previously mentioned cell can be seen in [Figure 4.13](#). The plot in the top box shows one of the edges of the sinogram fitted with the function

$$f(x) = A \sin(x + b) + B. \quad (4.9.1)$$

The fit error, defined as the difference between the actual edge and the fit, is displayed in the lower box. The edge of the sinogram appears approximately sinusoidal between $\pm 100^\circ$, but it deviates a lot from a sine wave beyond this which can also be seen in the fit error plot. A 5-pixel shift in the horizontal direction is a significant deviation from the nominal geometry as seen in [Figures 4.8a](#) and [4.9a](#), where the detector position is shifted with 5-pixels in horizontal direction. The artefacts we saw in the flow cell data were much less significant than what we see in [Figures 4.8a](#) and [4.9a](#) because only a subset of the projections are shifted.

The deviation of the sample in the horizontal plane is given by the error shown in [Figure 4.13](#), and it was used to mitigate the artefacts that arose from the movement of the sample. The correction was performed by shifting the position of the detector to cancel the deviation from zero given by the fit error. The effect of the correction can be seen in [Appendix B](#).

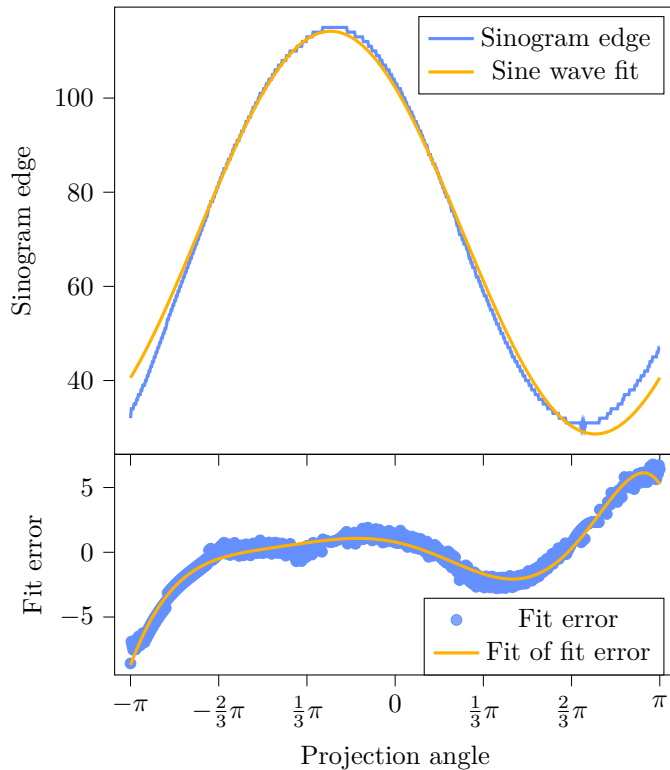


Figure 4.13: The upper plot shows the edge of a sinogram from Contribution C fitted with a sine wave (see (4.9.1)). The fit error is displayed below. The fit error has also been fitted.

4.10 Implementation Details

The choice of software is significant in tomographic reconstruction as it drastically affects computation time. The ASTRA toolbox has been used through this project as the forward and backward projection operators are implemented in CUDA which makes it possible to perform calculations on GPUs which substantially speeds up reconstruction speed (Palenstijn, Batenburg, and Sijbers, 2011; Aarle, Palenstijn, Beenhouwer, et al., 2015; Aarle, Palenstijn, Cant, et al., 2016). ASTRA also contains implementations of the most common reconstruction algorithms and scanner geometries. The software is called via a Python wrapper, which increases ease of use.

ASTRA is strictly designed for reconstruction and does not provide any tools to assist with the I/O of projection data or the preprocessing of it. It is therefore up to the users of ASTRA prepare the projection data before reconstruction. CT scanners often save a separate file for each projection and there is a significant overhead associated with reading multiple small files compared to reading a single large one.

I have, therefore, chosen to use the HDF5 data format as reading and writing to it is very fast and it supports very large files (The HDF Group, 1997; Collette et al., 2021). Additionally, it is possible to store metadata such as scanner parameters that are relevant to the reconstruction of the data.

The primary programming language during the project has been Python where NumPy and SciPy has used for most tasks (Harris et al., 2020; Virtanen et al., 2020). Computationally demanding tasks have performed on GPU using the Python package CuPy, which makes it possible to perform calculations on GPU (Okuta et al., 2017).

CHAPTER 5

Contribution A:

Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction

Contribution A represents my co-author and I's attempt to develop a reconstruction algorithm suited for dynamic tomography experiments that allow for a high-quality static scan of the experiment. The algorithm was designed such that it could be used in conjunction with the flow cell mentioned in [Section 1.3.2](#) to enable a wide range of dynamic experiments. We did not want a complicated regulariser but instead sought to leverage that most core flooding experiments are already constrained.

The original idea for the paper came from a poster I presented at American Geophysical Union's Fall Meeting in December 2019 in San Francisco, and the majority of the work was done in 2020.

The paper is included as published in its entirety below, and it can also be found online at <https://doi.org/10.1038/s41598-021-91776-1>.



OPEN

Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction

Peter Winkel Rasmussen¹, Henning Osholm Sørensen^{2,4}, Stefan Bruns^{3,4}, Anders BJORHOLM DAHL^{1,4} & Anders NYMARK CHRISTENSEN¹

Dynamic tomography has become an important technique to study fluid flow processes in porous media. The use of laboratory X-ray tomography instruments is, however, limited by their low X-ray brilliance. The prolonged exposure times, in turn, greatly limit temporal resolution. We have developed a tomographic reconstruction algorithm that maintains high image quality, despite reducing the exposure time and the number of projections significantly. Our approach, based on the Simultaneous Iterative Reconstruction Technique, mitigates the problem of few and noisy exposures by utilising a high-quality scan of the system before the dynamic process is started. We use the high-quality scan to initialise the first time step of the dynamic reconstruction. We further constrain regions of the dynamic reconstruction with a segmentation of the static system. We test the performance of the algorithm by reconstructing the dynamics of fluid separation in a multiphase system. The algorithm is compared quantitatively and qualitatively with several other reconstruction algorithms and we show that it can maintain high image quality using only a fraction of the normally required number of projections and with a substantially larger noise level. By robustly allowing fewer projections and shorter exposure, our algorithm enables the study of faster flow processes using laboratory tomography instrumentation but it can also be used to improve the reconstruction quality of dynamic synchrotron experiments.

For many years the primary technique to determine fluid flow properties of rocks was to perform classical core plug scale tests, where fluids, e.g. gases or liquids, were injected into natural porous media. The absolute permeability could then be established from Darcy's law¹. During the last 15 years, methods have been developed that estimate rock permeability by conducting computational fluid dynamics simulations of single or multiphase flow^{2–6}. These simulations are typically based on three-dimensional pore-scale models of the rocks obtained by X-ray tomography. In recent years in situ X-ray tomography has become one of the most popular methods to directly study dynamic processes in rocks^{7–9} such as fluid flow properties^{10–15} and reactive transport in rocks^{16–20}. To capture these phenomena in situ, X-ray tomography has to be performed at high spatial and temporal resolution. Therefore, most studies have been performed using synchrotron sources, which provide an extremely high X-ray beam brilliance, many magnitudes above laboratory X-ray sources²¹. Unfortunately, beamtime at synchrotron facilities is scarce and performing dynamic experiments require extensive preparation and a substantial amount of auxiliary equipment. Therefore, it is desirable to be able to perform some of the dynamic experiments using laboratory CT instrumentation.

The low photon flux of laboratory instruments leads to a compromise between image quality and the temporal resolution. Temporal resolution can be increased at the expense of image quality by decreasing scanning time. Scanning time is decreased by either reducing the exposure time of each projection, which decreases the signal-to-noise ratio or by reducing the number of projections gathered resulting in artefacts in the reconstruction^{9,22}. Bultreys et al.⁹ have built a laboratory instrument for in situ microtomography, where they managed to have an impressive time scale of just 12 seconds, by using a very short exposure time combined with a reduced number of projections^{9,23}.

¹Department of Applied Mathematics and Computer Science, Technical University of Denmark, 2800 Kongens Lyngby, Denmark. ²Department of Physics, Technical University of Denmark, 2800 Kongens Lyngby, Denmark. ³Helmholtz-Zentrum Hereon, Institute for Metallic Biomaterials, 21502 Geesthacht, Germany. ⁴These authors contributed equally: Henning Osholm Sørensen, Stefan Bruns and Anders BJORHOLM DAHL. ✉email: pwra@dtu.dk; anym@dtu.dk

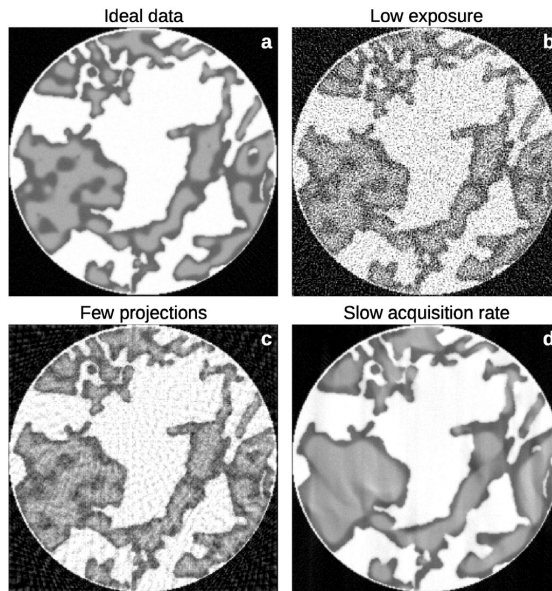


Figure 1. The reconstructions resulting from a tomographic experiment is highly affected by the experimental parameters. To visualise the potential effects that might occur in dynamic tomography, we have performed reconstructions of three data sets that are each limited in one experimental parameter. (a) The “ideal” reconstruction of the system, which is carbonate rock (white) filled with a fluid mixture of oil (dark grey) and water (light grey). (b) A reconstruction performed on data limited in the signal-to-noise ratio, i.e. short exposure or low X-ray brilliance. (c) A reconstruction from a data set with a low number of projections. (d) A reconstruction performed on data with high signal-to-noise (long exposure) and high number of radiographs, i.e. long data collection leading to low temporal resolution.

Figure 1 shows how three different data-deficiencies: limited exposure (high noise); a limited number of projections; and limited temporal resolution, affect the resulting reconstruction of a dynamic data set. The model system shown in Fig. 1 consists of a rock matrix (white) and two immiscible fluid phases, oil (dark grey) and water (light grey), that spontaneously separate over time. From this, we see that short exposures lead to a noisy reconstruction, few projections to line artefacts and long exposures to smeared fluid boundaries.

The most commonly used reconstruction techniques, filtered back projection (FBP) and its cone beam counterpart the Feldkamp, Davis, and Kress algorithm (FDK) are unsuited for data with the previously mentioned deficiencies^{24–26}.

This is because a good reconstruction using this type of algorithm requires a rather large number of projections N_{proj} , preferably $N_{\text{proj}} \gtrsim N_{\text{pix}}\pi/2$ where N_{pix} is the number of detector pixels²⁶. This means that thousands of low noise radiographs are needed to provide high-quality 3D reconstructions, eventually leading to high scan times – often in the order of hours²⁷.

It has previously been shown that iterative reconstruction techniques perform substantially better than FBP methods when the N_{proj} is limited—especially when prior knowledge about the object is leveraged²⁶.

Prior knowledge can be used to constrain the solution of the reconstruction algorithm to behave in a certain way. For instance, a solution can be encouraged to have a noise-free appearance by penalising the norm of the derivative of the reconstruction, which is the case in e.g. total variation regularisation²⁸.

Some simple examples of using prior knowledge are non-negativity constraints and box constraints. Non-negativity stems from the fact that attenuation coefficients are theoretically always positive. This can be extended to also include an upper limit to the values allowed in the reconstruction, i.e. box constraints. Setting the upper limit requires that the largest attenuation coefficient in the sample is known.

There have been several different attempts to leverage prior knowledge to improve the quality of reconstructions through iterative methods. Lin et al.²⁷ introduced a regularisation term during their minimisation similar to that of total variation regularisation. The Huber function is applied instead of the seminorm used in total variation, which preserves boundaries between different phases in the reconstruction²⁸. Lin et al.²⁷ tested their algorithm on a microCT data set of a Bentheimer sandstone, saturated with a mixture of brine and oil. They reported that their suggested algorithm provided a much-enhanced contrast between the reconstructed phases. Another approach was suggested by Myers et al.²⁹ who limited the number of unknowns in the equation by subtracting projections recorded on the initial static system from the projections of the dynamic system. This means that only the dynamic component is reconstructed. They used Simultaneous Iterative Reconstruction Technique

(SIRT) to reconstruct the difference projection data. Additionally, they encourage spatial localisation of changes between time steps, voxels within the static region are set to a fixed value and voxel values in the reconstruction are binarised, i.e. voxels are set equal to one of two values. In their case, this corresponds to either be empty or filled. Hence the reconstruction will also be automatically segmented.

That method was further developed into a Bayesian framework by Myers et al.³⁰. The Bayesian framework iteratively updates a solution such that the *maximum a posteriori* estimate of the solution is found. The solution is modelled as a sum of conditional probabilities, which ensures data fidelity, the physics of the system such as noise and correlations across time. Additionally, it is possible to add terms, which constrain the dynamic solution by using a static reconstruction, that directly segments the solution and terms that regularise it. The Bayesian algorithm presented in Myers et al.³⁰ is equivalent to the one presented in Myers et al.²⁹ if the assumptions such as binarisation and spatial localisation are applied to the Bayesian algorithm.

Binarisation or discretisation of attenuation values is commonly used to improve the reconstruction quality for samples with only a few unmixed well-defined phases. The discrete algebraic reconstruction technique (DART) and its extension total variation regularised discrete algebraic reconstruction technique (TVR-DART), presented by Batenburg et al.³¹ and Zhuge et al.³², are designed for such systems.

Van Eyndhoven et al.³³ has introduced a method, rSIRT-PWC, similar to the method by Myers et al.²⁹ i.e. they separated the dynamic system into two regions – a static and a dynamic. However, they take special care to handle pixels along the border of the dynamic and static regions. The attenuation value of pixels within the static region is set to zero while the attenuation value of pixels which are either partially or fully in the dynamic region is modelled as piecewise constant functions. This assumption is appropriate for their use case where a single fluid phase is propagating through a porous media. However, it is not appropriate for two-phase fluid flow cases, where the value of a voxel might change multiple times during the dynamic process.

In this paper, we present a method that is developed with the aim to reconstruct dynamic data from two-phase fluid flow experiments, but it can be used for any dynamic experiment, where it is possible to obtain a high-quality static data set of the initial system before initiating the dynamic experiment. This could for example be a core flooding experiment where projection images could be obtained from many angles and with long exposure times before the actual flooding experiment. With this large amount of low-noise data, a detailed image of the different parts of the sample such as rock-matrix and voids could be obtained. The information gathered from a high-quality reconstruction of the static system is the crux of our reconstruction algorithm. It is used to initialise an iterative reconstruction method, which will bring the algorithm closer to a desirable solution. The reconstruction of each time is initialised by the solution of the former step. Additionally, we constrain the solution with a segmentation of the static data set.

We have investigated the performance of our proposed approach by comparison to other SIRT based algorithms as well as the commonly used filtered-back projection (FBP) algorithm. The SIRT based algorithms we compare to are simpler versions of the algorithm we have developed. We compare the results of the different algorithms qualitatively by visual inspection and quantitatively using the ℓ_2 -norm of the residual between the reconstructions and the ground truth. Furthermore, we assess the resulting image contrast by comparing histograms of reconstructed voxel values.

Methodology

Reconstruction. An iterative reconstruction technique is used for this work. Typically, iterative reconstruction techniques attempt to solve the linear system

$$b = Ax, \quad (1)$$

where $x \in \mathbb{R}^n$ is the reconstructed volume stored as a vector, $b \in \mathbb{R}^m$ is the projection data or radiographs also stored as a vector, $A \in \mathbb{R}^{m \times n}$ is the forward projection operator or the system matrix. Determining an x that solves the equation is typically an ill-posed problem because there is either no solution or the solution is not unique. Hence, a direct inversion of Eq. (1) is not possible³⁴. We, like others, have chosen to employ the iterative reconstruction method, SIRT, because it is a robust technique and it allows us to incorporate prior knowledge when solving the linear set of equations^{29,33,35,36}.

The basic principle behind the SIRT algorithm is that it uses the residual between the forward projection of the current reconstruction and the radiographs to update the solution. The update step of the SIRT algorithm is given by

$$x^{(k+1)} = x^{(k)} + CA^T R(b - Ax^{(k)}), \quad (2)$$

where $x^{(k)} \in \mathbb{R}^n$ is the image obtained at the k th iteration, $A^T \in \mathbb{R}^{n \times m}$ is the backward projection operator, $C \in \mathbb{R}^{n \times n}$ is a diagonal matrix containing the inverse column sums of A i.e. $c_{jj} = 1/\sum_i a_{ij}$, and $R \in \mathbb{R}^{m \times m}$ is a diagonal matrix of the inverse row sums of A i.e. $r_{ii} = 1/\sum_j a_{ij}$ ³⁷.

The starting point of the reconstruction $x^{(0)}$ can be initialised with an arbitrary vector of real numbers. However, a vector where each element has the same value (normally zero) is generally used^{29,32,33}. In the present example, the rock matrix does not change during the experiment. Hence, all the voxels in the rock matrix should have constant intensity independent of the time step, and these voxels make up a large part of the sample. This means that we can initialise the first time step of the dynamic reconstruction with the high-quality reconstruction of the static sample. Additionally, for a time series of data, we suggest initialising $x^{(0)}$ for time step, t , with the solution of the previous time step, $t - 1$, since that reconstruction is expected to be closer to our solution than a vector of zeroes.

As mentioned the rock matrix should not change during the experiment. Hence, we can also use the high-quality static reconstruction to constrain our solution. We can determine the rock matrix voxels via segmentation of the high-quality static reconstruction, which we can use to force the algorithm to keep the voxel values of the rock matrix constant. Mathematically, this operation is equivalent to projecting the right hand side in Eq. (2) onto a convex set \mathcal{C} , which only contains allowed values, using the projection operator $\mathcal{P}_{\mathcal{C}}$ ³⁸. The projection operator is also used to apply the box constraints mentioned in the “Introduction” section, where the set would be given by

$$\mathcal{C} = [\mu_{\min}, \mu_{\max}]^n, \tag{3}$$

with μ_{\min} being the smallest attenuation value in the sample and μ_{\max} being the largest.

The set we project our solution onto depends on the classification of each voxel, which is derived from the segmentation. We obtain the segmentation by thresholding the static reconstruction to identify regions of either rock or fluid. Voxels with a value above the threshold are defined as rock and fixed at the expected value while voxels below the threshold might be fluid. A voxel is only defined as fluid if its value is between the attenuation values of oil and water. This leaves us with voxels which have a larger attenuation value than water but smaller than rock. Voxels within this interval cannot be uniquely assigned to either fluid or rock and are therefore subjected to regular box constraints shown in Eq. (3). Using this technique, the iterative updating step is given by

$$x^{(k+1)} = \mathcal{P}_{\mathcal{C}}\left(x^{(k)} + \mathbf{CA}^T \mathbf{R}(b - \mathbf{Ax}^{(k)})\right) = \mathcal{P}_{\mathcal{C}}\left(\text{SIRT}\left(x^{(k)}\right)\right). \tag{4}$$

The SIRT algorithm from the ASTRA toolbox is used because it provides highly optimised C++ and CUDA code that can be called via a Python (or Matlab) interface. This enables the use of one or more GPUs to perform the reconstructions, which is substantially more effective than using CPUs^{39–41}. The projection operation $\mathcal{P}_{\mathcal{C}}$ is performed with NumPy in Python.

Stopping criteria. A general problem associated with iterative reconstruction methods is to determine when the optimal solution is obtained. Ideally, we would like to stop iterating when the minimal ℓ_2 -norm of the residual between the ground truth and the reconstruction is reached i.e. we wish to minimise

$$\text{Figure of merit} = \|x^{(k)} - \bar{x}\|_2, \tag{5}$$

where \bar{x} is the ground truth. Due to noise in the projection data, the solution $x^{(k^*)}$, which minimises the figure of merit, might not be where Eq. (4) converges to as $k \rightarrow \infty$ ³⁸.

The ground truth, \bar{x} , is not known in a real experiment, so we have to find a way to minimise Eq. (5) without being able to compute it directly. Multiple stopping rules have been proposed in the literature, however, using the normalised cumulative periodogram (NCP) of the residual, $r^{(k)} = b - \mathbf{Ax}^{(k)} \in \mathbb{R}^n$, seems to stop the algorithm close to the optimal solution^{42–45}.

The NCP stopping rule is based on the assumption that the residual, $r^{(k)}$, will have an NCP similar to the NCP of white noise when Eq. (5) is minimised, because there should only be white noise left in the residual at this point. This means that all information have been extracted from the projection data and the reconstruction can therefore be terminated.

Seeing if the residual is consistent with white noise requires calculating the periodogram. A periodogram is defined as the absolute squared values of the discrete Fourier coefficients of a vector. The periodogram of the residual vector is given

$$\hat{p}_i = |\hat{r}_i|^2, \quad i = 1, 2, \dots, q \quad \text{with} \quad \hat{r} = \text{DFT}(r). \tag{6}$$

DFT denotes the discrete Fourier transform and $q = \lceil n/2 \rceil$. The reason why only approximately half of the elements of rr are used to calculate \hat{p} is because the Fourier coefficients in the power spectrum of a real vector are symmetric around the midpoint of the vector.

The normalised cumulative periodogram (NCP) is now defined as

$$c_j(r) = \frac{\hat{p}_2 + \dots + \hat{p}_{j+1}}{\hat{p}_2 + \dots + \hat{p}_{q+1}}, \quad j = 1, 2, \dots, q. \tag{7}$$

Note that the first element of \hat{p} , known as the DC-component, is excluded from the definition such that it starts in (0, 0). The NCP value of white noise is expected to be a straight line ranging from (0, 0) to (q, 1). This line c_{white} can be used for comparison with the NCP of the residual. This can be done using the ℓ_2 -norm $r_{\text{NCP}} = \|c(r) - c_{\text{white}}\|_2$.

A detailed description of how the NCP stopping rule is used can be found in⁴⁵. A major benefit of this method is that it adapts to noise level in the projection data.

We terminate the reconstruction in our implementation when two iterations on either side of $r_{\text{NCP}}^{(k)}$ are larger than $r_{\text{NCP}}^{(k)}$. We require two iterations to prevent small fluctuations of $r_{\text{NCP}}^{(k)}$ from terminating the reconstruction prematurely. It was found that r_{NCP} exhibited more than one minimum at low noise levels. The algorithm, therefore, iterates beyond the first detected minimum to inspect if the current is a local minimum, i.e. if there should exist a second r_{NCP} minimum.

The NCP stopping rule is, computationally, fairly demanding since $r^{(k)} = b - \mathbf{Ax}^{(k)}$ has to be calculated along with its discrete Fourier transform after every iteration. $r^{(k)}$ is calculated by ASTRA during the SIRT update step, however, only the norm of it can be retrieved which makes it necessary to calculate it explicitly after a SIRT

	Box constraints	Initialisation	Local box constraints
SIRT	✗	✗	✗
SIRT-BC	✓	✗	✗
SIRT-IC	✓	✓	✗
SIRT-LC	✓	✓	✓

Table 1. An overview of the different approaches used for the four SIRT reconstruction algorithms that were tested. See text or the “Reconstruction” section for a detailed explanation.

update. CuPy is utilised to speed up the computation of the NCP via CUDA as the Fourier transform especially can benefit from parallelisation⁴⁶. The forward projection is calculated using the ASTRA toolbox.

The implementation of the reconstruction algorithm is shown in Algorithm 1. The algorithm starts with the initialisation of the current time step using either the static reconstruction or the reconstruction of the previous time step.

This is followed by a loop where the actual reconstruction is performed. The loop is limited to N_{\max} iterations to prevent the algorithm from failing to terminate. The stopping criterion is simplified as the actual implementation can handle cases where only one iteration is needed before convergence. Additionally, the implementation also continues iterating beyond the first detected minimum to ensure it is not stopping prematurely.

Algorithm 1: Pseudo-code of the algorithm used for reconstruction.

```

Input:  $b, x_{\text{initial}}$ 
1  $x_0 \leftarrow x_{\text{initial}}$ 
2 for  $t$  from 1 to  $N_{\text{time}}$  do
3    $k \leftarrow 0$ 
4    $x_t^{(k)} \leftarrow x_{t-1}$ 
5   converged  $\leftarrow$  false
6   while  $N_{\text{itr}} < N_{\max}$  or converged do
7      $x_t^{(k)} \leftarrow \text{SIRT}(x_t^{(k)}, b_t)$  // Calculated with ASTRA on GPUs
8      $k \leftarrow k + 1$ 
9      $x_t^{(k)} \leftarrow \mathcal{P}_{\mathcal{C}}(x_t^{(k)})$  // Calculated with Python
10     $c^{(k)} \leftarrow \text{NCP}(x_t^{(k)})$  // Calculated with Python
11     $N_{\text{itr}} \leftarrow N_{\text{itr}} + 1$ 
12    if ( $k > 4$  and  $c^{(k-2)} = \min(c)$ ) then
13       $x_t \leftarrow x_t^{(k-2)}$ 
14      converged  $\leftarrow$  true
15    end
16  end
17 end
18 return  $x$ 

```

Reconstruction algorithms used. We have chosen four versions of the SIRT algorithm and the FBP method to test the performance of our algorithm. The latter will serve for comparison as it is the most commonly used algorithm for tomographic reconstruction²⁶. The SIRT algorithms are also compared to a FBP reconstruction and an ideal FBP which uses 720 projections and $\rho = 0.25\%$.

The differences between the SIRT methods used are shown in Table 1. *Box constraints* means that the attenuation coefficients of voxels in the reconstruction are truncated to the minimum and maximum values present in the sample. For the present case, this means a lower limit of 0 and an upper limit of 2.5. *Initialisation* refers to initialising time step $t = 0$ with the reconstruction of the static system and the remaining time steps with a reconstruction of the previous time step. *Local box constraints* refers to projecting the reconstruction onto the convex set \mathcal{C} created with a segmentation of the static reconstruction as explained in the “Reconstruction” section. This means the attenuation value of voxels identified as chalk are set to the same attenuation value as that of chalk and the attenuation value of voxels identified as fluid are confined to be within the interval of oil and water.

Results and discussion

Comparison of the reconstructions. The method is tested on a synthetic data set that consists of a rock matrix with a homogeneous mixture of water and oil (Fig. 2), that separate over time as they are immiscible. The details of this simulation can be found in the “Methods” section. Working with simulated data enables quantitative comparisons between the different reconstruction methods since we have the ground truth. We will, from now on, use the term *residual* as the difference between the ground truth and the reconstruction unless otherwise stated. We have chosen to quantitatively examine the reconstruction methods in four ways:




Figure 2. 3D visualisation of the simulation. The rock matrix have been removed from the upper part of the simulation along with upper front part of the fluid phase. The rock is brown, the water is blue and the oil is light grey. This figure is created with ParaView⁴⁷.

- The ℓ_1 -norm of the residual between ground truth and reconstruction. This can be found in the supplementary material.
- The ℓ_2 -norm of the residual.
- The distribution of voxel values in the reconstructions.
- The distribution of the residual. This can be found in the supplementary material.

Only the voxels within the sample area are used for the quantitative analysis i.e. the air surrounding the sample is ignored.

Visual appearance of the reconstructions. Figure 3 shows two reconstruction series, one in the top row with low noise ($\rho = 0.25\%$) and a large number of projections ($N_{\text{proj}} = 360$) and one in the bottom row with high noise ($\rho = 5.0\%$) and a low number of projections ($N_{\text{proj}} = 45$). N_{proj} refers to the number of projections in a data set and ρ refers to the relative noise level in a data set. A detailed explanation of the noise in the data sets can be found in the “Noise” section. The remaining reconstructions series can be found in Section S1 of the supplementary material. It is obvious from visual inspection that all reconstruction techniques used perform well when applied to the data set with a large number of low noise projections in the top row. In this case, the primary concern becomes computational speed.

For the other extreme, we have a data set with high noise (5%) and few projections (45), shown in the bottom row of Fig. 3, a significant difference is found in the obtained image qualities. Here the FBP reconstruction becomes very noisy. Almost to the point where it is impossible to differentiate between the two fluid phases. SIRT and SIRT-BC perform similarly, which indicates the addition of box constraints in SIRT-BC does not improve the reconstruction significantly. A major improvement is found when the reconstruction is initialised using the high-quality static data as described in the Reconstruction, which can be seen for the SIRT-IC and SIRT-LC reconstructions. The fluid phases are clearly visible using both, but SIRT-IC exhibits a fair bit of noise, which is eliminated by the local box constraints used in SIRT-LC.

ℓ_2 -norm of the residual. The performance of the algorithms has been quantified by calculating the ℓ_2 -norm of the residual between the reconstructions and the ground truth for each time step in the simulation. This has been plotted as a function of time in Fig. 4. The figure confirms that all algorithms provide good and similar results for data set reconstructed using the low noise $\rho = 0.25\%$ and a high number of projections, 360. Noticeably, they all perform almost as well as the FBP reconstruction with 720 projections and $\rho = 0.25\%$, the data set that represents a reconstruction under “ideal” conditions. It is apparent that FBP solution quickly deteriorates as noise increases and the number of projections is reduced. The same is partly true for SIRT and SIRT-BC, but it is not as significant. The addition of box constraints does improve the ℓ_2 -norm of the residual noticeably. However, this effect becomes increasingly less pronounced as the data degrades. The addition of initialisation substantially improves the reconstruction when the data quality degrades.

The ℓ_2 -norm for both SIRT-IC and SIRT-LC vary across time due to the initialisation which links the current time step with the previous. Both reconstructions initially improve slightly in the best data case after which their performance slightly degrades. This behaviour gets less pronounced as the quality of the data deteriorates. We performed SIRT reconstructions on a special data set where the simulation was frozen such that the first time step in the simulation was repeated for all time steps. The noise in each time step is unique. This was done to ensure that the deterioration of performance seen in Fig. 4 of SIRT-IC and SIRT-LC across time is not because




Figure 3. An example of the five different reconstruction algorithms for the best (upper row) and worst (lower row) data cases. Slice 171/256 at time step 51/100 is shown in the figure. Note that scale bar is truncated to [0, 2.5]. This makes the effect of box constraints present in SIRT-BC, SIRT-IC and SIRT-LC less pronounced.




Figure 4. The ℓ_2 -norm of the residual as function of the time step for the three different cases of noise and number of projections. The ideal FBP reconstruction ($\rho = 0.25\%$, 720 projections) is shown for comparison. Notice that the y-axis range is different on the three plots.

the algorithms diverge. This test showed that SIRT-IC and SIRT-LC improves across time. The results from this test can be found in Fig. S8 in the supplementary material.

NCP stopping criteria and convergence. The challenge when using iterative techniques to solve the linear set of equations is to determine when the optimal solution has been obtained. Here we will analyse the performance of the NCP criteria, which is used to terminate the iterative algorithms. In Fig. 4 we observed that the ℓ_2 -norm increased for later time steps. This behaviour seems to be related to the performance of the NCP stopping rule, which terminates prematurely for low noise data. In general, the method seems less suited for low noise data. This is especially true for the initialised algorithms. The number of iterations taken before the NCP stopping rule is met for each time step is shown in Fig. 5.

A general trend for all methods is that the number of iterations needed decreases as the quality of the data decreases. This is because the residual will resemble white noise more quickly as the noise level increases. The necessary number of iterations depends more on the noise level than the number of projections. This can be deduced by examining Fig. S11 in the supplementary material which shows the iterations needed for all the iterative algorithms on all data sets.

SIRT-IC and SIRT-LC show some variability in the number of iterations required compared to the two other algorithms. Initially, we see a sharp increase in the number of iterations required which is followed by a long decay. When starting, few iterations are needed because the algorithms are initialised with a reconstruction that already has converged according to the NCP criterion. The simulation changes most rapidly for the first time steps which means more iterations are needed in this period of the simulation compared to later on where the dynamics of the simulation slow down.

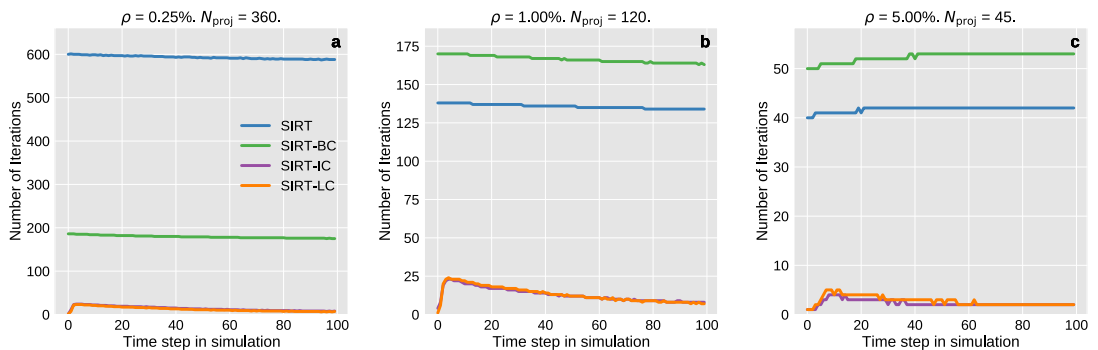


Figure 5. The number of iterations required in each time step before the NCP stopping criteria is met for the three different data cases. SIRT-IC and SIRT-LC nearly coincides in all cases.

This is confirmed by reversing the dynamics and performing the reconstruction on this reversed data set. The iterations needed for the reversed reconstruction is shown in Fig. 6, where we see the number of iterations needed gradually increase as the dynamics of the simulation increase.

The availability of the ground truth makes it possible to evaluate how well the NCP algorithm is at terminating at correct iteration number. This is done by comparing the solution achieved using the NCP stopping criteria with an “ideal” solution which minimises Eq. (5). The comparison consists of calculation the differences between the number of iterations used by the two stopping criteria and the difference between the ℓ_2 -norm of the two stopping criteria. This is shown for the first 20 time steps of the simulation using 120 projections with a noise level of 1.0% in Fig. 7 for the SIRT and SIRT-LC algorithms. In plot a we see that the NCP criteria with SIRT in general overestimates the number of iterations needed which results in a slight increase in the ℓ_2 -norm when compared to the ideal case which can be seen in plot b of Fig. 7. The SIRT-LC algorithm initially underestimates the number of iterations needed after which it remains fairly close to the ideal solution. In general, the NCP stopping criteria works best when the noise level is 1.0% or above and the number of projections is 120 or below. The behaviour of SIRT and SIRT-BC is very similar and the same is true for SIRT-IC and SIRT-LC.

Histograms of voxel values. A more direct way to compare the performance of the reconstruction algorithms is to examine the distribution of reconstructed voxel values compared to the actual voxel values in the simulation across all time steps. Some of these results are shown in Fig. 8.

The black line represents the distribution of voxel values found in the simulation. There are three distinct peaks which correspond to three phases, oil at 1.0, water at 1.7 and rock at 2.5. Values between 1.0 and 1.7 are primarily related to the mixture of oil and water, however, it can also be related to the partial volume effects at the interface between fluid and mineral, which can range from 1.0 to 2.5.

The plot in Fig. 8 a shows the ideal data case and confirms that all algorithms give similar results for this data set as was found analysing Figs. 3 and 4. The effect of box constraints is noticeable as both FBP and SIRT have a high amount of voxels with values that far exceed the upper limit of 2.5. We also see that SIRT-IC has a tail towards 0 in plot c that could be a result of the limited amount of iterations used by the algorithm for that specific data set. Looking at the worst data case in c SIRT-IC and SIRT-LC are the only algorithms that keep having noticeable peaks, although SIRT-LC does appear a bit sharper. This increase in contrast fits well with the difference in visual appearance between in SIRT-IC and SIRT-LC as seen in Fig. 3.

Challenging regions in the reconstructions. In Fig. 9 the residual is shown for the best and worst data case for slice 171 at time step 51. In the upper row, we see that FBP and SIRT has most of their errors spread out compared to the remaining algorithms. SIRT-BC has most of its errors at the transition between the rock and fluid phase. In contrast SIRT-IC and SIRT-LC do fairly well in general. There are, however, large regions within the fluid phase in both of the reconstructions that are either overestimated or underestimated. This is again caused by the underestimation of iterations needed.

The bad data case shows there is no discernible area which the FBP fails to reconstruct, unlike the SIRT algorithms where there is a definite structure in the plots. SIRT-IC and SIRT-LC still handle the reconstruction fairly well with SIRT-LC being a bit more smooth.

Global performance of the algorithms.

The ℓ_1 -norm and ℓ_2 -norm of the residual for the entire 4D reconstruction is shown in Table 2. The table shows that SIRT-LC is superior to the other algorithms in the bad data case and slightly inferior to SIRT-IC for the best data case for the ℓ_2 -norm. This was also expected based on Fig. 4 where the values for SIRT-IC are below the values of SIRT-LC. A table of all data cases is available in the supplementary material where it is seen that SIRT-LC is the best algorithm in general.




Figure 6. The number of iterations required in each time step before the NCP stopping criteria is met for the case with $\rho = 1.00\%$ and $N_{\text{proj}} = 120$ for the SIRT-LC algorithm.

Conclusions

We have developed a reconstruction algorithm for dynamic tomography, based on SIRT. Our algorithm targets experiments where it is possible to collect an initial high-quality tomography data set before the dynamic experiment is initiated. The reconstruction of the static system is used to initialise and constrain the reconstructions of the dynamic data via a segmentation of the static system in order to strongly regularise the solution. Additionally, we use the NCP stopping criterion to optimise the number of iterations used. We have shown using simulated data that this procedure significantly improves the quality of the reconstruction of data with a minimum number of projections and high noise levels to a point, where it is comparable to an ideal traditional reconstruction even when using poor data.

Methods

Computational fluid dynamics simulation. We test our reconstruction algorithm using a synthetic dynamic data set as the ground truth. The data set consists of a rock matrix with a homogeneous mixture of two immiscible fluids, modelled as an emulsion of water and oil, that separate over time while being driven upward by a small body force. The separation is initially fairly vigorous, i.e. the dynamics during this period of the experiment is much faster than later in the separation process, yielding a data set that mimics an experimental two-phase fluid system.

A segmented nanoCT data set collected on a piece of chalk, a fine-grained carbonate rock, provided a realistic environment for simulating a dynamic data set. The nanotomography measurements were performed at BL47XU, SPring-8, Japan⁴⁸, providing a voxel size of 38 nm. 1800 projections were recorded while rotating the sample 180° with an exposure time of 150 ms. The projection data were dark current and bright field corrected. The truncated sinogram, due to a smaller FOV than the sample dimension, were completed⁴⁹ and to avoid ring artefacts in the reconstructed image stripe artefacts were reduced in the sinogram⁵⁰ before the 3D volume was reconstructed using the GridRec algorithm in TomoPy⁵¹. Noise in the 3D image was reduced using our iterative nonlocal means method⁵². A cylindrical rock matrix was made by taking a subvolume of 256³ voxels whereafter voxels outside a radius of 124 voxels were removed slice by slice. We mirror the rock matrix along its vertical axis to allow for vertical periodic boundary conditions of the simulation domain, i.e. the resulting cylindrical volume has a diameter of 248 voxels and length of 512 voxels.

Multiphase flow simulations were conducted following the formulation of a phase-field Lattice Boltzmann method for isothermal and incompressible fluid systems as given by Fakhari et al.^{53,54} with a custom CUDA implementation. Implementation details and parameter settings, that have been used but are not essential to our findings here, are presented in Table S2 of the supplementary material.

The initial system contains a fluid mixture of equal amounts of oil and water in every wet node that separates into an equivolumentric mixture of two separate phases with a density ratio of about 4:3, a dynamic viscosity ratio of about 3:4 and a three-phase contact angle of 90° at the rock matrix interface. The differentiability of the phase-field over the course of the simulation was ensured by modelling fluid-fluid interfaces with a three voxel wide smooth transition. Snapshots of the multiphase dynamics were generated by exporting the phase-field first after running the simulation for 3000 steps and then after every additional 100 steps until 100 frames were collected that are subsequently called time step 0 to 99. The top half of the simulation i.e. the “mirrored” part was excluded from the volume used for the simulation of the tomography experiment.




Figure 7. Evaluation of the performance of the NCP stopping criteria. In plot (a) the difference between the number of iterations used with the NCP stopping criteria and the ideal number of iterations a function of time step in the simulation. In plot (b) the difference between the ℓ_2 -norm of residual when using the NCP stopping criteria and the ideal ℓ_2 -norm is shown as a function of time step in the simulation.




Figure 8. Histograms of the voxel values on a logarithmic scale of the ground truth image and the reconstructions. Plots (a–c) show histograms for reconstructions of data sets with 360 number of projections with 0.25% noise, 120 number of projections with 1.00% noise and 45 number of projections with 5.00% noise respectively.

The numerical value of the phase-field voxels exported from the simulation was set equal to values measured experimentally with a laboratory CT scanner in a two-phase system presented in Lin et al.²⁷. However, they used a Bentheimer sandstone instead of a carbonate. Their scan was performed at 80 KeV and both the brine and the decane used to saturate the sandstone were doped with 3.5 wt% potassium iodide. Using these measured values ensures that the contrast between the different phases of the system is comparable to a real experiment. The interface between the rock and fluid was smeared using a Gaussian filter to emulate partial volume effects, i.e. voxels, which are composed of both rock and fluid. The first recorded time frame of the simulation can be seen in Fig. 2 where the rock matrix is shown in brown and the water and oil are shown in blue and white respectively. The top part of the rock matrix along with half of the fluid phase is transparent in the figure.

Simulation of a dynamic X-ray experiment. *Forward projection.* The fluid dynamic simulation is forward projected using a parallel beam geometry with the ASTRA tomography toolbox. The forward projection operator of ASTRA does not reflect the energy spectrum of a laboratory X-ray source and can be viewed as perfectly monochromatic. Projection angles are distributed uniformly between 0° and 180° as angles between 180° and 360° are redundant when using a parallel beam setup. The detector response is modelled as perfect and with a width of 300 pixels, to ensure that the full sample width is covered.

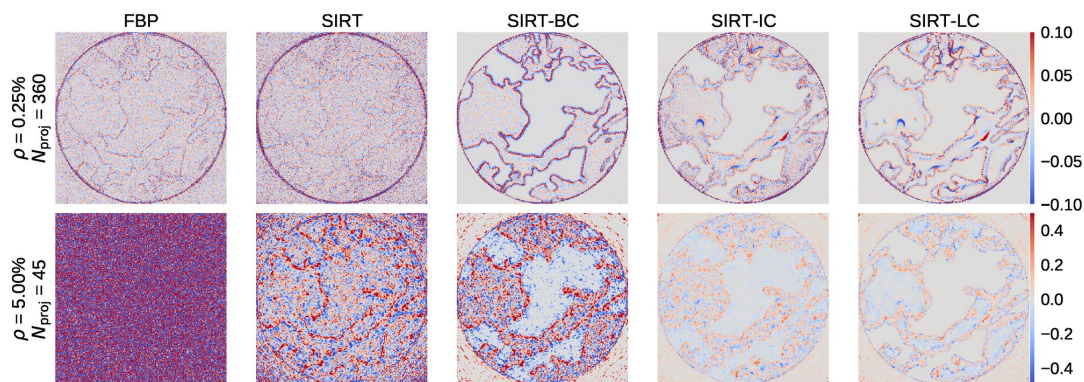


Figure 9. An example of the residual of the five different algorithms for the best (upper row) and worst (lower row) data cases. Slice 171/256 at time step 51/100 is shown in the figure. Pixel values are constrained to be within ± 0.1 in the upper row and within ± 0.5 in the lower row.

	$N_{proj} = 360$ $\rho = 0.25\%$		$N_{proj} = 45$ $\rho = 5.00\%$	
	$\ell_1 (\cdot 10^7)$	$\ell_2 (\cdot 10^3)$	$\ell_1 (\cdot 10^7)$	$\ell_2 (\cdot 10^3)$
FBP	4.13	1.51	179.59	63.21
SIRT	4.76	1.70	32.49	11.49
SIRT-BC	2.78	1.54	25.77	11.06
SIRT-IC	2.74	1.34	8.68	3.69
SIRT-LC	2.31	1.43	6.25	3.34

Table 2. Table of the ℓ_1 - and ℓ_2 -norms for the best and worst data case. Bold numbers indicate the best performing algorithm.

Each time step of the simulation is forward projected independently, which means that we approximate individual steps as static. An alternative to this would be to include the dynamics in the forward projection such that the simulation develops between each projection. This would also make it possible to account for the time it takes for the gantry to rotate the sample which is negligible in our local CT scanner but considerable in other CT scanners. Simulating the acquisition of radiographs makes it much harder to perform quantitative analysis as a time step in the reconstruction will be composed of multiple time steps in the simulation and is therefore not done.

If we use the same geometric matrix, A , to perform the forward projection as for the reconstruction, we will commit the so-called inverse crime, i.e. that we use the exact same discretisation both ways⁵⁵. To avoid committing the inverse crime the forward projected data are rotated with respect to the grid of the reconstruction.

The number of projections needed for a good reconstruction using standard FBP should be larger than $N_{pix}\pi/2^{26}$. Our detector size of 300 pixels, means that at least 471 projections are required to perform an FBP reconstruction of high quality. To be a bit conservative 720 projections are used for both the high-quality static prior and an ideal FBP reconstruction. 45, 120 and 360 projections are used for the numerical experiments, which represent experiments with a low, a moderate and a high number of projections.

Noise. In real experiments, the recorded projection data will be affected by noise. The data obtained from an X-ray detector can often be assumed to follow Poisson statistics, i.e. the variance of the signal is equal to the signal itself. To do this b have to be converted from the negative logarithmic scale to photon counts (Step 3 in Algorithm 2). The next step of the algorithm is applying noise to the rescaled data. This is done by sampling a Poisson distribution where the intensity in each detector pixel is used as the mean of the distribution (Step 4 in Algorithm 2). Since the noise operation can only be applied on integers the floor function is applied first. The noisy projection data is scaled back to the negative logarithmic scale and returned along with the noise vector e .

Algorithm 2: Pseudo-code of the algorithm used to apply noise to the projection data.

Input: $A\bar{x}$, I_0
 1 $b \leftarrow \text{copy}(A\bar{x})$
 2 $b_{\max} \leftarrow \max(b)$
 3 $b \leftarrow \lfloor I_0 \exp(-b/b_{\max}) \rfloor$
 4 $b \leftarrow \text{Poisson}(b)$
 5 $b \leftarrow -b_{\max} \log(b/I_0)$
 6 $e \leftarrow b - A\bar{x}$
 7 **return** b , e

Modelling the noise as a single Poisson distribution is not entirely accurate as laboratory X-ray sources provide a wider spectrum of X-ray energies, which each differ in transmission through the sample. Hence, it in principle gives rise to multiple Poisson distributions with different means. Additionally, the X-ray detector will exhibit electrical noise which can be modelled as Gaussian noise^{56–59}.

The relative noise level in the projection data is calculated using

$$\rho = \frac{\|e\|_2}{\|A\bar{x}\|_2} \quad (8)$$

where e is the noise vector added to the forward projection of the ground truth $A\bar{x}$. We have chosen to use three noise levels which represent low, moderate and high-level noise, which is equivalent to 0.25%, 1% and 5%. These levels were based on a qualitative comparison between the noisy simulated projection data and data acquired by our local CT scanner where a noise 0.25% is generally the noise level of a high-quality scan.

Simulated experiments. The numerical experiments were performed using the three different image noise levels (0.25%, 1% and 5%) and with three different number of projections (45, 120 and 360), i.e. nine simulated experiments will be reconstructed.

The nine different data sets are reconstructed using the five algorithms described in the “[Reconstruction algorithms used](#)” section.

Data availability

All code used for the paper along with the data sets of the attenuation coefficients are available at <https://gitlab.gbar.dtu.dk/pwra/NumericalExperiments> and <https://doi.org/10.11583/DTU.c.5448594>. All figures are created with Matplotlib⁶⁰ except where noted otherwise.

Received: 22 February 2021; Accepted: 31 May 2021

Published online: 14 June 2021

References

- Hu, X., Hu, S., Jin, F. & Huang, S. (eds.). *Physics of Petroleum Reservoirs*, chap. 2.4 Permeability of Reservoir Rocks (Springer-Verlag, 2016).
- Blunt, M. J. *et al.* Pore-scale imaging and modelling. *Adv. Water Resour.* **51**, 197–216. <https://doi.org/10.1016/j.advwatres.2012.03.003> (2013).
- Arns, C. H., Knackstedt, M. A., Val Pinczewski, W. & Lindquist, W. B. Accurate estimation of transport properties from microtomographic images. *Geophys. Res. Lett.* **28**, 3361–3364. <https://doi.org/10.1029/2001GL012987> (2001).
- Misztal, M. K., Hernandez-Garcia, A., Matin, R., Sorensen, H. O. & Mathiesen, J. Detailed analysis of the lattice Boltzmann method on unstructured grids. *J. Comput. Phys.* **297**, 316–339. <https://doi.org/10.1016/j.jcp.2015.05.019> (2015).
- Tölke, J. & Krafczyk, M. Second order interpolation of the flow field in the lattice Boltzmann method. *Comput. Math. Appl.* **58**, 898–902. <https://doi.org/10.1016/j.camwa.2009.02.012> (2009).
- Ramstad, T., Idowu, N., Nardi, C. & Øren, P. E. Relative permeability calculations from two-phase flow simulations directly on digital images of Porous rocks. *Trans. Porous Media* **94**, 487–504. <https://doi.org/10.1007/s11242-011-9877-8> (2012).
- Cnudde, V. & Boone, M. N. High-resolution X-ray computed tomography in geosciences: a review of the current technology and applications. *Earth Sci. Rev.* **123**, 1–17. <https://doi.org/10.1016/j.earscirev.2013.04.003> (2013).
- Wildenschild, D. & Sheppard, A. P. X-ray imaging and analysis techniques for quantifying pore-scale structure and processes in subsurface porous medium systems. *Adv. Water Resour.* **51**, 217–246. <https://doi.org/10.1016/j.advwatres.2012.07.018> (2013).
- Bultreys, T. *et al.* Fast laboratory-based micro-computed tomography for pore-scale research: Illustrative experiments and perspectives on the future. *Adv. Water Resour.* **95**, 341–351. <https://doi.org/10.1016/j.advwatres.2015.05.012> (2016).
- Berg, S. *et al.* Real-time 3D imaging of Haines jumps in porous media flow. *Proc. Natl. Acad. Sci.* **110**, 3755–3759. <https://doi.org/10.1073/pnas.1221373110> (2013).
- Singh, K. *et al.* Dynamics of snap-off and pore-filling events during two-phase fluid flow in permeable media. *Sci. Rep.* **7**, 1–13. <https://doi.org/10.1038/s41598-017-05204-4> (2017).
- Lin, Q. *et al.* Minimal surfaces in porous media: Pore-scale imaging of multiphase flow in an altered-wettability Bentheimer sandstone. *Phys. Rev. E* **99**, 1–13. <https://doi.org/10.1103/PhysRevE.99.063105> (2019).
- Spurin, C., Bultreys, T., Bijeljic, B., Blunt, M. J. & Krevor, S. Intermittent fluid connectivity during two-phase flow in a heterogeneous carbonate rock. *Phys. Rev. E* **100**, 1–10. <https://doi.org/10.1103/PhysRevE.100.043103> (2019).
- Scanziani, A., Singh, K., Bultreys, T., Bijeljic, B. & Blunt, M. J. In situ characterization of immiscible three-phase flow at the pore scale for a water-wet carbonate rock. *Adv. Water Resour.* **121**, 446–455. <https://doi.org/10.1016/j.advwatres.2018.09.010> (2018).
- Scanziani, A., Singh, K., Menke, H., Bijeljic, B. & Blunt, M. J. Dynamics of enhanced gas trapping applied to CO₂ storage in the presence of oil using synchrotron X-ray micro tomography. *Appl. Energy* **259**, 114136. <https://doi.org/10.1016/j.apenergy.2019.114136> (2020).
- Menke, H. P., Bijeljic, B., Andrew, M. G. & Blunt, M. J. Dynamic three-dimensional pore-scale imaging of reaction in a carbonate at reservoir conditions. *Environ. Sci. Technol.* **49**, 4407–4414. <https://doi.org/10.1021/es505789f> (2015).

17. Menke, H. P., Bijeljic, B. & Blunt, M. J. Dynamic reservoir-condition microtomography of reactive transport in complex carbonates: Effect of initial pore structure and initial brine pH. *Geochim. Cosmochim. Acta* **204**, 267–285. <https://doi.org/10.1016/j.gca.2017.01.053> (2017).
18. Yang, Y. *et al.* Retraction of the dissolution front in natural porous media. *Sci. Rep.* **8**, 5693. <https://doi.org/10.1038/s41598-018-23823-3> (2018).
19. Yang, Y. *et al.* Direct Observation of Coupled Geochemical and Geomechanical Impacts on Chalk Microstructure Evolution under Elevated CO₂. *ACS Earth Space Chem.* **2**, 618–633. <https://doi.org/10.1021/acsearthspacechem.8b00013> (2018).
20. Yang, Y. *et al.* Effect of cumulative surface on pore development in chalk. *Water Res. Res.* <https://doi.org/10.1029/2018WR023756> (2019).
21. Willmott, P. *An Introduction to Synchrotron Radiation* (Wiley, New York, 2019).
22. Davis, G. R. & Elliott, J. C. Artefacts in X-ray microtomography of materials. *Mater. Sci. Technol.* **22**, 1011–1018. <https://doi.org/10.1179/174328406X114117> (2006).
23. Bultreys, T. *et al.* Real-time visualization of Haines jumps in sandstone with laboratory-based microcomputed tomography. *Water Resour. Res.* **51**, 8668–8676. <https://doi.org/10.1002/2015WR017502> (2015).
24. Feldkamp, L. A., Davis, L. C. & Kress, J. W. Practical cone-beam algorithm. *J. Opt. Soc. Am. A* **1**, 612. <https://doi.org/10.1364/josaa.1.000612> (1984).
25. Boas, F. E. & Fleischmann, D. CT artifacts: causes and reduction techniques. *Imag. Med.* **4**, 229–240. <https://doi.org/10.2217/iim.12.13> (2012).
26. Maire, E. & Withers, P. J. Quantitative X-ray tomography. *Int. Mater. Rev.* **59**, 1–43. <https://doi.org/10.1179/1743280413y.0000000023> (2013).
27. Lin, Q., Andrew, M., Thompson, W., Blunt, M. J. & Bijeljic, B. Optimization of image quality and acquisition time for lab-based X-ray microtomography using an iterative reconstruction algorithm. *Adv. Water Resour.* **115**, 112–124. <https://doi.org/10.1016/j.advwatres.2018.03.007> (2018).
28. Rudin, L. I., Osher, S. & Fatemi, E. Nonlinear total variation based noise removal algorithms. *Phys. D Nonlinear Phenomena* **60**, 259–268. [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F) (1992).
29. Myers, G. R., Kingston, A. M., Varslot, T. K., Turner, M. L. & Sheppard, A. P. Dynamic tomography with a priori information. *Appl. Opt.* **50**, 3685. <https://doi.org/10.1364/ao.50.003685> (2011).
30. Myers, G. R., Geleta, M., Kingston, A. M., Recur, B. & Sheppard, A. P. Bayesian approach to time-resolved tomography. *Opt. Express* **23**, 20062. <https://doi.org/10.1364/OE.23.020062> (2015).
31. Batenburg, K. J. & Sijbers, J. DART: a practical reconstruction algorithm for discrete tomography. *IEEE Trans. Image Process. Publ. IEEE Signal Processing Soc.* **20**, 2542–2553. <https://doi.org/10.1109/TIP.2011.2131661> (2011).
32. Zhuge, X., Palenstijn, W. J. & Batenburg, K. J. TVR-DART: A more robust algorithm for discrete tomography from limited projection data with automated gray value estimation. *IEEE Trans. Image Process.* **25**, 455–468. <https://doi.org/10.1109/TIP.2015.2504869> (2016).
33. Van Eyndhoven, G. *et al.* An iterative CT reconstruction algorithm for fast fluid flow imaging. *IEEE Trans. Image Process.* **24**, 4446–4458. <https://doi.org/10.1109/TIP.2015.2466113> (2015).
34. Hansen, P. C. *Rank-Deficient and Discrete Ill-Posed Problems* (Society for Industrial and Applied Mathematics, 1998).
35. Hara, A. K. *et al.* Iterative reconstruction technique for reducing body radiation dose at CT: Feasibility study. *Am. J. Roentgenol.* **193**, 764–771. <https://doi.org/10.2214/AJR.09.2397> (2009).
36. Brabant, L., Vlassenbroeck, J., De Schryver, T., Boone, M. & Van Hoorebeke, L. The advantages of iterative reconstruction methods for high resolution x-ray tomography. In *12th International conference on X-Ray Microscopy: conference program handbook*, 273–273. hdl.handle.net/1854/LU-5778935 (2014).
37. Gregor, J. & Benson, T. Computational analysis and improvement of SIRT. *IEEE Trans. Med. Imag.* **27**, 918–924. <https://doi.org/10.1109/TMI.2008.923696> (2008).
38. Elfving, T., Hansen, P. C. & Nikazad, T. Semiconvergence and relaxation parameters for projected SIRT algorithms. *SIAM J. Sci. Comput.* **34**, A2000–A2017. <https://doi.org/10.1137/110834640> (2012).
39. Palenstijn, W., Batenburg, K. & Sijbers, J. Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs). *J. Struct. Biol.* **176**, 250–253. <https://doi.org/10.1016/j.jsb.2011.07.017> (2011).
40. van Aarle, W. *et al.* The ASTRA toolbox: a platform for advanced algorithm development in electron tomography. *Ultramicroscopy* **157**, 35–47. <https://doi.org/10.1016/j.ultramic.2015.05.002> (2015).
41. van Aarle, W. *et al.* Fast and flexible X-ray tomography using the ASTRA toolbox. *Opt. Express* **24**, 25129–25147. <https://doi.org/10.1364/OE.24.025129> (2016).
42. Rust, B. W. Parameter selection for constrained solutions to ill-posed problems. *Comput. Sci. Stat.* **32**, 333–347 (2000).
43. Hansen, P. C., Kilmer, M. E. & Kjeldsen, R. H. Exploiting residual information in the parameter choice for discrete ill-posed problems. *BIT Numer. Math.* **46**, 41–59. <https://doi.org/10.1007/s10543-006-0042-7> (2006).
44. Rust, B. W. & O’Leary, D. P. Residual periodograms for choosing regularization parameters for ill-posed problems. *Inverse Prob.* **24**, 034005. <https://doi.org/10.1088/0266-5611/24/3/034005> (2008).
45. Hansen, P. C., Sørensen, H. O., Sükösd, Z. & Poulsen, H. F. Reconstruction of single-grain orientation distribution functions for crystalline materials. *SIAM J. Imag. Sci.* **2**, 593–613. <https://doi.org/10.1137/080726021> (2009).
46. Okuta, R., Unno, Y., Nishino, D., Hido, S. & Loomis, C. C. Cupy: CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, http://learningsys.org/nips17/assets/papers/paper_16.pdf (2017).
47. Ahrens, J., Geveci, B. & Law, C. *ParaView: An End-User Tool for Large Data Visualization, Visualization Handbook* (Elsevier, 2005).
48. Takeuchi, A., Uesugi, K. & Suzuki, Y. Zernike phase-contrast x-ray microscope with pseudo-Kohler illumination generated by sectorized (polygon) condenser plate. *J. Phys. Conf. Ser.* **186**, 012020. <https://doi.org/10.1088/1742-6596/186/1/012020> (2009).
49. Chitalya, R., Hoffmann, K. R., Rudin, S. & Bednarek, D. R. Artifact reduction in truncated CT using Sinogram completion. *Proc. SPIE Int. Soc. Opt. Eng.* <https://doi.org/10.1117/12.595450> (2005).
50. Münch, B., Trtik, P., Marone, F. & Stampanoni, M. Stripe and ring artifact removal with combined wavelet - Fourier filtering. *Opt. Express* **17**, 8567–8591. <https://doi.org/10.1364/OE.17.008567> (2009).
51. Gürsoy, D., De Carlo, F., Xiao, X. & Jacobsen, C. TomoPy: a framework for the analysis of synchrotron tomographic data. *J. Synchrotron Radiat.* **21**, 1188–1193. <https://doi.org/10.1107/S1600577514013939> (2014).
52. Bruns, S., Stipp, S. L. S. & Sørensen, H. O. Looking for the signal: a guide to iterative noise and artefact removal in X-ray tomographic reconstructions of porous geomaterials. *Adv. Water Resour.* **105**, 96–107. <https://doi.org/10.1016/j.advwatres.2017.04.020> (2017).
53. Fakhari, A., Mitchell, T., Leonardi, C. & Bolster, D. Improved locality of the phase-field lattice-Boltzmann model for immiscible fluids at high density ratios. *Phys. Rev. E* **96**, 053301. <https://doi.org/10.1103/PhysRevE.96.053301> (2017).
54. Fakhari, A., Bolster, D. & Luo, L.-S. A weighted multiple-relaxation-time lattice Boltzmann method for multiphase flows and its application to partial coalescence cascades. *J. Comput. Phys.* **341**, 22–43. <https://doi.org/10.1016/j.jcp.2017.03.062> (2017).
55. Colton, D. & Kress, R. *Inverse Acoustic Obstacle Scattering* Vol. 93 (Springer, New York, 2019).
56. Gravel, P., Beaudoin, G. & De Guise, J. A. A method for modeling noise in medical images. *IEEE Trans. Med. Imag.* **23**, 1221–1232. <https://doi.org/10.1109/TMI.2004.832656> (2004).

57. Whiting, B. R. *et al.* Properties of preprocessed sinogram data in x-ray computed tomography. *Med. Phys.* **33**, 3290–3303. <https://doi.org/10.1118/1.2230762> (2006).
58. Buzug, T. *Computed tomography: From photon statistics to modern cone-beam CT* (Springer, Berlin, 2008).
59. Ma, J. *et al.* Variance analysis of x-ray CT sinograms in the presence of electronic noise background. *Med. Phys.* **39**, 4051–4065. <https://doi.org/10.1118/1.4722751> (2012).
60. Hunter, J. D. Matplotlib: a 2d graphics environment. *Comput. Sci. Eng.* **9**, 90–95. <https://doi.org/10.1109/MCSE.2007.55> (2007).

Acknowledgements

We would like to thank the Danish Hydrocarbon Research and Technology Centre for financial support and the Danish Agency for Science, Technology, and Innovation for funding the instrument center DanScatt. We thank the Japan Synchrotron Radiation Research Institute for the allotment of beam time on beamline BL47XU of SPring-8 (Proposal 2016A1459) and N. Bovet and D. Mütter for assisting at the experiment. We are grateful to P. C. Hansen for discussions related to the NCP stopping rule. We also thank the Center for Quantification of Imaging Data from MAX IV (QIM) funded by the Capital Region of Denmark.

Author contributions

P.W.R. got the idea of using a static prior image to constrain a dynamic reconstruction, performed the numerical experiment and wrote the first draft of the manuscript. P.W.R., H.O.S., A.B.D., A.N.C. designed the numerical experiment and worked extensively on the manuscript. S.B. performed the simulation used for the numerical experiment. All authors reviewed the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-021-91776-1>.

Correspondence and requests for materials should be addressed to P.W.R. or A.N.C.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021

CHAPTER 6

Contribution B: Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography

Contribution B is in many ways a result of the work done during **Contribution A**. We needed an automatic way to terminate our reconstructions during that project as the optimal number of iterations for our reconstruction algorithms varied with data quality. I had previously learned of different “stopping rules” during a PhD course on computed tomography by Per Christian Hansen. We, therefore, established contact with him, and he provided us with the literature needed to implement the NPC stopping rule used in **Contribution A**. We met again during a summer school on inverse problems where we discussed writing a paper on different stopping rules, which eventually resulted in **Contribution B**.

In **Contribution B**, we sought to compare multiple stopping rules both using the ubiquitous Shepp-Logan phantom but also using a large scale example more akin to what is seen in modern tomography (Hansen, Jørgensen, and Rasmussen, 2021). Implementing the stopping rules such that they could be applied to large problems constituted the majority of my contribution. The paper was submitted to the 21st edition Computational Science and Its Applications (ICCSA) Conference in April 2021 and was accepted after peer-review in June 2021.

Unfortunately, delays at the publisher mean that the work has not been published yet. The version of the paper submitted to arXiv.org is, therefore, included below. It can also be found online at <https://arxiv.org/abs/2106.10053>.

Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography

Per Christian Hansen, Jakob Sauer Jørgensen, Peter Winkel Rasmussen
Department of Applied Mathematics and Computer Science
Technical University of Denmark, Kgs. Lyngby, Denmark
ORCID 0000-0002-7333-7216, 0000-0001-9114-754X, 0000-0002-0823-0316
Email {pcha, jakj, pwra}@dtu.dk

Abstract—Algebraic models for the reconstruction problem in X-ray computed tomography (CT) provide a flexible framework that applies to many measurement geometries. For large-scale problems we need to use iterative solvers, and we need stopping rules for these methods that terminate the iterations when we have computed a satisfactory reconstruction that balances the reconstruction error and the influence of noise from the measurements. Many such stopping rules are developed in the inverse problems communities, but they have not attained much attention in the CT world. The goal of this paper is to describe and illustrate four stopping rules that are relevant for CT reconstructions.

Index Terms—tomographic reconstruction, iterative methods, stopping rules, semi-convergence

I. INTRODUCTION

This paper considers large-scale methods for computed tomographic (CT) based on a discretization of the problem followed by solving the system of linear equations by means of an iterative solver. These methods are quite generic in the sense that they do not assume any specific scanning geometry, and they tend to produce good reconstructions with few artifacts in the case of limited-data and/or limited-angle problems.

In CT, a forward projection maps the object to the data in the form of projections of the object onto the detector planes for various scan positions. In the case of parallel-beam CT the forward projection is known as the Radon transform [21].

In practise, data consists of noisy measurements of the attenuation of the X-rays that pass through the object. The discretization of the reconstruction problem takes the form

$$\mathbf{A}\mathbf{x} \approx \mathbf{b}, \quad \mathbf{b} = \mathbf{A}\bar{\mathbf{x}} + \mathbf{e}, \quad (1)$$

where the “system matrix” $\mathbf{A} \in \mathbb{R}^{m \times n}$ is a discretization of the forward projector, $\mathbf{b} \in \mathbb{R}^m$ is a vector with the measured data, and $\mathbf{x} \in \mathbb{R}^n$ is a vector that holds the pixels of the reconstructed image of the object’s interior. Moreover, $\bar{\mathbf{x}}$ represents the exact object and \mathbf{e} represents the measurement noise. A number of discretization schemes are available for computing the matrix \mathbf{A} , see, e.g., [15], [17].

There are no restrictions on the dimensions m and n of the matrix \mathbf{A} ; both over-determined and under-determined systems

are common, depending on the measurement setup. The matrix \mathbf{A}^T represents the so-called back projector which maps the data back onto the solution domain [21]; it plays a central role in filtered back projection and similar methods.

In large-scale CT problems, the matrix \mathbf{A} – in spite of the fact that it is sparse – is often too large to store explicitly. For this reason we must use iterative solvers that only access the matrix via functions that compute the multiplications with \mathbf{A} and \mathbf{A}^T in a matrix-free fashion, often using GPUs or other hardware accelerators. In CT these iterative solvers are collectively referred to as *algebraic iterative reconstruction methods* which includes well-known methods such as ART [11] and SIRT (also known as SART) [3].

Common for all these methods is that they, from an initial vector $\mathbf{x}^{(0)}$ (often the zero vector) produce a sequence of iteration vectors $\mathbf{x}^{(k)}$, $k = 1, 2, \dots$ which, in the ideal situation, converge to the ground truth $\bar{\mathbf{x}}$. In practise, however, when noise is present in the measured data we experience a phenomenon called *semi-convergence*:

- During the initial iterations, the iteration vector $\mathbf{x}^{(k)}$ approaches the desired but un-obtainable ground truth $\bar{\mathbf{x}}$.
- During later iterations, $\mathbf{x}^{(k)}$ converges to the undesired noisy solution associated with the particular iterative method (e.g., $\mathbf{A}^{-1}\mathbf{b}$ if the system matrix is invertible).

This is illustrated in Fig. 1 which shows the error history, i.e., the reconstruction error $\|\mathbf{x}^{(k)} - \bar{\mathbf{x}}\|_2$ versus the number of iterations k , together with selected iterates $\mathbf{x}^{(k)}$ shown as images. The error history has the characteristic form associated with semi-convergence.

If we can stop the iterations just when the convergence behavior changes from the former to the latter, then we achieve an approximation to $\bar{\mathbf{x}}$ that is not too perturbed by the noise in the data. *This paper describes four such methods based on certain statistical properties of the noise.*

Sections II and III describe four state-of-the-art stopping rules as well as two methods to efficiently estimate a trace-term that is needed on some of these rules; all numerical experiments in these sections were performed by means of the AIR Tools II software package [14]. In Section IV we illustrate these techniques with a large-scale example.

This work was partially funded by a Villum Investigator grant (no. 25893) from The Villum Foundation.

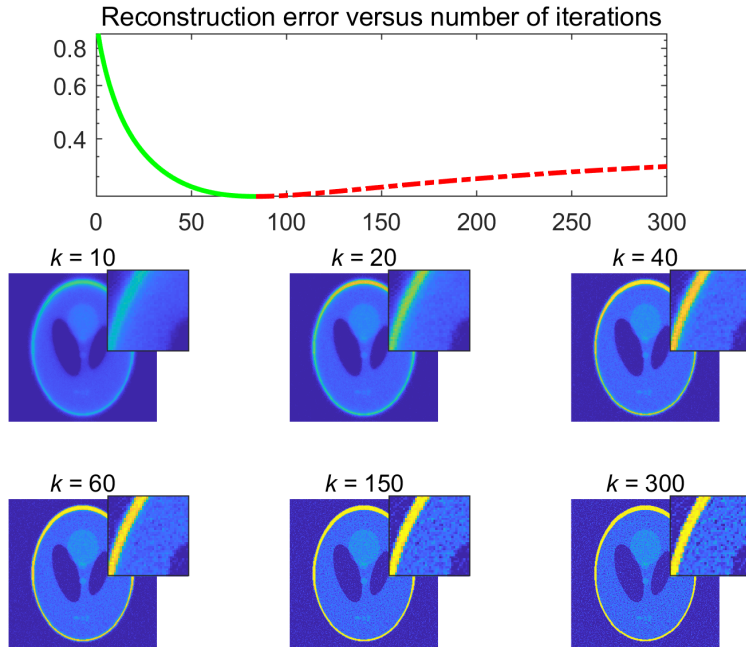


Fig. 1. Illustration of semi-convergence for Landweber's method applied to a small noisy test problem. Top: the error history, i.e., the reconstruction error $\|\mathbf{x}^{(k)} - \bar{\mathbf{x}}\|_2$ as a function of the number of iterations k . The solid green line shows the part where $\mathbf{x}^{(k)}$ approaches $\bar{\mathbf{x}}$ while the red dash-dotted line shows the part where $\mathbf{x}^{(k)}$ becomes dominated by the noise. Bottom: selected iterations with inserts that zoom in on a small region; the increasing amount of noise is clearly visible.

II. FOUR STOPPING RULES

Our stopping rules apply to methods of the general form

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \mathbf{D} \mathbf{A}^T \mathbf{M} (\mathbf{b} - \mathbf{A} \mathbf{x}^{(k)}), \quad (2)$$

where different choices of the diagonal matrices \mathbf{D} and \mathbf{M} lead to different methods – see, e.g., [5] for an overview. To simplify the presentation, we focus on the simple case where \mathbf{D} and \mathbf{M} are identity matrices, in which case we obtain *Landweber's method* (which is equivalent to the steepest descent method applied to the least squares problem):

Landweber's method

$\mathbf{x}^{(0)}$ = initial vector

for $k = 0, 1, 2, \dots$

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \omega \mathbf{A}^T (\mathbf{b} - \mathbf{A} \mathbf{x}^{(k)})$$

end

We will frequently refer to the *residual* for the k th iterate, defined as the vector

$$\mathbf{r}^{(k)} = \mathbf{b} - \mathbf{A} \mathbf{x}^{(k)}, \quad k = 1, 2, 3, \dots \quad (3)$$

Moreover, after a bit of algebraic manipulations it turns out that we can write the k th iteration vector as

$$\mathbf{x}^{(k)} = \sum_{j=0}^{k-1} (\mathbf{I} - \omega \mathbf{A}^T \mathbf{A})^j \omega \mathbf{A}^T \mathbf{b} = \mathbf{A}_k^\# \mathbf{b}, \quad (4)$$

which defines the matrix $\mathbf{A}_k^\#$ such that we can write the k th iterate as $\mathbf{x}^{(k)} = \mathbf{A}_k^\# \mathbf{b}$. This is convenient as a theoretical tool, but $\mathbf{A}_k^\#$ is never computed explicitly.

To set the stage, we need to introduce a small amount of statistical framework and notation. We will often need the exact *noise-free data* that corresponds to the ground truth image, and we write

$$\bar{\mathbf{b}} = \mathbf{A} \bar{\mathbf{x}}. \quad (5)$$

We can then write the measured data as $\mathbf{b} = \bar{\mathbf{b}} + \mathbf{e}$. The elements of the noise vector \mathbf{e} are random variables, i.e., their values depend on a set of well-defined random events. The vector of expected values $\mathcal{E}(\mathbf{e})$ and the covariance matrix $\text{Cov}(\mathbf{e})$ are defined as

$$\mathcal{E}(\mathbf{e}) = \begin{pmatrix} \mathcal{E}(e_1) \\ \mathcal{E}(e_2) \\ \vdots \end{pmatrix}, \quad (6)$$

$$\text{Cov}(\mathbf{e}) = \mathcal{E} \left((\mathbf{e} - \mathcal{E}(\mathbf{e})) (\mathbf{e} - \mathcal{E}(\mathbf{e}))^T \right). \quad (7)$$

To simplify our discussion and make the ideas clearer, throughout this section we will restrict our analysis to *white Gaussian noise* with zero mean:

$$\mathcal{E}(e) = \mathbf{0}, \quad \text{Cov}(e) = \eta^2 \mathbf{I}, \quad \mathcal{E}(\|e\|_2^2) = m\eta^2, \quad (8)$$

where η is the standard deviation of the noise and m is the number of elements in e . Noise in tomographic problems is rarely strictly Gaussian, but sometimes this is a reasonable assumption.

A. Fitting to the Noise Level

Our description of this stopping rule is based on [15, §11.2.3]. A simple idea is to choose the number of iterations k such that the residual $\boldsymbol{\varrho}^{(k)}$ is “of the same size” as the noise vector e . Specifically, such that $\|\boldsymbol{\varrho}^{(k)}\|_2$ approximates the expected value $\mathcal{E}(\|e\|_2)$ of the latter:

$$\|\boldsymbol{\varrho}^{(k)}\|_2 \approx \eta\sqrt{m}. \quad (9)$$

In the literature this is referred to as the *discrepancy principle* [6]. Since $\|\boldsymbol{\varrho}^{(k)}\|_2$ takes discrete values for $k = 1, 2, 3, \dots$ we cannot expect to find a k such that the above holds with equality.

It is common to include a constant τ slightly larger than 1, say, $\tau = 1.02$, such that the above condition takes the form $\|\boldsymbol{\varrho}^{(k)}\|_2 \leq \tau\eta\sqrt{m}$. This constant can be useful as a “safety factor” when we have only a rough estimate of the noise.

If we replace $\boldsymbol{x}^{(k)}$ with the ground truth $\bar{\boldsymbol{x}}$ then the residual is $\mathbf{b} - \mathbf{A}\bar{\boldsymbol{x}} = \mathbf{e}$ and the residual norm obviously equals $\|e\|_2$. However, this is not a sound statistical argument that the norm of the residual $\boldsymbol{\varrho}^{(k)}$ in Eq. (3) should be equal to $\|e\|_2$ for the optimal iterate $\boldsymbol{x}^{(k)}$.

Here we present an alternative that is based on statistical principles. To motivate this stopping rule, we split the residual vectors as follows:

$$\begin{aligned} \boldsymbol{\varrho}^{(k)} &= \mathbf{b} - \mathbf{A}\boldsymbol{x}^{(k)} = \mathbf{b} - \mathbf{A}\mathbf{A}_k^\# \bar{\mathbf{b}} \\ &= (\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}} + (\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}. \end{aligned}$$

The heuristic insight is then as follows:

- When k is too small then $\mathbf{A}\boldsymbol{x}^{(k)}$ is not a good approximation the exact data $\bar{\mathbf{b}}$. Hence, the residual $\boldsymbol{\varrho}^{(k)}$ is dominated by $(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}}$ and $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}}\|_2$ is larger than $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}\|_2$.
- When k is “just about right” then $\mathbf{A}\boldsymbol{x}^{(k)}$ approximates $\bar{\mathbf{b}}$ as well as possible; the norm $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}}\|_2$ has now become smaller and it is of the same size as the norm $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}\|_2$.
- When k is too large then the residual $\boldsymbol{\varrho}^{(k)}$ is dominated by the noise component $(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}$, and therefore $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}\|_2$ dominates the residual norm.

According to these observations we should therefore choose k such that $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}}\|_2 \approx \|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}\|_2$. Unfortunately both these are unknown.

The above heuristic reasoning has been formalized in [13], [18] and [25], and we will summarize the main results as they

apply here. The key points are that $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \bar{\mathbf{b}}\|_2$ is an approximation to the prediction error $\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{b}\|_2$ and that

$$\mathcal{E}(\|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#) \mathbf{e}\|_2^2) = \eta^2(m - t_k)$$

in which

$$t_k = \text{trace}(\mathbf{A}\mathbf{A}_k^\#). \quad (10)$$

Hence, at the optimal k we have

$$\mathcal{E}(\|\boldsymbol{\varrho}^{(k)}\|_2^2) \approx \eta^2(m - t_k). \quad (11)$$

Here, k is “optimal” in the sense that it is the largest iteration number for which we cannot reject $\boldsymbol{x}^{(k)}$ – computed from the noisy data \mathbf{b} – as a possible solution to the noise-free system, cf. [25, p. 93].

The real number $m - t_k$ is sometimes referred to as the effective (or equivalent) degrees of freedom [27] in the residual. An exact computation of t_k is cumbersome for most methods, but it can be approximated quite efficiently as described in §III. We have thus arrived at the following “fit-to-noise-level” (FTNL) stopping rule where, again, we include the “safety factor” τ :

Stop rule: FTNL

Stop at the smallest k

for which $\|\boldsymbol{\varrho}^{(k)}\|_2 \leq \tau\eta\sqrt{m - t_k}$.

Example 1. We illustrate the FTNL “fit-to-noise-level” stopping rule with two small parallel-beam CT problems with image size 64×64 and 91 detector pixels. The projection angles are, respectively, $3^\circ, 6^\circ, 9^\circ, \dots, 180^\circ$ (giving an over-determined system) and $8^\circ, 16^\circ, 24^\circ, \dots, 180^\circ$ (giving an under-determined system). In both cases we removed zero rows from the system matrix.

We used Landweber’s method to solve these two problems. Figure 2 shows the reconstruction errors $\|\boldsymbol{x}^{(k)} - \bar{\boldsymbol{x}}\|_2$ and the norms $\|\boldsymbol{\varrho}^{(k)}\|_2$ versus k , together with the threshold $\eta\sqrt{m}$ and the function $\eta\sqrt{m - t_k}$, i.e., here we use $\tau = 1$. The graphs confirm the monotonic decrease of the residual norm. For both problems, the “fit-to-noise-level” stopping rule terminates the iterations close to the optimal number of iterations. A stopping rule involving $\eta\sqrt{m}$, on the other hand, would terminate the iterations much too early. \square

B. Minimization of the Prediction Error – UPRE

The key idea is to find the number of iterations that minimizes the prediction error, i.e., the difference between the noise-free data $\bar{\mathbf{b}} = \mathbf{A}\bar{\boldsymbol{x}}$ and the predicted data $\mathbf{A}\boldsymbol{x}^{(k)}$. Statisticians refer to various measures of this difference as the predictive risk, and the resulting method for choosing k is often called the *unbiased predictive risk estimation (UPRE)* method.

Here we present the results specifically in the framework of iterative reconstruction methods and using the matrix $\mathbf{A}_k^\#$ defined in Eq. (4). Following [26, §7.1], where all the details




Fig. 2. Illustration of the FTNL stopping rule for Landweber's method, with two parallel-beam tomographic problems. The smallest reconstruction error is marked with the black dot, and the residual norms that satisfies the stopping rules are marked with red circles. The FTNL rule works well, while stopping at that k for which $\|\mathbf{q}^{(k)}\|_2 \approx \eta\sqrt{m}$ terminates the iterations much too early.

can be found, the expected squared norm of the prediction error (the risk) is

$$\mathcal{E}(\|\bar{\mathbf{b}} - \mathbf{A}\mathbf{x}^{(k)}\|_2^2) = \|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#)\bar{\mathbf{b}}\|_2^2 + \eta^2 \text{trace}((\mathbf{A}\mathbf{A}_k^\#)^2)$$

while the expected squared norm of the residual can be written as

$$\mathcal{E}(\|\mathbf{b} - \mathbf{A}\mathbf{x}^{(k)}\|_2^2) = \|(\mathbf{I} - \mathbf{A}\mathbf{A}_k^\#)\bar{\mathbf{b}}\|_2^2 + \eta^2 \text{trace}((\mathbf{A}\mathbf{A}_k^\#)^2) - 2\eta^2 \text{trace}(\mathbf{A}\mathbf{A}_k^\#) + \eta^2 m .$$

Combining these two equations we can eliminate one of the trace terms and arrive at the following expression for the risk:

$$\mathcal{E}(\|\bar{\mathbf{b}} - \mathbf{A}\mathbf{x}^{(k)}\|_2^2) = \mathcal{E}(\|\mathbf{b} - \mathbf{A}\mathbf{x}^{(k)}\|_2^2) + 2\eta^2 \text{trace}(\mathbf{A}\mathbf{A}_k^\#) - \eta^2 m .$$

Substituting the actual squared residual norm $\|\mathbf{q}^{(k)}\|_2^2 = \|\mathbf{b} - \mathbf{A}\mathbf{x}^{(k)}\|_2^2$ for its expected value, we thus define the UPRE risk as a function of k :

$$U^{(k)} = \|\mathbf{q}^{(k)}\|_2^2 + 2\eta^2 t_k - \eta^2 m \quad (12)$$

with t_k given by (10). A minimizer of $U^{(k)}$ will then give an approximation to a minimizer of the prediction error. We note that $U^{(k)}$ may not have a unique minimizer, and we therefore choose the smallest k at which $U^{(k)}$ has a local minimum. Thus we arrive at the following stopping rule:

Stop rule: UPRE

$$\text{Minimize } U^{(k)} = \|\mathbf{q}^{(k)}\|_2^2 + 2\eta^2 t_k - \eta^2 m .$$

C. Another Rule Based on the Prediction Error – GCV

The above UPRE stopping rule depends on an estimate of the standard deviation η of the noise – which may or may not be a problem in practise. We shall now describe an alternative method for minimization of the prediction error, derived by Wahba [27], that does not depend on knowledge of η .

The outset for this method is the principle of cross validation. Assume that we remove the i th element b_i from the right-hand side (the noisy data), compute a reconstruction $\mathbf{x}_{[i]}^{(k)}$, and then use this reconstruction to compute a prediction $\hat{b}_i = \mathbf{r}_i^T \mathbf{x}_{[i]}^{(k)}$ of the missing data b_i , where

$$\mathbf{r}_i^T = \mathbf{A}(i, :) = i\text{th row of } \mathbf{A} .$$

The goal would then be to choose the number of iterations k that minimizes the mean of all the squared prediction errors:

$$\hat{G}^{(k)} = \frac{1}{m} \sum_{i=1}^m (b_i - \hat{b}_i)^2 = \frac{1}{m} \sum_{i=1}^m (b_i - \mathbf{r}_i^T \mathbf{x}_{[i]}^{(k)})^2 .$$

Then it is proved in [27, Thm. 4.2.1] that we can avoid the vectors $\mathbf{x}_{[i]}^{(k)}$ and write $\hat{G}^{(k)}$ directly in terms of $\mathbf{x}^{(k)}$:

$$\hat{G}^{(k)} = \frac{1}{m} \sum_{i=1}^m \left(\frac{b_i - \mathbf{r}_i^T \mathbf{x}^{(k)}}{1 - \alpha_i^{(k)}} \right)^2, \quad (13)$$

where $\alpha_i^{(k)}$ is the i th diagonal element of the matrix product $\mathbf{A}\mathbf{A}_k^\#$ associated with $\mathbf{x}^{(k)}$.

At this stage, recall that the 2-norm is invariant under an orthogonal transformation, of which a permutation is a special case. Specifically, if \mathbf{Q} is an orthogonal matrix then

$$\|\mathbf{Q}(\mathbf{A}\mathbf{x} - \mathbf{b})\|_2 = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$$

which means that the reconstruction $\mathbf{x}^{(k)}$ is invariant to such a transformation. Unfortunately it can be proved [27] that the minimizer of $\hat{G}^{(k)}$ is not invariant to an orthogonal transformation of the data. In particular, it is inconvenient that a stopping rule based on $\hat{G}^{(k)}$ would produce a k that depends on the particular ordering of the data.

The *generalized cross validation (GCV)* method circumvents this problem by replacing all $\alpha_i^{(k)}$ with their average

$$\mu^{(k)} = \frac{1}{m} \sum_{i=1}^m \alpha_i^{(k)} = \frac{1}{m} \text{trace}(\mathbf{A}\mathbf{A}_k^\#) = \frac{t_k}{m} ,$$

leading to the modified measure

$$\begin{aligned}\tilde{G}^{(k)} &= \frac{1}{m} \frac{1}{(1 - \mu^{(k)})^2} \sum_{i=1}^m (b_i - \mathbf{r}_i^T \mathbf{x}^{(k)})^2 \\ &= \frac{\|\bar{\mathbf{b}} - \mathbf{A} \mathbf{x}^{(k)}\|_2^2}{m(1 - t_k/m)^2} = m \frac{\|\boldsymbol{\varrho}^{(k)}\|_2^2}{(m - t_k)^2}.\end{aligned}\quad (14)$$

The minimizer of $\tilde{G}^{(k)}$ is, of course, independent of the factor m and hence we choose to define the GCV risk as a function of k as

$$G^{(k)} = \frac{\|\boldsymbol{\varrho}^{(k)}\|_2^2}{(m - t_k)^2}.\quad (15)$$

We have thus arrived at the following η -free stopping rule where again, in practice, we need to estimate the quantity t_k :

Stop rule: GCV

$$\text{Minimize } G^{(k)} = \|\boldsymbol{\varrho}^{(k)}\|_2^2 / (m - t_k)^2.$$

The above presentation follows [27, §4.2–3]. A different derivation of the GCV method was presented in [10]; here the coordinate system for \mathbb{R}^m is rotated such that the corresponding influence matrix becomes a circulant matrix with identical elements along all its diagonals. This approach leads to the same GCV risk $G^{(k)}$ as above.

Perhaps the most important property of the GCV stopping rule is that the value of k which minimizes $G^{(k)}$ in (15) is also an estimate of the value that minimizes the prediction error. Specifically, if k_{GCV} minimizes the GCV risk $G^{(k)}$ and k_{PE} minimizes the prediction error $\|\bar{\mathbf{b}} - \mathbf{A} \mathbf{x}^{(k)}\|_2^2$, then it is shown in [27, §4.4] that

$$\begin{aligned}\mathcal{E}(\|\bar{\mathbf{b}} - \mathbf{A} \mathbf{x}^{(k_{\text{GCV}})}\|_2^2) &\rightarrow \mathcal{E}(\|\bar{\mathbf{b}} - \mathbf{A} \mathbf{x}^{(k_{\text{PE}})}\|_2^2) \\ &\text{for } m \rightarrow \infty.\end{aligned}$$

The UPRE and GCV stopping rules have the slight inconvenience that we need to take at least one iteration too many in order to detect a minimum of $U^{(k)}$ and $G^{(k)}$, respectively. In practise, this is not really a problem. For tomography problems the iteration vector $\mathbf{x}^{(k)}$ does not change very much from one iteration to the next, and hence the minimum of the error history $\|\bar{\mathbf{x}} - \mathbf{x}^{(k)}\|_2$ is usually very flat. Hence it hardly makes any difference if we implement the UPRE and GCV stopping rules such that we terminate the method one iteration (or a few iterations) after the actual minimum of $U^{(k)}$ or $G^{(k)}$.

Example 2. We illustrate the UPRE and GCV stopping rules applied to Landweber’s method with the two CT problems from Example 1. In both cases we removed zero rows from the system matrix. Figure 3 shows $U^{(k)}$ and $G^{(k)}$ from Eqs. (12) and (15) versus k , together with the error histories. The two stopping rules terminate the iterations at approximately the same number of iterations – not too far from the minimum of the error history. Note how flat the error history is: in practise it makes no difference if we terminate the iterations exactly at the minimum of $U^{(k)}$ and $G^{(k)}$ or a few iterations later. \square

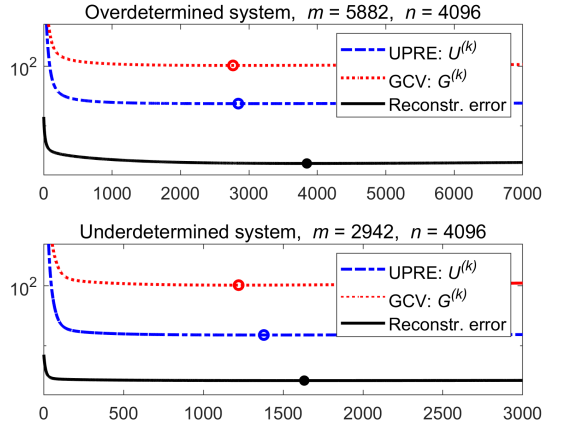


Fig. 3. Illustration of the UPRE and GCV stopping rules for Landweber’s method applied to the two parallel-beam CT problems in Example 1 and 2.

D. Stopping When All Information is Extracted — NCP

The above stopping rules include the trace term t_k in Eq. (10). This term can be estimated at additional cost as discussed in §III below, but it is also worthwhile to consider a stopping rule that needs neither the trace term t_k nor the standard deviation η of the noise. The so-called NCP criterion from [16] and [23] is one such method. The considerations that underly this method are as follows:

- 1) noisy data only contain partial information about the reconstruction,
- 2) in each iteration we extract more information from the data, and
- 3) eventually we have extracted all the available information in the noisy data.

Therefore we want to monitor the properties of the residual vector. During the initial iterations we have not yet extracted all information present in the data and the residual still resembles a meaningful signal, while at some stage – when all information is extracted – the residual starts to appear like noise. When we iterate beyond this point, we solely extract noise from the data (we “fit the noise”) and the residual vector will appear as filtered noise where some of the noise’s spectral components are removed.

To formalize this approach, in the white-noise setting of this presentation, we need a computational approach to answering the questions: when does the residual vector look the most like white noise? To answer this question, statisticians introduced the so-called normalized cumulative periodogram.

In the terminology of signal processing, a periodogram is identical to a discrete power spectrum defined as the squared absolute values of the discrete Fourier coefficients. Hence the periodogram for an arbitrary vector $\mathbf{v} \in \mathbb{R}^m$ is given by

$$\hat{p}_i = |\hat{v}_i|^2, \quad i = 1, 2, \dots, q, \quad \hat{\mathbf{v}} = \text{DFT}(\mathbf{v}).\quad (16)$$




Fig. 4. Illustration of NCP vectors $\mathbf{c}(\mathbf{v}) \in \mathbb{R}^{256}$ for vectors \mathbf{v} that are white noise (left), dominated by low-frequency components (middle), and dominated by high-frequency components (right).

Here, DFT denotes the discrete Fourier transform (computed by means of the FFT algorithm) and $q = \lfloor m/2 \rfloor$ denotes the largest integer such that $q \leq m/2$. The reason for including only about half of the Fourier coefficients in the periodogram/power spectrum is that the DFT of a real vector is symmetric about its midpoint. We then define the corresponding *normalized cumulative periodogram (NCP)* for the vector \mathbf{v} as the vector $\mathbf{c}(\mathbf{v})$ of length q with elements, for $j = 1, 2, \dots, q$,

$$c_j(\mathbf{v}) = \frac{\widehat{p}_2 + \dots + \widehat{p}_{j+1}}{\widehat{p}_2 + \dots + \widehat{p}_{q+1}} = \frac{\|\widehat{\mathbf{v}}_{2:j+1}\|_2^2}{\|\widehat{\mathbf{v}}_{2:q+1}\|_2^2}. \quad (17)$$

White noise is characterized by having a flat power spectrum (similar to white light having equal amounts of all colors), and thus the expected value of its power spectrum components is a constant independent of i . Consequently, the expected value of the NCP for a white-noise vector \mathbf{v}_w is the vector

$$\mathcal{E}(\mathbf{c}(\mathbf{v}_w)) = \mathbf{c}_w = \left(\frac{1}{q}, \frac{2}{q}, \dots, 1 \right).$$

How much a given vector \mathbf{v} deviates from being white noise can be measured by the deviation of the corresponding $\mathbf{c}(\mathbf{v})$ from \mathbf{c}_w , e.g., as measured by the norm $\|\mathbf{c}(\mathbf{v}) - \mathbf{c}_w\|_2$.

Example 3. Figure 4 illustrates the appearance of NCP vectors $\mathbf{c}(\mathbf{v})$ for vectors \mathbf{v} of length $m = 256$ with different spectra. The completely flat spectrum for white noise corresponds to a straight line from $(0, 0)$ to $(q, 1)$ with $q = \lfloor 256/2 \rfloor = 128$. The left plot shows NCPs for 10 random realizations of white noise, and they are all close to the ideal white-noise NCP \mathbf{c}_w . The middle and right plots show NCPs for random vectors that are dominated by low-frequency and high-frequency components, respectively; their systematic deviation from \mathbf{c}_w is obvious. \square

To utilize the NCP framework in the algebraic iterative methods for tomographic reconstruction, a first idea might be to terminate the iterations when the deviation measured by $\|\mathbf{c}(\mathbf{q}^{(k)}) - \mathbf{c}_w\|_2$ exhibits a minimum. However, this would be a bit naive since the residual vector does not really correspond to a 1D signal of length m . Rather, the right-hand side \mathbf{b} consists of a number of projections, one for each angle of the measurements – and the residual vector inherits this structure. Hence, a better approach is to apply an NCP analysis to each projection’s residual, and then combine this information into a simple measure.

Depending on the CT scanner, each projection is either a 1D or 2D image, when we perform 2D and 3D reconstructions,

respectively. To simplify our presentation, we assume that our data consists of m_θ 1D projections, one for each projection angle $\theta_1, \theta_2, \dots, \theta_{m_\theta}$. We also assume that the data are organized such that we can partition the right-hand side \mathbf{b} and the residual vector into m_θ sub-vectors,

$$\mathbf{b} = \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_{m_\theta} \end{pmatrix}, \quad \mathbf{q}^{(k)} = \begin{pmatrix} \mathbf{q}_1^{(k)} \\ \mathbf{q}_2^{(k)} \\ \vdots \\ \mathbf{q}_{m_\theta}^{(k)} \end{pmatrix}, \quad (18)$$

with each sub-vector corresponding to a single 1D projection. Now define the corresponding quantities

$$\nu_\ell^{(k)} = \|\mathbf{c}(\mathbf{q}_\ell^{(k)}) - \mathbf{c}_w\|_2, \quad \ell = 1, 2, \dots, m_\theta \quad (19)$$

that measure the deviation of each residual sub-vector from being white noise. Then for the k th iteration we propose to measure the residual’s deviation from being white noise by averaging the above quantities, i.e., by means of the “NCP-number”

$$N^{(k)} = \frac{1}{m_\theta} \sum_{\ell=1}^{m_\theta} \nu_\ell^{(k)}. \quad (20)$$

This multi-1D approach for 2D reconstruction problems leads to the following stopping rule:

Stop rule: NCP

$$\text{Minimize } N^{(k)} = \frac{1}{m_\theta} \sum_{\ell=1}^{m_\theta} \|\mathbf{c}(\mathbf{q}_\ell^{(k)}) - \mathbf{c}_w\|_2.$$

In the case of 3D reconstructions, where the data consist of a collection of 2D images, the computation of $\nu_\ell^{(k)}$ should take this into consideration. In particular we need to define the NCP vector $\mathbf{c}(\mathbf{q}_\ell^{(k)})$ when the residual sub-vector $\mathbf{q}_\ell^{(k)}$ represents an image; how to do this is explained in [16].

Similar to the previous stopping rules, in practise it is more convenient to implement the NCP stopping rule such that we terminate the iterations at the first iteration k for which $N^{(k)}$ increases. There is no theory to guarantee that $N^{(k)}$ will behave smoothly, and we occasionally see that $N^{(k)}$ exhibits a minor zig-zag behavior. Hence it may be necessary to apply the NCP stopping rule to a smoothed version of the NCP-numbers, obtained by applying a “local” low-pass filter to the $N^{(k)}$ -sequence.

Example 4. We illustrate the NCP stopping rule with a parallel-beam CT problem with image size 256×256 and with 362 detector pixels and projection angles $1^\circ, 2^\circ, \dots, 180^\circ$. The performance is shown in Fig. 5 together with surface plots of the matrix $[\mathbf{c}(\mathbf{q}_1^{(k)}), \mathbf{c}(\mathbf{q}_2^{(k)}), \dots, \mathbf{c}(\mathbf{q}_{m_\theta}^{(k)})]$ for selected iterations k . We clearly see the changing shape of the NCP vectors $\mathbf{c}(\mathbf{q}_\ell^{(k)})$ as k increases. The minimum of $N^{(k)}$ is obtained at $k_{\text{NCP}} = 179$. This is somewhat early, considering that the minimum reconstruction error is obtained at $k = 497$ iterations – but on the other hand, the reconstruction and the error changes only little between iterations 150 and 700. \square

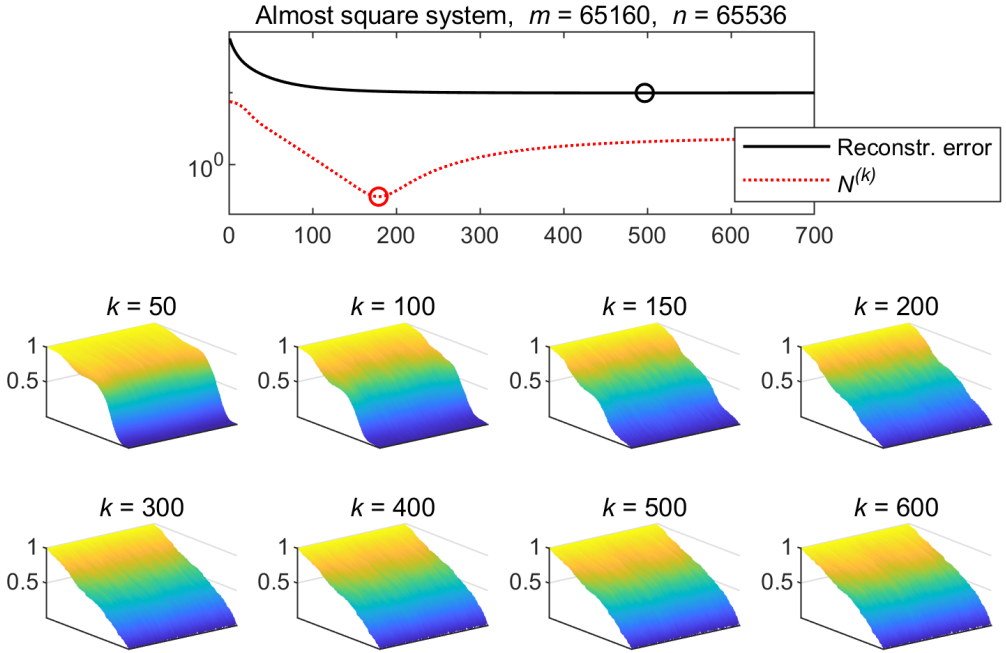


Fig. 5. Illustration of the NCP stopping rules for Landweber's method applied to a parallel-beam test problem. We also show surface plots of the matrix $[\mathbf{c}(\rho_1^{(k)}), \mathbf{c}(\rho_2^{(k)}), \dots, \mathbf{c}(\rho_{m_a}^{(k)})]$ for selected iterations k . This stopping rule leads to a somewhat premature termination of the iterations at $k_{\text{NCP}} = 179$ (the minimum error occurs for $k = 497$ iterations), but it should be noted that the error does not change much between iterations 179 and 700.

III. ESTIMATION OF THE TRACE TERM

The FTNL, UPRE and GCV stopping rules include the term $t_k = \text{trace}(\mathbf{A} \mathbf{A}_k^\#)$. To make these methods practical to use, we need to be able to estimate this trace term efficiently, without having to compute the SVD of the system matrix \mathbf{A} or form the influence matrix $\mathbf{A} \mathbf{A}_k^\#$. The most common way to compute this estimate is via a Monte Carlo approach.

Underlying this approach is the following result from [9]. If $\bar{\mathbf{w}} \in \mathbb{R}^m$ is a random vector with elements $\bar{w}_i \sim \mathcal{N}(0, 1)$, and if $\mathbf{S} \in \mathbb{R}^{m \times m}$ is a symmetric matrix, then $\bar{\mathbf{w}}^T \mathbf{S} \bar{\mathbf{w}}$ is an unbiased estimate of $\text{trace}(\mathbf{S})$. Therefore $\bar{t}_k^{\text{est}} = \bar{\mathbf{w}}^T \mathbf{A} \mathbf{A}_k^\# \bar{\mathbf{w}}$ is an unbiased estimator of $t_k = \text{trace}(\mathbf{A} \mathbf{A}_k^\#)$.

To compute this estimate we need to compute the matrix-vector product $\mathbf{A}_k^\# \bar{\mathbf{w}}$ efficiently. Recalling the definition of $\mathbf{A}_k^\#$ in Eq. (4), this can be done simply by applying the algebraic iterative method to the system $\mathbf{A} \bar{\boldsymbol{\xi}} = \bar{\mathbf{w}}$ which, after k iterations, produces the iteration vector $\bar{\boldsymbol{\xi}}^{(k)} = \mathbf{A}_k^\# \bar{\mathbf{w}}$. The resulting estimate

$$\bar{t}_k^{\text{est}} = \bar{\mathbf{w}}^T \mathbf{A} \bar{\boldsymbol{\xi}}^{(k)} = (\mathbf{A}^T \bar{\mathbf{w}})^T \bar{\boldsymbol{\xi}}^{(k)} \quad (21)$$

is the standard Monte Carlo trace estimate from [9]. In an efficient implementation of (21) the vector $\mathbf{A}^T \bar{\mathbf{w}}$ is pre-computed and stored.

An alternative approach was presented in [24]. This approach also applies to the general method in (2) with $\mathbf{D} = \mathbf{I}$

and with a general $m \times m$ matrix \mathbf{M} (it is not required to be symmetric). When we apply such a method with an arbitrary nonzero starting vector $\boldsymbol{\xi}^{(0)}$ to the system $\mathbf{A} \boldsymbol{\xi} = \mathbf{0}$, then it follows from Eq. (4) that the iterates are

$$\boldsymbol{\xi}^{(k)} = (\mathbf{I} - \omega \mathbf{A}^T \mathbf{B} \mathbf{A})^k \boldsymbol{\xi}^{(0)}.$$

Then it is shown in [24] that if we use a random starting vector $\boldsymbol{\xi}^{(0)} = \mathbf{w} \in \mathbb{R}^n$ with elements $w_i \sim \mathcal{N}(0, 1)$, and if $\boldsymbol{\xi}^{(k)}$ denotes the corresponding iterations for the system $\mathbf{A} \boldsymbol{\xi} = \mathbf{0}$, then $\mathbf{w}^T \boldsymbol{\xi}^{(k)}$ is an unbiased estimator of $n - \text{trace}(\mathbf{A} \mathbf{A}_k^\#)$. This leads to the alternative trace estimate

$$t_k^{\text{est}} = n - \mathbf{w}^T \boldsymbol{\xi}^{(k)}. \quad (22)$$

In order to use either of these trace estimates instead of the exact t_k , we must simultaneously apply the iterative method to two right-hand sides, which essentially doubles the amount of work. The Landweber method with the two different trace estimation schemes are shown below.

If we are willing to increase the overhead further, we can compute a more robust estimate of t_k by applying the above idea to several random vectors and computing the mean or median of the t_k^{est} -values.




Fig. 6. Comparison of the two trace estimates \bar{t}_k^{est} and t_k^{est} for Landweber's method applied to the over-determined test problem from Example 1. The thick red line is the exact trace t_k , and the thin black lines are the trace estimates for 10 different random vectors \bar{w} and w .

Landweber method with (21) trace estimator

\bar{w} = random m -vector for trace estimation
 $\mathbf{x}^{(0)}$ = initial vector
 $\bar{\xi}^{(0)}$ = $\mathbf{0}$ initial zero vector for trace estimation
 $\mathbf{z} = \mathbf{A}^T \bar{w}$
 for $k = 0, 1, 2, \dots$
 $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \omega \mathbf{A}^T (\mathbf{b} - \mathbf{A} \mathbf{x}^{(k)})$
 $\bar{\xi}^{(k+1)} = \bar{\xi}^{(k)} + \omega \mathbf{A}^T (\bar{w} - \mathbf{A} \bar{\xi}^{(k)})$
 $\bar{t}_{k+1}^{\text{est}} = \mathbf{z}^T \bar{\xi}^{(k+1)}$ trace estimate
 stopping rule goes here
 end

Landweber method with (22) trace estimator

w = random n -vector
 $\mathbf{x}^{(0)}$ = initial vector
 $\xi^{(0)}$ = w initial vector for for trace estimation
 for $k = 0, 1, 2, \dots$
 $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \omega \mathbf{A}^T (\mathbf{b} - \mathbf{A} \mathbf{x}^{(k)})$
 $\xi^{(k+1)} = \xi^{(k)} + \omega \mathbf{A}^T (\mathbf{0} - \mathbf{A} \xi^{(k)})$
 $t_{k+1}^{\text{est}} = n - w^T \xi^{(k+1)}$ trace estimate
 stopping rule goes here
 end

Example 5. We illustrate the two trace estimates \bar{t}_k^{est} and t_k^{est} for Landweber's method applied to the over-determined test problem from Example 1. Figure 6 shows the trace estimates for 10 different realizations of the random vectors \bar{w} and w , together with the exact trace t_k . We see that the estimate t_k^{est} , shown in the bottom plot, has the smallest variance. We




Fig. 7. Illustration of the use of the trace estimate t_k^{est} in the FTNL stopping rule for Landweber's method applied to the over-determined test problem from Examples 1. We used 10 different random vectors w in (22) and the corresponding 10 intersections between $\|e^{(k)}\|_2^2$ (thick red line) and $\eta^2(m - t_k^{\text{est}})$ (thin blue lines) are shown by the red circles. The black dot shows the intersection with the exact $\eta^2(m - t_k)$.

are not aware of theory that supports this observation. \square

Example 6. Continuing from the previous example, Fig. 7 illustrates the use of the trace estimate t_k^{est} in the FTNL stopping rule. To show the variability of the stopping rule we used 10 different random vectors w , leading to 10 different realizations of $\eta^2(m - t_k^{\text{est}})$. Their intersections with $\|e^{(k)}\|_2^2$ are shown by the red circles, corresponding to stopping the iterations at

$$k = 3100, 3112, 3421, 3512, 3722, \\ 3875, 4117, 4133, 5553, 7000.$$

The black dot marks the intersection of the exact of $\|e^{(k)}\|_2^2$ with $\eta^2(m - t_k)$, corresponding to iteration $k = 3846$. \square

IV. LARGE-SCALE NUMERICAL EXAMPLE

In this section we use a simulated large-scale CT reconstruction problem to illustrate the use of the GCV and NCP stopping rules described above. We focus on an application in dynamic tomography where the time scale of the process being examined dictates the use of a small number of projections as well as short exposure times of each projection. This leads to challenging reconstruction problems where it is critical to use a stopping rule that terminates the iterations such that $\mathbf{x}^{(k)}$ is as close as possible to \bar{x} and without having knowledge of the noise level in the data.

Specifically we study the use of the GCV and NCP stopping rules applied to the reconstruction of a single time step in a simulation of a dynamic CT experiment. The dynamic process under study is the separation of an emulsion of oil and water in a porous rock; the components separate vigorously over time, due to the two fluids being immiscible.

The basis of our simulation is a segmentation of a nano-CT scan of a piece of chalk from the Hod field in the North Sea Basin (sample id HC #15) which was scanned, reconstructed and segmented as described in [4], [20]. A subset consisting of $200 \times 256 \times 256$ voxels is chosen for the fluid simulation. Pixels outside a radius of 124 pixels from the center axis are set to zero to form a cylinder, which is mirrored along its vertical axis to ensure that the multiphase flow simulation has periodic boundary conditions. The flow simulation is done




Fig. 8. A single slice of the volume being examined. White corresponds to rock while light and dark grey correspond to water and oil respectively.

with a phase-field Lattice Boltzmann method for systems that are isothermal and incompressible [7], [8]. The simulation produces phase values for each voxel that describes the fraction of oil and water in the voxel. These phase values are converted to attenuation values based on values measured in [19] where a sandstone filled with a brine and oil is imaged with X-rays at 80 keV.

Figure 8 shows a slice from a single time frame in the simulation; the rock matrix is white, the water phase is light grey, and the oil phase is dark grey. The time frame which was chosen for testing the stopping rules is fairly early in the simulation where multiple interfaces between the two fluids are present.

Forward projection: The forward projection of the volume is performed using the ASTRA toolbox with a parallel beam geometry [1], [2], [22]. We use 362 detector pixels and 360 projection angles. The forward projection in ASTRA can be considered an ideal experiment with monochromatic X-rays and infinite brilliance, i.e., without any noise.

Noise: We create noisy data from the above clean data in such a way that we emulate the noise present in X-ray tomography as a result of the finite count of photons, cf. [15, §4.4]. Specifically, if $\bar{\mathbf{b}} = \mathbf{A}\bar{\mathbf{x}}$ denotes the clean data computed by means of ASTRA, then the corresponding X-ray intensities at the detector are given by

$$\bar{I}_i = I_0 \exp(-\bar{b}_i), \quad i = 1, \dots, m,$$

where I_0 is the source's intensity. We then use \bar{I}_i as the expected value in a Poisson distribution to obtain noisy intensities

$$I_i = \mathcal{P}(\bar{I}_i) = \mathcal{P}(I_0 \exp(-\bar{b}_i)), \quad i = 1, \dots, m.$$

Finally, we convert these noisy intensities back to the noisy data vector \mathbf{b} via the relation

$$b_i = -\log(I_i/I_0), \quad i = 1, \dots, m.$$

We use three different noise levels 0.25%, 1% and 5% which visually corresponds to low, moderate and high noise. The noise level is given by

$$\rho = \frac{\|\mathbf{e}\|_2}{\|\bar{\mathbf{b}}\|_2}, \quad \mathbf{e} = \mathbf{b} - \bar{\mathbf{b}}, \quad (23)$$

where \mathbf{e} denotes the measurement error in Eq. (1). This noise does not exactly fit with the assumption of white Gaussian noise which is used for the previous derivations, but it is a good approximation to the noise present in CT experiments.

Reconstruction: We compute reconstructions from the simulated projection data with the ASTRA toolbox by using the Simultaneous Iterative Reconstruction Technique (SIRT) iterative method. This is a special case of the general method in Eq. (2) where the diagonal matrices \mathbf{D} and \mathbf{M} contain the inverse column and row sums of \mathbf{A} :

$$d_{jj} = 1 / \sum_{i=1}^m a_{ij} \quad \text{and} \quad m_{ii} = 1 / \sum_{j=1}^n a_{ij}.$$

We perform 1000 SIRT iterations and for UPRE and GCV the trace $t^{(k)}$ in $U^{(k)}$ and $G^{(k)}$ is estimated using Eq. (21). Note that only a single random vector \mathbf{w} is used to reduce computation time. This is not a concern in this specific case as the estimation of $t^{(k)}$ proved very stable for different random vectors.

The forward projection of $\mathbf{x}^{(k)}$ used in the calculation of $N^{(k)}$ is computed with ASTRA, and the remaining part of the algorithm is calculated with CuPy, a Python package which makes it possible to offload calculations to a CUDA-compatible GPU to improve the computation time of $N^{(k)}$. The vector \mathbf{v} in (16) is padded with zeros such that its length can be written in the form $n = 2^a + 3^b + 5^c + 7^d$ as this substantially speeds up the calculation when the DFT is calculated with CUDA.

Results: As previously mentioned, we reconstruct the data at three different noise levels. Moreover, we subsample the number of projections used for the reconstruction such that it is performed with 360, 120 and 45 projection angles. This leads to 9 different data sets.

Figure 9 shows graphs of $\|\mathbf{g}^{(k)}\|_2$, $\tau \eta \sqrt{m - t_k}$ (called FTNL), $U^{(k)}$, $G^{(k)}$, $N^{(k)}$ and $\|\mathbf{x}^{(k)} - \bar{\mathbf{x}}\|_2$ along with their local minima. In general we see that the FTNL and UPRE stopping rules perform well in this simulated example; but they depend on knowledge of the noise level η which is rarely known for real data. An advantage of the GCV and NCP stopping rules is that they do not rely on an estimate of η . GCV performs well in the case with 1% and 5% noise, but it overestimates the number of iterations with the small noise level $\rho = 0.25\%$. The NCP stopping rule is also a robust method and it performs well for all noise levels. It is worth noting that the reconstruction error is very flat for this noise level, which means the exact amount of iterations used is less critical.

Figure 10 shows the effect of semi-convergence on the data set when it is reconstructed with $N_{\text{proj}} = 360$ and $\rho = 1.00\%$. A single slice of the reconstruction is shown in all subfigures.

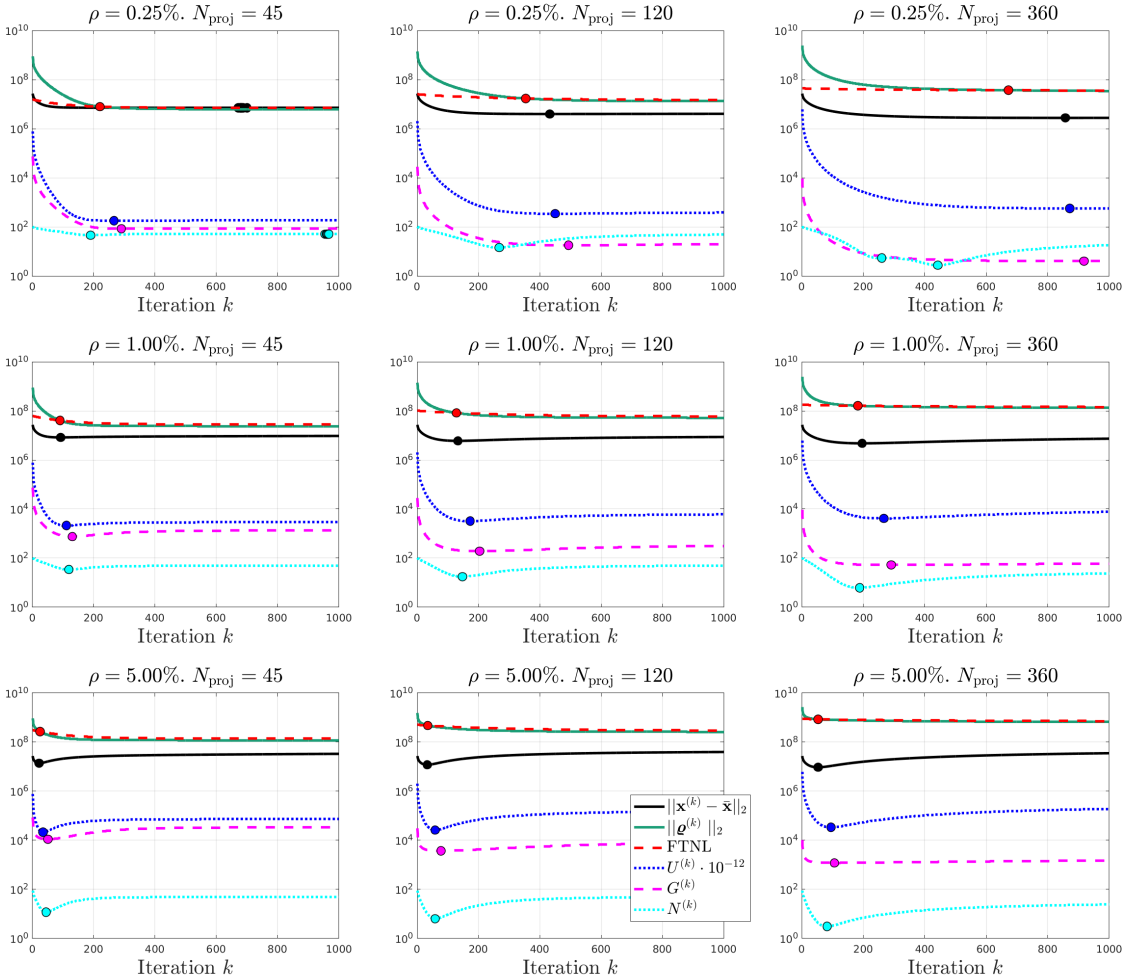


Fig. 9. Illustration of the four different stopping rules for the large-scale example. The filled circles on each curve represent the minimum. Each column has a varying number of projections and each row has a varying amount of noise, as shown in the titles of each subplot.

All images in Fig. 10 are truncated such that their intensities are between 0 and 25,000. Image **a** is the reconstruction after $k = 100$ iteration where it still has a blurred appearance, showing that more iterations are necessary. Images **b** and **c** are the reconstructions at the number of iterations which minimize the reconstruction error and $N^{(k)}$, respectively. The appearance of these reconstructions is very similar, but **c** has a slight increase in noise. The rightmost image **d** is the reconstruction after 1000 iterations where it is noticeable more noisy than **b** and **c**.

V. CONCLUSION

We surveyed several state-of-the-art stopping rules, based on statistical considerations, that are useful for algebraic iterative reconstruction methods in X-ray computed tomography (CT). Common for these stopping rules is that they seek to terminate the iterations at the optimal point where the reconstruction error and the noise error balance each other. They are easy to use and they are also easy to integrate in existing software. We also illustrated the use of two of these methods for a large-scale CT problem related to the study of multiphase flow in chalk. Our numerical experiments show that especially the NCP stopping rule – which is based on statistical properties of




Fig. 10. Illustration of semi-convergence for the large-scale example with 360 projections and noise level $\rho = 1\%$. Results for four different iteration numbers are shown. Image **a** is the reconstruction after 100 iterations. Image **b** is the reconstruction at $k_{\min} = 157$ which is the number of iterations that minimizes the reconstruction error. Image **c** is the reconstruction at $k_{\text{NCP}} = 189$ which is the number of iterations that minimizes $N^{(k)}$. Finally, image **d** is the reconstruction after $k = 1000$ iterations.

the residual and does not depend on knowledge of the noise level – works well for this problem.

ACKNOWLEDGEMENT

We thank the referees for valuable comments.

REFERENCES

- [1] W. van Aarle, W. J. Palenstijn, J. De Beenhouwer, T. Altantzis, S. Bals, K. J. Batenburg, and J. Sijbers, *The ASTRA Toolbox: a platform for advanced algorithm development in electron tomography*, *Ultramicroscopy*, 157 (2015), 35–47.
- [2] W. van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabrovolski, J. De Beenhouwer, K. J. Batenburg, and J. Sijbers, *Fast and flexible X-ray tomography using the ASTRA Toolbox*, *Optics Express* (2016), 24(22), 25129–25147.
- [3] A. H. Andersen and A. C. Kak, *Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm*, *Ultrasonic Imaging*, 6 (1984), pp. 81–94.
- [4] S. Bruns, S. L. S. Stipp, and H. O. Sørensen, *Statistical representative elementary volumes of porous media determined using greyscale analysis of 3D tomograms*, *Adv. Water Resour.*, 107 (2017), pp. 32–42.
- [5] T. Elfving, T. Nikazad, and C. Popa, *A class of iterative methods: semi-convergence, stopping rules, inconsistency, and constraining*; in Y. Censor, M. Jiang, and G. Wang (Eds.), *Biomedical Mathematics: Promising Directions in Imaging, Therapy Planning, and Inverse Problems*, Medical Physics Publishing, Madison, WI, 2010, pp. 157–184.
- [6] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Springer, 1996.
- [7] A. Fakhari, T. Mitchell, C. Leonardi, and D. Bolster, *Improved locality of the phase-field lattice-Boltzmann model for immiscible fluids at high density ratios*, *Physical review. E*, 96(5-1) (2017), 053301.
- [8] A. Fakhari, D. Bolster and L. Luo, *A weighted multiple-relaxation-time lattice Boltzmann method for multiphase flows and its application to partial coalescence cascades*, *J. Computational Physics*, 341 (2017), pp. 22-43.
- [9] A. Girard, *A fast ‘Monte-Carlo cross-validation procedure’ for large least squares problems with noisy data*, *Numer. Math.*, 56 (1989), pp. 1–23.
- [10] G. H. Golub, M. T. Heath, and G. Wahba, *Generalized cross-validation as a method for choosing a good ridge parameter*, *Technometrics*, 21 (1979), pp. 215–223.
- [11] R. Gordon, R. Bender, and G. T. Herman, *Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography*, *J. Theoretical Biology*, 29 (1970), pp. 471–481.
- [12] P. Gravel, G. Beaudoin and J. A. De Guise, *A method for modeling noise in medical images*, *IEEE Trans. Medical Imaging*, vol. 23 (2004), pp. 1221–1232.
- [13] P. Hall and D. M. Titterton, *Common structure of techniques for choosing smoothing parameters in regression problems*, *J. Royal Stat. Soc.: Series B (Methodological)*, 49 (1987), pp. 184–198.
- [14] P. C. Hansen, J. S. Jørgensen, and M. Saxild-Hansen, *AIR Tools II: algebraic iterative reconstruction methods, improved implementation*, *Numerical Algorithms*, 79 (2018), pp. 107–137.
- [15] P. C. Hansen, J. S. Jørgensen, and W. R. B. Lionheart (Eds.), *Computed Tomography: Algorithms, Insight and Just Enough Theory*, SIAM, Philadelphia, 2021.
- [16] P. C. Hansen, M. E. Kilmer, and R. H. Kjeldsen, *Exploiting residual information in the parameter choice for discrete ill-posed problems*, *BIT Numerical Mathematics*, 46 (2006), pp. 41–59.
- [17] K. Hahn, H. Schöndube, K. Stierstorfer, J. Hornegger, and F. Noo, *A comparison of linear interpolation models for iterative CT reconstruction*, *Medical Physics*, 43 (2016), pp. 6455–6473.
- [18] J. Kay, D. Titterton, J. Brown and A. Thompson, *A Study of Methods of Choosing the Smoothing Parameter in Image Restoration by Regularization*, *IEEE Trans. Pattern Anal. Machine Intelligence*, 13 (1991), pp. 326–339.
- [19] Q. Lin, M. Andrew, W. Thompson, M. J. Blunt and B. Bijeljic, *Optimization of image quality and acquisition time for lab-based X-ray microtomography using an iterative reconstruction algorithm*, *Advances in Water Resources*, 115 (2018), pp. 112–124.
- [20] D. Mütter, H. O. Sørensen, D. Jha, R. Harti, K. N. Dalby, H. Suhonen, R. Feidenhans’l, F. Engström, and S. L. S. Stipp, *Resolution dependence of petrophysical parameters derived from X-ray tomography of chalk*, *Appl. Phys. Lett.*, 105 (2014), 032108.
- [21] F. Natterer, *The Mathematics of Computerized Tomography*, SIAM, Philadelphia, 2001 (reprint).
- [22] W. J. Palenstijn, K. J. Batenburg, and J. Sijbers, *Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs)*, *J. Structural Biology*, 176 (2011), pp. 250–253.
- [23] B. W. Rust and D. P. O’Leary, *Residual periodograms for choosing regularization parameters for ill-posed problems*, *Inverse Problems*, 24 (2008), 034005.
- [24] R. J. Santos and Á. R. De Pierro, *A cheaper way to compute generalized cross-validation as a stopping rule for linear stationary iterative methods*, *J. Comput. Graphical Statistics*, 12 (2003), pp. 417–433.
- [25] V. F. Turchin, *Solution of the Fredholm equation of the first kind in a statistical ensemble of smooth functions*, *USSR Comp. Math. and Math. Physics*, 7 (1967), pp. 79–96.
- [26] C. Vogel, *Computational Methods for Inverse Problems*, SIAM, Philadelphia, 2002.
- [27] G. Wahba, *Spline Models for Observational Data*, SIAM, Philadelphia, 1990.

CHAPTER 7

Contribution C:

Crystal formation from high resolution 4D X-ray μ CT: Initiation, expansion, and densification

The two previous contributions represent the theoretical work done during the project. However, we always had the intention that the algorithms developed in **Contribution A** should be used experimentally as this would verify their performance. Unforeseen events delayed the production of the tri-axial flow cell, so we decided to broaden the scope of the project such that the focus shifted from flow in porous media to dynamic processes in general.

One of my co-supervisors, Henning Osholm Sørensen, therefore, introduced me to PhD student Isaac Appelquist Løge, whose project is to develop a method to investigate the kinetics of scale formation. We decided to perform dynamic in situ experiments that examine the crystallisation process as there have been few of these in the literature. We performed the initial test experiments in the spring of 2021 and gathered the final data in the fall of 2021.

During the experiment, we flooded a carbon steel flow cell with water that was supersaturated with BaSO_4 for over 150 hours while continuously imaging it in a CT scanner. We used a modified version of one of the algorithms presented in **Contribution A** to reconstruct the projections. This reduced image noise considerable and allowed us to analyse the scaling process in detail.

The contribution is still in preparation, and a draft of it is, therefore, included below.

Crystal formation from high resolution 4D X-ray μ CT: Initiation, expansion, and densification

Isaac Appelquist Løge^{a,1,*}, Peter Winkel Rasmussen^{b,1,**}, Henning Osholm Sørensen^c, Anders Nymark Christensen^b, Anders Bjorholm Dahl^b, Philip Loldrup Fosbøl^a

^a*Technical University of Denmark, Department of Chemical Engineering., Søtofts Plads, Building 228A, Kgs. Lyngby, DK-2800, Denmark*

^b*Technical University of Denmark, Department of Applied Mathematics and Computer Science., Richard Petersens Plads, Building 324, Kgs. Lyngby, DK-2800, Denmark*

^c*Technical University of Denmark, Department of Physics., Fysikvej, Building 307, Kgs. Lyngby, DK-2800, Denmark*

Abstract

Crystallisation fouling presents a costly challenge, as it lowers the efficiency of heat exchangers, decreases the throughput of desalination membranes and clogs underground pipelines. Optimal maintenance requires accurate knowledge of the kinetics of the formation of crystallisation fouling. Current methods for comprehensive in situ investigations are limited. The most commonly used ex situ or 2D experiments cannot capture dynamic effects or describe formation kinetics. This work uses X-ray computed tomography (CT) scanning to address this challenge. We recorded the formation of crystalline fouling in situ by installing a custom-built flow loop inside a CT cabinet. To record the temporal changes in the crystalline fouling, we had to sparsely sample and underexpose the specimen, so we employed a custom reconstruction algorithm ensured accurate reconstructions. The high spatiotemporal resolution of the experiment made it possible to study the processes of crystal formation. Our findings show how the growth rate can vary 300 % in the initial phases of growth. Even though we performed measurements for 150 hours, the rate at which growth formed was not constant. This non-equilibrium effect reflects the complexity of crystallisation formation in natural systems. Our findings provides a foundation for unveiling the dynamics of fouling mechanisms, which will aid in developing more accurate prediction models.

Keywords: Crystallisation fouling, scale formation, 4D x-ray CT, In situ tomography, tomographic reconstruction, graph cut segmentation

*Corresponding author

**Second corresponding author

Email addresses: isacl@kt.dtu.dk (Isaac Appelquist Løge), pwra@dtu.dk (Peter Winkel Rasmussen), osho@fysik.dtu.dk (Henning Osholm Sørensen), anym@dtu.dk (Anders Nymark Christensen), abda@dtu.dk (Anders Bjorholm Dahl), plf@kt.dtu.dk (Philip Loldrup Fosbøl)

¹These authors contributed equally.

1. Introduction

Crystallisation fouling, i.e. the growth of unwanted crystals on a solid surface, presents both technical and economic challenges in many industries [1, 2, 3, 4, 5, 6]. Optimal mitigation strategies hinges on accurate prediction models that are established through developing insights to the process at hand. Decoupling the build-up allows for functional mitigation strategies, capable of targeting the particular problem. Fouling can be divided into five separate stages occurring throughout the build-up; nucleation, transport, attachment, detachment and ageing [7]. Numerous experiments on the separate stages of fouling have revealed some of the underlying mechanics.

Surface nucleation of insoluble crystals depends on the number of kink sites as kink sites lower the surface energy [8]. Transport governs both bulk mixing and influx of reactive species to the surface; fluid turbulence can cause complex anisotropic mixing [9], and accelerate wall transport in systems where surface features modulate fluid flow [10]. Attachment is important both from a molecular and macroscopic viewpoint. Molecular attachment in stagnant hydrodynamic surface layers can cause unexpected induction times [11]. Macroscopic adhesion of pre-crystallised particles is dependent on gravitational forces, and the rate of deposition or particle sedimentation becomes dependent on crystal density [12]. Crystal detachment is dependent on the forces from the liquid [10], and of the substrate on which they are formed [13]. Modulation of substrate properties can both increase [14] and decrease growth rates [13]. The individual phases described above can be characterised in-depth; however, piecing the processes together is difficult due to experimental limitations. Nucleation requires a high spatiotemporal resolution, as these processes are rapid, while transport, attachment, detachment and ageing are occurring over long periods.

Investigating fouling formation has, throughout the literature, been performed with three prototypical setups; the bulk jar test experiment, the rotating electrode experiment, and the flow loop [15] where fouling was estimated by gained mass or increased pressure. To bridge the gap, there has been a call for new approaches [16, 17]. Examples of newly developed methods used to examine fouling formation are:

- Scanning electron microscopy paired with energy dispersive X-ray spectroscopy [16].
- Quartz crystal micro balance (QCM) in combination with atomic force microscopy [18].
- In situ synchrotron CT scanning [19].
- X-ray fluorescence in conjunction with μ CT scanning [20].
- In situ X-ray diffraction [21, 22].

Insights into the surface attachment mechanism have been gained from QCM, which showed larger crystal growth at higher saturation index [18]. The growth kinetics of BaSO_4 was investigated through in situ synchrotron CT scanning, permitting a more in-depth mechanistic understanding of the growth in

porous media [19]. However, development of an experimental method capable of balancing a sufficiently high spatiotemporal resolution, with a long experimental time is missing [23].

The once-through flow loop poses as a good template for simulating real-world conditions of residence time, hydrodynamics, and physio-chemical conditions, relevant for crystal growth. Furthermore, there has been recent development showcasing the possibilities of designing these once-through systems, such that surface adhesion was decoupled from bulk precipitation [24, 25].

In this study, we combined a once-through flow loop with spatiotemporally resolved CT scanning to measure the fouling formation process throughout multiple stages. In situ investigation of fouling formation allowed for a detailed analysis of different stages in growth and the transition between modes of fouling. We decoupled three growth mechanisms present in crystallisation fouling build-up and showed how they affect volume and mass deposition independently. We observed an initiation phase, characterised by many crystal clusters forming; an expansion phase, where the volumetric growth rate is larger than the gravimetric rate; lastly, we show densification, where incoming ions are closer packed, and the gravimetric rate is larger than the volumetric.

Optimal mitigation strategies are based mechanistic insights, as these allow for protocols with the correct functionalities to be developed. Predicting crystallisation formation is based on accurate knowledge of the growth rate in the system. We show a growth rate that varies non-linearly with a span of 300 % from 0 to 150 hours. Non-equilibrium effect reflects the complexity of crystallisation formation in natural systems. Our findings allow for comprehensive prediction models to be developed, decoupling the various stages present in crystallisation build-up, and implement functionality based intervention strategies.

2. Materials and Method

2.1. Materials

BaCl₂ (99.9 %, Sigma Aldrich), and Na₂SO₄ (99.0 %, Merck Millipore), were used as received. Two separate solutions were prepared: 0.2 mmolal BaCl₂ and 0.27 mmolal Na₂SO₄. The saturation indices (SI) for the solution were calculated at 60 °C after mixing in a 1:1 ratio using the PHREEQC geochemical simulator. Pitzer parameters were implemented [26, 27] and SI is calculated as:

$$SI = \log_{10} \left(\frac{a_i a_j}{K_{SP}} \right) \quad (1)$$

where a_i and a_j are the activity coefficients of the reactants and K_{SP} is the solubility constant. The SI was calculated to be 1.5 for the BaSO₄ solution.

2.2. The experimental procedure

The experimental set-up is identical to the one described and used by Løge et al. for ex situ fouling investigation [10]. We used single flow-through mode in a deposition cell with an internal diameter of 3 mm

and a length of 3 cm. The cells were fabricated on a CNC machine in a workshop affiliated with the authors, based on the deposition cell design from Løge et al. [10]. The reaction temperature was fixed to 60 °C, and the injected flow rate for each brine was 36 mL/min, which gave a Reynolds number of 1080.

2.3. In situ X-ray tomography experiment


We imaged the set-up using the ZEISS XRadia Versa 410 instrument at the 3D Imaging Centre, DTU. Before and after the in situ scanning procedure, high-quality scans of the sample were acquired. These scans consisted of 2401 projections acquired over a 360° rotation. Acquisition was done using a beam energy of 150 kV and an exposure time of 5 s. Beam hardening was accounted for by filtering low energy radiation away using an HE2 filter. Optical magnification of 4× was used yielding a pixel size of 4.6 μm. We used binning of factor two which resulted in a volume with 1015³ voxels. The number of projections per scan was reduced to 401 during the dynamic imaging to decrease acquisition time. The reduction resulted in an acquisition time of approximately one hour, which ensured that we had a sufficient time resolution to capture the dynamics of the scaling process. A total of 154 scans were gathered.

2.4. Reconstruction method

Artefacts and high noise levels are generally obtained using the standard FDK reconstruction algorithm when few projections are recorded [28, 29, 30]. To avoid these we used the SIRT-BC algorithm as described by Rasmussen et. al [31]. The SIRT-BC algorithm is based on the Simultaneous Iterative Reconstruction Technique [32, 33] and is here used adding box constraints, i.e. an upper and a lower limit for the reconstructed values is given in the reconstruction procedure.

The reconstruction algorithm leverages a high-quality scan of the sample acquired before dynamic data-taking. The high-quality scan is without noise or other artefacts as it consists of more projections. The first time step of the dynamic reconstruction is initialised with this high-quality reconstruction and subsequent time steps are initialised with the reconstruction of the previous time step. The initialisation improves both the quality of the reconstruction and speeds up the convergence of the algorithm. The difference between FDK and SIRT-BC can be seen in Fig. 1, which shows an orthoslice of the cell after 50 hours of flow. The magnified area shows a noticeably larger noise level in the FDK reconstruction. The noise makes distinguishing the scale from the background difficult. Relying on this reconstruction would inhibit further analysis.

Algorithms such as SIRT-BC does not necessarily converge when the reconstruction that resembles the actual object the most is obtained. Iterating beyond the optimal point decreases the signal to noise ratio of the reconstruction, which makes it desirable to estimate this point. We attempted to do this with the NCP stopping rule [34], but the NCP stopping rule terminated the reconstruction prematurely. Therefore, we elected to use 100 iterations, based on experiments, for each time step.



(a) FDK

(b) SIRT-BC

Figure 1: A comparison of a reconstruction performed with the FDK algorithm and the SIRT-BC algorithm after 50 hours of flow.


Six data sets were lost during acquisition, and an additional two data sets exhibited a large amount of sample movement and were not reconstructed. The first five reconstructions were excluded from the analysis as the sample shifted slightly during this period. The remaining 141 data sets, excluding the initial and final high-quality scans, were used for analysis throughout the study.

2.5. Segmentation

Quantitative information about the crystallisation fouling process were extracted by segmenting both steel and BaSO_4 in the reconstructions. Threshold segmentation was not feasible as the formed crystals were faint at early stages of the experiment and at later stages the pixel intensity of the steel and BaSO_4 is approximately identical. To create a satisfying segmentation, we used a graph cut method described by Jeppesen et. al [35, 36].

The graph cut method requires flat surfaces, which is not the case for the circular cell. The surface was unrolled by a geometric transformation from Cartesian coordinates into polar coordinates around its centre. Such a transformation will make a circular object appear flat. The graph cut algorithm minimises the energy of the graph cut it performs. The energy minimisation occurs when an image is cut into separate regions along the most distinct surface present. The algorithm cuts the unrolled data in the darkest line, thus yielding a division that provides the segmentation. To improve the performance of the algorithm the contrast between the steel and the surrounding area should be maximised. We obtain a high contrast by using the derivative and negative derivative of the image as it provides a single sharp transition for the

outside and inside of the cell. An example of the graph cut method can be seen in Fig. 2. In Fig. 2a we see a single unrolled slice of the cell for time step 100 along with the segmentation shown in cyan. From the negative and positive derivatives the energies are created (Fig. 2b and Fig. 2c). The edge of the steel is very dark compared to the background in these images, which is why the graph cut is performed there. The segmentation gives us labels for the air surrounding the cell, the cell itself and the insides of the cell. The segmentation algorithm is fast but memory intensive. Due to memory limitations only fifteen time steps can be segmented simultaneously.



(a) An example of the unrolled data. (b) The energy used to find the inner edge of the steel. (c) The energy used to find the outer edge of the steel.

Figure 2: An example of the segmentation for the central slice of the cell after 100 hours of injection. The cyan lines indicate the surface detected by the graph cut algorithm. The area above the upper line is the inside of the cell where the scale is growing. Below this line is the cell wall, and below the lower line, is the air surrounding the cell.

We would also like to segment the surface of the BaSO_4 . This surface is much fainter, especially during the first 50 hours of the experiment, which makes it far more difficult to segment.

We use the previously obtained segmentation to remove the air and steel surrounding the inside of the cell before using the graph cut method. For decent segmentation of the faint scale formation it is not sufficient to use only the the derivative. To enhance the contrast between the scale and the background the unrolled data must be preprocessed. Preprocessing is initiated by using an iterative nonlocal means denoiser to increase the signal to noise ratio of the image [37]. Secondly, the denoised data is corrected for beam hardening. The data values are then centre around zero by subtraction the average of the nonzero elements of the unrolled data. The final step of the preprocessing, is filtering the data along the radial axis with a median filter and with a Gaussian filter along both the radial and the angular axis. The final energy is given by maximal image value minus the derivative and double derivative of the preprocessed image.

2.6. Surface analysis

We calculated the texture parameters from Gadelmawla et al. [38] with the methodology developed by Klingaa et al. [39, 40]. The algorithm used the segmented orthoslices of the reconstruction as input. The profile texture parameters used in this work are the profile skewness, Psk, calculated as:

$$\text{Psk} = \frac{1}{\text{Pq}^3} \left(\frac{1}{l} \int_0^l Z(x)^3 dx \right) \quad (2)$$

where Pq is the variation of the surface height $Z(x)$, l is the length of measured surface in the axial direction, and x is position on the surface. The skewness describes whether a height distribution are distributed above (predominately consists of peaks) or below zero (predominately consists of valleys). The kurtosis, Pku, is calculated as:

$$\text{Pku} = \frac{1}{\text{Pq}^4} \left(\frac{1}{l} \int_0^l Z(x)^4 dx \right) \quad (3)$$

and describes the spikiness of a surface. Values above three are generally considered "spiky" surfaces.

Finally, the average height of the 10 largest distances from valley to peak, P10z, is calculated as:

$$\text{P10z} = \frac{1}{10} \left| \sum_{i=1}^{10} \max_{10}[Z(x)]_i - \min_{10}[Z(x)]_i \right| \quad (4)$$

where \max_{10} and \min_{10} are operators which return the ten largest and the ten smallest values of a vector. P10z gives an indication of the overall size of the surfaces peaks and valleys.

2.7. Surface coverage and cluster distribution

We used an in-house program written in MATLAB to calculate the surface coverage. The program used orthoslices of the reconstructions as input. Firstly, the algorithm identifies the inner diameter of the cell. Secondly, the line integral along the inner diameter of the cell toward its centre is calculated. The line integral should be sufficiently long such that all crystals that grow on the cell are included in the line integral. The perimeter was divided into 500 points with an integration length of 0.3 radii. We used a geometric transformation similar to the one described in Sec. 2.5 on the line integrals resulting in a planar surface. The surface was subjected to a threshold based on the mean value of the line integrals, which divided the surface into areas defined as covered and uncovered. The ratio of the covered and the total surface areas defines the surface coverage.

Based on the segmented surfaces, we use the MATLAB function `bwconncomp` with its default settings to calculate the number of clusters and the size of them. A covered pixel is defined as a cluster if its four nearest neighbours (or more) are covered as well. The number of cluster at different sizes allows us to calculate the distribution of cluster sizes.

3. Results and Discussion

Crystallisation fouling formation is a complex process, going through various stages throughout its build-up. We have identified three growth modes through spatiotemporal measurement of the formation process; initiation, expansion, and densification. The surface area of the formed crystals indicates how three phases are present (Fig. 4b). The initiation phase was characterised local roughening arising from corrosion, which started the crystal formation, and the overall surface area decreased. The volumetric growth rate was larger than the gravimetric rate in the expansion phase, while the opposite was true in the densification phase. As spiky features are formed in the expansion phase (Fig. 4a), the surface area increases, whilst in the densification phase, the crystals grow in density, while the surface area remains constant. The following section characterises and quantifies the process in the identified regimes.

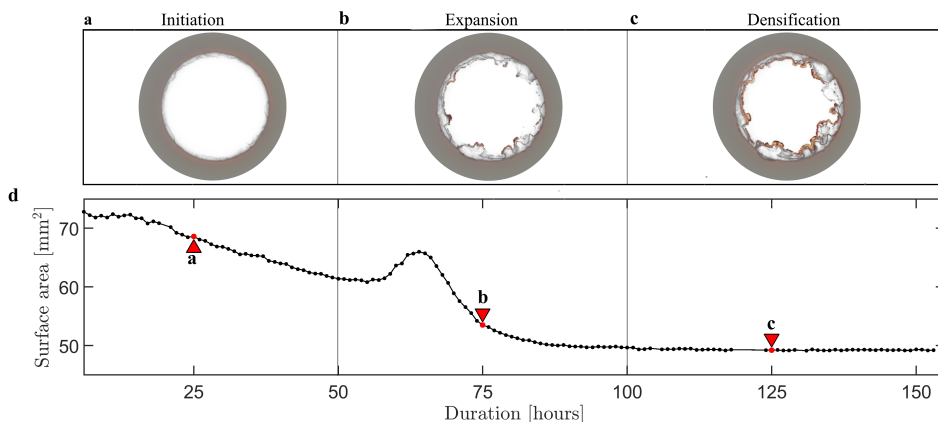


Figure 3: Observed mechanisms for crystallisation fouling. **a-c** Cross-sections of the cell at **a** 25 hours, **b** 75 hours and **c** 125 hours. **d** The surface area of the BaSO_4 as a function of time.

3.1. Initiation: Crystal formation on pristine steel

Initiation of growth on a pristine surface through the nucleation and formation of clusters is a transient stage of the overall formation process. However, without overcoming the initial barrier for growth, a surface will stay unfouled (Fig. 4a). The required energy to overcome this initial barrier for growth depends on the condition of the surface[41]. A surface with either preexisting scale or roughness would have a smaller barrier for surface [42, 43, 14]. Corrosion processes could create roughness capable of initiating surface growth, which is something we observed during the initial phase of the scale formation. Corrosion processes are initiated by the presence of dissolved O_2 in the solution[44]. As corrosion process are initiated, crystallisation fouling will attach to the local rough points (Fig. 4b). The surface is in pristine condition at the start of the experiment.

After 5-10 hours corrosion roughens the surface, seen in the increased Pku. Surface morphologies with a Pku above three are deemed spiky. The roughness act as nucleation sites for crystallisation fouling, which initiates the growth. After 15 hours, growth occurs in multiple corroded areas, and it spreads out in the plane of the surface. The formation of surface roughness due to corrosion is coincident with the formation

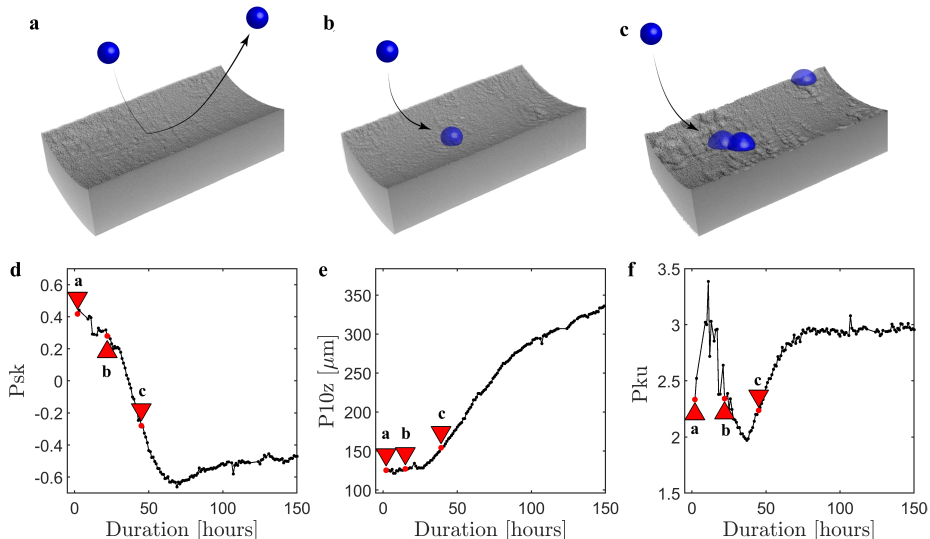


Figure 4: Initiation of crystallisation fouling. **a-c** Visualisation of initiation mechanism. The steel surfaces corresponds to measurements at **a** 6 hours, **b** 15 hours, **c** 45 hours. The corrosion behaviour of the steel surface is characterised from **d** the skewness of the height distribution (Psk), **e** P10z and **f** the spikiness of the height distribution (Pku)

of the first clusters, which lead us to believe that corrosion and scale formation are intimately linked. The formation of the most clusters can be observed to occur during 20-50 hours, and it is in the period that the local corroded areas are formed.

The formation of clusters is also witnessed in the cluster distribution, which is shown in Fig. 5b. In the first 20 hours, nucleation seeds grow on the whole surface before reaching a detectable size with our methods. This is followed by a rapid increase in the number of clusters which is a sign of widespread surface nucleation. None of these initial formations are taller than $30 \mu\text{m}$ after 50 hours of growth (Fig. 6e). The low height is because the clusters are, at this point, primarily growing laterally. The height of the clusters increases as they start to coalesce during the expansion phase of the fouling.

As the clusters expand, the surface coverage increases to 20% at 50 hours (Fig. 6e). After the initiation period at 0-20 hours, the rate of surface covering is linear at $0.5\% \text{ hour}^{-1}$. The constant surface coverage




Figure 5: Evolution of clusters. **a** Bivariate histogram of cluster evolution. **b** Amount of clusters present.

rate also indicates that formed clusters are not coalescing. The clusters form near each other and create the outline of a superstructure. These larger structures of higher intensity are faint but can be observed in Fig. 6a-c. These show how a porous layer is being formed.

The deposition of crystals formed from 0-50 hours is low. Even though the crystals cover up to 20% of the surface, and more than 70 clusters have formed, this causes little change in the intensity maps.

Curiously, after overcoming the barrier for surface growth, all presented parameters grow linearly. A constant rate suggests that from 20 to 50 hours, the growth mechanism is unchanged in this period. Therefore, until approximately 20% of the surface is covered with scale growth, the formation rate is caused by crystals growing directly on steel. The clusters that form in this period are growing in the proximity of each other, and are forming superstructures, however not much crystalline material is formed yet.

3.2. Expansion: Scale formation on a partially covered steel surface

After 50 hours of crystal formation, a new growth behaviour starts. We observe that instead of formation of new clusters, the existing deposited crystals expand rapidly in volume. The expansion, is observed both in the surface coverage, height of the crystals, and by the fact that crystals start to coalesce (see Fig. 6e, Fig. 6f and Fig. 5b). The growth mechanism changes as the surface growth has covered 20% of the surface. A rapid expansion of the scale growth can be observed in both the volumetric deposition and spikiness shown in Fig. 7a and Fig. 7b. At 60 hours, the growth behaviour changes; there is a stark increase in the deposited volume. The surface area, see Fig. 3, further supports this with a decrease from 20 to 60 hours and then a sudden increase at 60-70 hours. A decrease in the surface area indicates that uniform deposition occurs, while an increase in surface area results from features that extend away from the surface. Expansion occurs on the whole surface as the surface area decreases again. Therefore, we observe that in the period from 50-100 hours, the crystals undergo a growth process associated with a significant expansion. Further

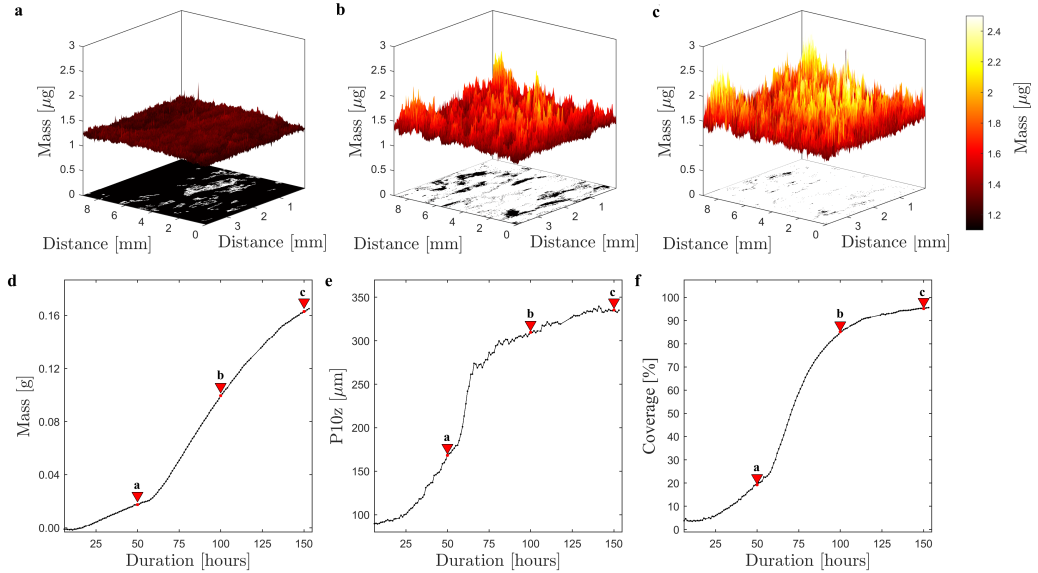


Figure 6: Evolution of scale formation. **a-c** Surface map of the fouling formation for **a** 50 hours, **b** 100 hours, **c** 150 hours. **d** Intensity of deposited crystals. **e** P10z for BaSO_4 segmented surface. **f** Surface coverage.

evidence for this phase to be a transition, is in the spikiness (Pku) of the scale. The value of Pku is around 3.2 from 0-50 hours; however, at 50 hours Pku increases to 3.6 and decreases to 2.8 at 70 hours where it settles. This shift in the surface morphology indicates that two different growth behaviours were present before 50 hours, and after 70 hours, with a transition in between.

At 50 hours, clusters start to coalesce, as they expand in the lateral dimension of the surface and the available surface area of pristine steel decreases. The number of clusters between 50 and 100 hours is shown in Fig. 5b, which shows the competition between formation and coalescence. After 70 hours, the net cluster formation rate is negative, which can be ascribed to two effects. Firstly, less area for new clusters to grow exists, thereby decreasing the formation rate. Secondly, as the clusters grow larger, the distance between them decreases, increasing the rate of coalescence. The rate of coalescence can be observed in Fig. 5a. The number of small clusters have decreased between 50 hours and 100 hours and have instead been replaced by larger clusters.

Larger crystals, in general, arise both from a larger surface attachment area and the increased crystal height. During 60-70 hours, the crystal height increases from 50 μm to 70 μm . The height increase corresponds to the increase in surface area in Fig. 3b. This rapid height increase is not a global phenomenon but happening in local points; otherwise, the surface area would decrease. From 70 to 100 hours, the crystal

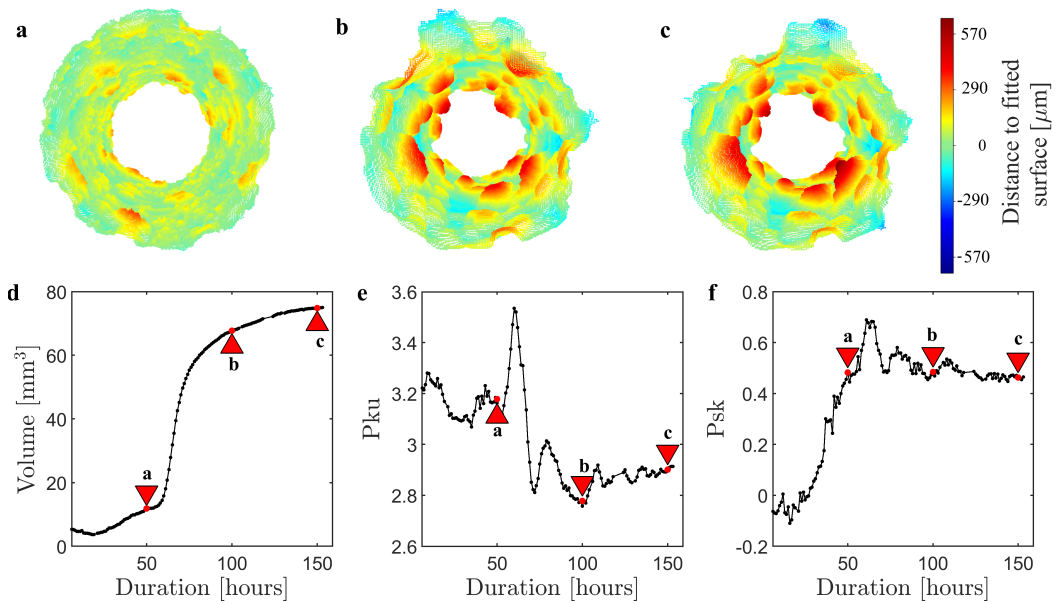


Figure 7: Characteristics of the surface growth for the BaSO₄ segmented surface. **a-c** Quantified height distribution at **a** 50 hours, **b** 100 hours, **c** 150 hours. **d** Volumetric deposition. The volumetric deposition is determined from the internal void volume of the deposition chamber. **e** Pku. The data is presented after filtering with a 3 point moving average. **f** Psk.

height only increases with 10 μm .

From 60 hours to 90 hours, the largest increase in surface coverage occurs. In this period the surface coverage goes from 20% to 80%. Until now, the surface coverage occurs linearly; however, at 60 hours it changes behaviour. The rate of covering slows down as less area is available. The surface coverage after 100 hours of deposition is visualised in Fig. 6b. The scale now covers 85% of the surface, and most of the growth is interconnected. The intensity surface map in Fig. 6b shows clusters growing in the previously observed superstructures, with high-intensity areas present. Both the segmented surface map and the intensity map at 100 hours.

In this phase, the deposited mass changes linearly based on Fig. 6b, even though there is an exponential growth in the lateral and horizontal dimensions (Fig. 6e and Fig. 6f).

Albeit, the rate is much larger than the initiation phase, the rate is constant from 60-100 hours. This indicates that it is the same mechanism concerning the deposition of crystals.

At this point, growth occurs on pristine steel, on top of already formed crystals, and finally on the edges between formed scale and pristine steel. This configuration of growth modes is stable still after 100 hours, even though the surface coverage decreases to 85%, indicating that the contribution of the rate associated with initiation is low. Furthermore, the growth on the edges of the existing crystals and steel is the most dominant, as only a few areas have densified.

3.3. Densification: Scale formation on a covered surface

After initiation and expansion, the crystal growth enters the next phase, densification. Contrary to the expansion phase, the crystal growth does not change its volumetric parameters after 100 hours (Fig. 6d-f). However, as deposition is still occurring it is the density and not the volume of the deposited material that increases. All clusters have merged into one single structure at 125 hours (Fig. 5), and the surface is almost completely covered (Fig. 6e). The surface coverage rate decreases steadily as the surface coverage approaches 100%. Growth can now only occur on the already formed crystal layer at this point. The clusters stop growing in height at 87 μm , having only increased 5 μm from 100 to 150 hours.

While the surface coverage has almost reached 100 %, all clusters have both merged, and stopped growing in height, there is still significant mass deposition. The three different stages are also presented through the instantaneous growth rate (Fig. 8d). In the initiation phase, the mass change was almost constant from 20 to approximately 55 hours. In the expansion phase the growth rate decreased increased rapidly. Lastly, during densification, the instantaneous growth rate decreases however, is still double that in the initiation phase (Fig. 8).

At 100 hours the BaSO_4 is deposited on a low-density phase (Fig. 8b), while at 150 hours, an outline has formed around this phase (Fig. 8c). This outline is the crystal layer, which increased its density. The densified layer is even clearer in Fig. 9. Here the thickness of the layers vary at different positions.

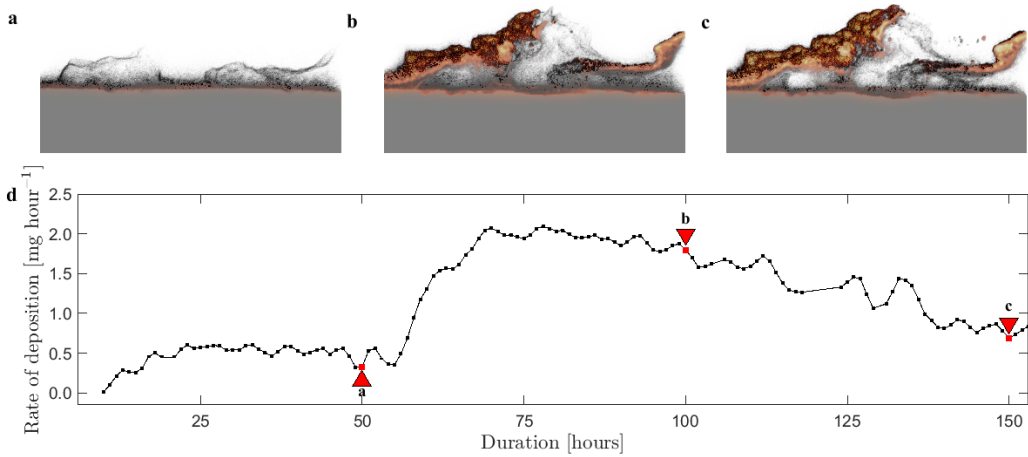


Figure 8: Orthoslices depicting the evolution of scale formation after **a** 50 hours **b** 100 hours **c** 150 hours. **d** The instantaneous rate of mass deposition. The growth rate is presented with a 3 point averaging filter.




Figure 9: Pseudo 3D representation of fouling formation at 150 hours growth. Periodic structures have formed, which would indicate a resonance modulation of the liquid flow field. Hypothesised flow field lines are drawn.

4. Conclusion

Through spatiotemporally resolved measurement of crystal formation, we present new insights into the formation of crystal build-up. We grew crystals in a plug-flow reactor installed in a high-resolution X-ray micro-computed tomography scanner. To mimic realistic fouling conditions, we performed experiments with a low saturation index, which revealed the implications of kinetics. A novel reconstruction method, based on the SIRT-BC algorithm, was used to improve reconstruction quality. The SIRT-BC algorithm uses a high-quality pre-experiment scan to improve reconstruction quality and speed. Analysis was performed of two segmented areas to characterise surface growth and corrosion features. The characterisation was done through volume deposition and morphological parameters, such as surface coverage, cluster size distribution, cluster height, spikiness of the height distribution.

We decouple three growth mechanisms present in crystallisation fouling build-up and show how they independently affect volume and mass deposition. We discovered an initiation phase, characterised by the formation of many crystal clusters, an expansion phase, where the volumetric growth rate is larger than the gravimetric rate. Lastly, we show a densification phase, where incoming ions are closer packed, and the gravimetric rate is larger than the volumetric.

The crystal formation was visualised through in-situ X-ray CT scanning and characterised through the total surface area, surface coverage, crystal height, deposited mass, deposited volume, and a range of texture parameters. The total surface area revealed three phases of growth. Firstly the surface area gradually shrank as the uniform deposition occurred. Secondly, the surface area rose as spikey features were formed. Lastly, the surface area stagnated as growth was incorporated into existing crystals rather than new structures. Deposition rates were identified for these phases, where growth initiated with a rate of 4 mg per hour, increased to 15 mg per hour and steadily decreased to 6 mg per hour 80 to 150 hours.

Predicting crystallisation formation is based on accurate knowledge of the growth rate in the system. Our findings show how the growth rate can vary 300 % from initiation to the end. Even though we performed measurements for 150 hours, the rate at which growth formed was not constant. This non-equilibrium effect reflects the complexity of crystallisation formation in natural systems. These results allow for comprehensive prediction models to be developed, capable of accounting for the various stages present in crystallisation build-up.

Acknowledgements

The authors acknowledge the Danish Underground Consortium (TotalEnergies, E&P Denmark, Noreco and Nordsofonden) for funding the Danish Hydrocarbon Research and Technology Centre (DHRTC). The Advanced Water Flooding and the Extended Service Life of Wells programmes of DHRTC funded the current work. The 3D Imaging Centre at The Technical University of Denmark is gratefully acknowledged

for providing access to X-ray tomography equipment. We would like to thank N. Jeppesen for extensive assistance with graph cut methods.

References

- [1] K. M. Sassi, I. M. Mujtaba, Optimal design and operation of reverse osmosis desalination process with membrane fouling, *Chemical Engineering Journal* 171 (2) (2011) 582–593. doi:10.1016/j.cej.2011.04.034.
URL <http://dx.doi.org/10.1016/j.cej.2011.04.034>
- [2] R. Sharifian, R. M. Wagterveld, I. A. Digdaya, C. Xiang, D. A. Vermaas, Electrochemical carbon dioxide capture to close the carbon cycle, *Energy & Environmental Science* 14 (2) (2021) 781–814. doi:10.1039/D0EE03382K.
URL <http://xlink.rsc.org/?DOI=D0EE03382K>
- [3] J. Li, M. Tang, Z. Ye, L. Chen, Y. Zhou, Scale formation and control in oil and gas fields: A review, *Journal of Dispersion Science and Technology* 38 (5) (2017) 661–670. doi:10.1080/01932691.2016.1185953.
URL <https://www.tandfonline.com/action/journalInformation?journalCode=ldis20http://dx.doi.org/10.1080/01932691.2016.1185953>
- [4] Q. Zhao, Y. Liu, C. Wang, S. Wang, H. Müller-Steinhagen, Effect of surface free energy on the adhesion of biofouling and crystalline fouling, *Chemical Engineering Science* 60 (17) (2005) 4858–4865. doi:10.1016/j.ces.2005.04.006.
- [5] J. Berce, M. Zupančič, M. Može, I. Golobič, A review of crystallization fouling in heat exchangers, *Processes* 9 (8) (2021). doi:10.3390/pr9081356.
- [6] R. Boch, A. Leis, E. Haslinger, J. E. Goldbrunner, F. Mittermayr, H. Fröschl, D. Hippler, M. Dietzel, Scale-fragment formation impairing geothermal energy production: interacting H₂S corrosion and CaCO₃ crystal growth, *Geothermal Energy* 5 (1) (2017) 4. doi:10.1186/s40517-017-0062-3.
- [7] N. EPSTEIN, Thinking about Heat Transfer Fouling: A 5 × 5 Matrix, *Heat Transfer Engineering* 4 (1) (1983) 43–56. doi:10.1080/01457638108939594.
URL <http://www.tandfonline.com/doi/abs/10.1080/01457638108939594>
- [8] J. Li, C. J. Tilbury, M. N. Joswiak, B. Peters, M. F. Doherty, Rate Expressions for Kink Attachment and Detachment during Crystal Growth, *Crystal Growth and Design* 16 (6) (2016) 3313–3322. doi:10.1021/acs.cgd.6b00292.
- [9] B. U. Anabaraonye, J. R. Bentzon, I. Khaliqdad, K. L. Feilberg, S. I. Andersen, J. H. Walther, The influence of turbulent transport in reactive processes: A combined numerical and experimental investigation in a Taylor-Couette reactor, *Chemical Engineering Journal* (2021) 129591doi:10.1016/j.cej.2021.129591.
URL <https://doi.org/10.1016/j.cej.2021.129591https://linkinghub.elsevier.com/retrieve/pii/S1385894721011785>
- [10] I. A. Løge, J. R. Bentzon, C. G. Klingaa, J. H. Walther, B. U. Anabaraonye, P. L. Fosbøl, Scale attachment and detachment: The role of hydrodynamics and surface morphology, *Chemical Engineering Journal* 430 (P2) (2021) 132583. doi:10.1016/j.cej.2021.132583.
URL <https://doi.org/10.1016/j.cej.2021.132583>
- [11] A. Yi-Tsung Lu, A. T. Kan, M. B. Tomson, Nucleation and Crystallization Kinetics of Barium Sulfate in the Hydrodynamic Boundary Layer: An Explanation of Mineral Deposition, *Crystal Growth and Design* 21 (3) (2021) 1443–1450. doi:10.1021/acs.cgd.0c01027.
- [12] M. M. Vazirian, T. V. Charpentier, M. de Oliveira Penna, A. Neville, Surface inorganic scale formation in oil and gas industry: As adhesion and deposition processes, *Journal of Petroleum Science and Engineering* 137 (2016) 22–32. doi:10.1016/j.petrol.2015.11.005.
URL <http://dx.doi.org/10.1016/j.petrol.2015.11.005>

- [13] I. A. Løge, B. U. Anabaraonye, P. L. Føsbøl, Growth mechanisms for composite fouling: The impact of substrates on detachment processes, in review (2022).
- [14] A. Quddus, I. Allam, BaSO₄ scale deposition on stainless steel, *Desalination* 127 (3) (2000) 219–224. doi:10.1016/S0011-9164(00)00012-6.
URL <https://linkinghub.elsevier.com/retrieve/pii/S0011916400000126>
- [15] G. M. Graham, I. R. Collins, R. Stalker, I. J. Littlehales, The Importance of Appropriate Laboratory Procedures for the Determination of Scale Inhibitor Performance, in: *Society of Petroleum Engineers - International Symposium on Oilfield Scale 2002*, 2002, pp. 1–14. doi:10.2523/74679-ms.
- [16] M. Barber, S. Heath, A New Approach to Testing Scale Inhibitors in Mild Scaling Brines – Are Dynamic Scale Loop Tests Needed?, in: *SPE*, 2019, pp. 1–14. doi:10.16309/j.cnki.issn.1007-1776.2003.03.004.
- [17] Y. M. Al Rawahi Feroz Shaik, L. Nageswara Rao, Studies on Scale Deposition in Oil Industries & Their Control, *IJIRST-International Journal for Innovative Research in Science & Technology* 3 (2017).
URL www.ijirst.org
- [18] E. Mavredaki, A. Neville, K. S. Sorbie, Initial stages of barium sulfate formation at surfaces in the presence of inhibitors, *Crystal Growth and Design* 11 (11) (2011) 4751–4758. doi:10.1021/cg101584f.
URL <https://pubs.acs.org/sharingguidelines>
- [19] J. R. Godinho, K. M. Gerke, A. G. Stack, P. D. Lee, The dynamic nature of crystal growth in pores, *Scientific Reports* 6 (2016). doi:10.1038/srep33086.
URL www.nature.com/scientificreports
- [20] D. F. Oliveira, R. S. Santos, A. S. Machado, A. S. Silva, M. J. Anjos, R. T. Lopes, Characterization of scale deposition in oil pipelines through X-Ray Microfluorescence and X-Ray microtomography, *Applied Radiation and Isotopes* 151 (2019) 247–255. doi:10.1016/j.apradiso.2019.06.019.
URL <https://doi.org/10.1016/j.apradiso.2019.06.019>
- [21] N. A. Webster, B. K. Gan, I. Livk, In situ X-ray diffraction analysis of the onset of mineral scale deposition from synthetic oilfield processing waters, *Fuel* 137 (2014) 211–215. doi:10.1016/j.fuel.2014.07.089.
URL <http://dx.doi.org/10.1016/j.fuel.2014.07.089>
- [22] E. Mavredaki, A. Neville, K. S. Sorbie, Using Synchrotron X-Ray Diffraction (SXRD) for studying the BaSO₄ formation kinetics and the effect of inhibitors on barite formation, in: *AIP Conference Proceedings*, Vol. 1202, 2009, pp. 94–101. doi:10.1063/1.3295618.
URL <https://doi.org/10.1063/1.3295618>
- [23] R. Barker, D. Burkle, T. Charpentier, H. Thompson, A. Neville, A review of iron carbonate (FeCO₃) formation in the oil and gas industry, *Corrosion Science* 142 (October) (2018) 312–341. doi:10.1016/j.corsci.2018.07.021.
- [24] O. Sanni, O. Bukuaghangin, M. Huggan, N. Kapur, T. Charpentier, A. Neville, Development of a novel once-Through flow visualization technique for kinetic study of bulk and surface scaling, *Review of Scientific Instruments* 88 (10) (2017) 103903. doi:10.1063/1.4991729.
URL <https://doi.org/10.1063/1.4991729>
- [25] O. Bukuaghangin, O. Sanni, N. Kapur, M. Huggan, A. Neville, T. Charpentier, Kinetics study of barium sulphate surface scaling and inhibition with a once-through flow system, *Journal of Petroleum Science and Engineering* 147 (2016) 699–706. doi:10.1016/j.petrol.2016.09.035.
URL <http://dx.doi.org/10.1016/j.petrol.2016.09.035>
- [26] K. S. Pitzer, *Activity Coefficients in Electrolyte Solutions*, 2nd Edition, CRC press, 1991. doi:10.1201/9781351069472.
- [27] D. Appelo, C. Parkhurst, User's Guide to PHREEQC (Version 2) A Computer Program for Speciation, Batch Reaction, One-Dimensional Transport, and Inverse Geochemical Calculations, Tech. Rep. Version 2, Toxic Substances Hydrology

- Program (1999).
- [28] L. A. Feldkamp, L. C. Davis, J. W. Kress, Practical cone-beam algorithm, *Journal of the Optical Society of America A* 1 (6) (1984) 612. doi:10.1364/josaa.1.000612.
 - [29] F. E. Boas, D. Fleischmann, et al., Ct artifacts: causes and reduction techniques, *Imaging Med* 4 (2) (2012) 229–240.
 - [30] E. Maire, P. J. Withers, Quantitative X-ray tomography, *International Materials Reviews* 59 (1) (2013) 1–43. doi:10.1179/1743280413y.0000000023.
 - [31] P. W. Rasmussen, H. O. Sørensen, S. Bruns, A. B. Dahl, A. N. Christensen, Improved dynamic imaging of multiphase flow by constrained tomographic reconstruction, *Scientific Reports* 11 (1) (2021) 1–14. doi:10.1038/s41598-021-91776-1. URL <http://www.nature.com/articles/s41598-021-91776-1>
 - [32] P. Gilbert, Iterative methods for the three-dimensional reconstruction of an object from projections, *Journal of Theoretical Biology* 36 (1) (1972) 105–117. doi:10.1016/0022-5193(72)90180-4.
 - [33] J. Gregor, T. Benson, Computational analysis and improvement of SIRT, *IEEE Transactions on Medical Imaging* 27 (7) (2008) 918–924. doi:10.1109/TMI.2008.923696. URL <http://ieeexplore.ieee.org/document/4492754/>
 - [34] P. C. Hansen, M. E. Kilmer, R. H. Kjeldsen, Exploiting residual information in the parameter choice for discrete ill-posed problems, *BIT Numerical Mathematics* 46 (1) (2006) 41–59. doi:10.1007/s10543-006-0042-7. URL <http://link.springer.com/10.1007/s10543-006-0042-7>
 - [35] N. Jeppesen, A. N. Christensen, V. A. Dahl, A. B. Dahl, Sparse layered graphs for multi-object segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020*, pp. 12777–12785. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Jeppesen_Sparse_Layered_Graphs_for_Multi-Object_Segmentation_CVPR_2020_paper.html
 - [36] N. Jeppesen, P. M. Jensen, A. N. Christensen, A. B. Dahl, V. A. Dahl, Faster multi-object segmentation using parallel quadratic pseudo-boolean optimization, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2021*, pp. 6260–6269. URL https://openaccess.thecvf.com/content/ICCV2021/html/Jeppesen_Faster_Multi-Object_Segmentation_Using_Parallel_Quadratic_Pseudo-Boolean_Optimization_ICCV_2021_paper.html
 - [37] S. Bruns, S. L. Stipp, H. O. Sørensen, Looking for the Signal: A guide to iterative noise and artefact removal in X-ray tomographic reconstructions of porous geomaterials, *Advances in Water Resources* 105 (2017) 96–107. doi:10.1016/j.advwatres.2017.04.020. URL <https://linkinghub.elsevier.com/retrieve/pii/S030917081630598X>
 - [38] E. S. Gadelmawla, M. M. Koura, T. M. Maksoud, I. M. Elewa, H. H. Soliman, Roughness parameters, *Journal of Materials Processing Technology* 123 (1) (2002) 133–145. doi:10.1016/S0924-0136(02)00060-2. URL <https://www.taylorfrancis.com/books/9781420069099https://linkinghub.elsevier.com/retrieve/pii/S0924013602000602>
 - [39] C. G. Klingaa, M. K. Bjerre, S. Baier, L. De Chiffre, S. Mohanty, J. H. Hattel, Roughness investigation of slm manufactured conformal cooling channels using x-ray computed tomography, in: *Proceedings of the 9th Conference on Industrial Computed Tomography (ICT 2019), Padova, Italy, 2019*, pp. 13–15.
 - [40] C. G. Klingaa, T. Dahmen, S. Baier, S. Mohanty, J. H. Hattel, X-ray CT and image analysis methodology for local roughness characterization in cooling channels made by metal additive manufacturing, *Additive Manufacturing* 32 (February) (2020) 101032. doi:10.1016/j.addma.2019.101032. URL <https://doi.org/10.1016/j.addma.2019.101032>
 - [41] R. L. Townsin, C. D. Anderson, Fouling control coatings using low surface energy, foul release technology, Woodhead Publishing Limited, 2009. doi:10.1533/9781845696313.4.693.

URL <http://dx.doi.org/10.1533/9781845696313.4.693>

- [42] N. Gathimba, Y. Kitane, T. Yoshida, Y. Itoh, Surface roughness characteristics of corroded steel pipe piles exposed to marine environment, *Construction and Building Materials* 203 (2019) 267–281. doi:10.1016/j.conbuildmat.2019.01.092. URL <https://doi.org/10.1016/j.conbuildmat.2019.01.092>
- [43] H. Azaza, L. Mechi, A. Doggaz, V. Optasanu, M. Tlili, M. B. Amor, Calcite and barite precipitation in CaCO₃-BaSO₄-NaCl and BaSO₄-NaCl-CaCl₂ aqueous systems: kinetic and microstructural study, *Arabian Journal of Geosciences* 10 (10) (2017). doi:10.1007/s12517-017-3005-1.
- [44] A. Ismail, N. Adan, Effect of Oxygen Concentration on Corrosion Rate of Carbon Steel in Seawater, *American Journal of Engineering Research* 03 (01) (2014) 64–67.

Conclusion

This thesis represents the conclusion of my PhD studies, which constitutes more than three years of research. The project was carried out with the Danish Offshore Technology Centre (DTU Offshore) as a part of their research on multiphase flow in porous media. In particular, we sought to assist DTU Offshore with carrying out dynamic X-ray computed tomography (CT) experiments. Dynamic CT experiments have multiple uses within the oil and gas industry but are especially suited for core flooding and scale formation experiments.

We have made significant contributions during this project that enables future research within these fields. Additionally, the methods developed are not only useful for the gas industry but are also relevant for fundamental research within carbon storage, compaction and faulting of porous media and scaling processes.

In **Contribution A**, we developed a range of tomographic reconstruction algorithms designed to preserve image quality for dynamic CT experiments despite noisy or sparsely sampled projection data. The robustness of the algorithm was achieved in two ways. Firstly, all dynamic reconstructions are initialised with a reconstruction of the previous time step. The first dynamic reconstruction is initialised with a high-quality scan of the experiment taken before the dynamic experiment is started. Secondly, the reconstruction is constrained such that pixels are only allowed to take on specific values or ranges of values. The constraints are derived from a segmentation of the high-quality scan. The algorithm performed excellently compared to traditional reconstruction algorithms, based on numerical experiments, and using it rather than traditional algorithms can greatly increase the temporal resolution of a CT scanner.

In **Contribution B**, we presented and tested four stopping rules that provide a heuristic for when to terminate algebraic reconstruction algorithms. This is necessary as these algorithms do not converge at the point when the optimal reconstruction is reached, and iterating beyond this point only increases the noise in the reconstruction. Numerical experiments revealed that all methods we tested provided an accurate estimate of the optimal stopping point. Two of the four methods do not require any information about the noise level in the data and can therefore be readily applied.

In **Contribution C** crystallisation fouling was examined using 4D X-ray tomogra-

phy. A flow cell consisting of a small carbon steel cylinder was injected with distilled water supersaturated with BaSO_4 for more than 150 hours while being imaged in a CT scanner. We used a modified version of one of the algorithms developed in [Contribution A](#) to reconstruct the dynamic data. Using this algorithm improved the signal to noise ratio of the reconstructions, which allow further quantitative analysis. Based on the reconstructions, we characterised the growth and morphology of the scale as it changed during the experiment. This further enhances the applicability of [contribution A](#), as we now have a detailed study of how terminate the reconstruction algorithms.

To summerise, [Contribution A](#) and [Contribution B](#) provide tools that make it possible to perform dynamic CT experiments with laboratory sources. These methods are used in [Contribution C](#), which made it possible to obtain novel insights in the scale formation process. Based on these contributions, we conclude that this thesis shows that dynamic X-ray tomography is a valuable metrology that can provide novel insights in a wide range of research fields.

APPENDIX **A**

CT Scanners at DTU Offshore

DTU Offshore has two CT scanners available for imaging, and one of these scanners have been used extensively throughout this project. This scanner is colloquially referred to as the nanoCT. The source of the scanner is a Hamamatsu microfocus X-ray source L10711-03. The source has two cathodes, one made of LaB₆ and another in tungsten, that allows for a tube voltage of either 20 kV to 100 kV or 20 kV to 160 kV. Switching between these cathodes allow for either high resolution (LaB₆ cathode) or high intensity (tungsten cathode).

The X-ray radiation is detected with a Varian PaxScan 4343CB flat-panel detector. The detector has 3072×3072 pixels, each with a pitch of $0.139 \mu\text{m}$ that results in a total pixel area of $42.7 \text{ cm} \times 42.7 \text{ cm}$. The detected radiation is converted to visible light with a thallium doped caesium iodide (CsI:TI) scintillator that is detected by amorphous silicon pixels.

A full list of the scanner components is shown below:

- X-ray radiation Source: [Hamamatsu microfocus X-ray source – L10711-03](#)
- Source translation: Newport linear stage (unknown model – possibly a UTS stage)
- Rotation stage: [Aerotech ABRT-150](#)
- Detector: [Varian PaxScan 4343CB](#)
- Detector stage: [Newport M-IMS400LM-S](#)
- Translation stage: [Airpel-AB Air Bearing Cylinders](#)
- Newport XPS motion controller

APPENDIX B

Sample Movement Corrections for Contribution C

In Figures B.1 and B.2 we see the a cross-section of the flow cell used in Contribution C before and after the reconstruction accounted for the movement of the sample. The difference is subtle, but there are streak artefacts along the left and right sides of the flow cell. These artefacts are not present in the corrected reconstruction.




Figure B.1: Reconstruction of the flow cell performed without sample movement corrections.




Figure B.2: Reconstruction of the flow cell performed with sample movement corrections.

Bibliography

- Aarle, W., W. J. Palenstijn, J. D. Beenhouwer, et al. (2015). “The ASTRA Toolbox: A platform for advanced algorithm development in electron tomography.” In: *Ultramicroscopy* 157, pages 35–47. ISSN: 0304-3991. DOI: [10.1016/j.ultramic.2015.05.002](https://doi.org/10.1016/j.ultramic.2015.05.002).
- Aarle, W., W. J. Palenstijn, J. Cant, et al. (October 2016). “Fast and flexible X-ray tomography using the ASTRA toolbox.” In: *Opt. Express* 24.22, pages 25129–25147. DOI: [10.1364/OE.24.025129](https://doi.org/10.1364/OE.24.025129).
- Ambrose, J. (December 1973). “Computerized transverse axial scanning (tomography): II. Clinical application.” In: *British Journal of Radiology* 46.552, pages 1023–1047. ISSN: 00071285. DOI: [10.1259/0007-1285-46-552-1023](https://doi.org/10.1259/0007-1285-46-552-1023).
- American Physical Society (2001). *This Month in Physics History: November 8, 1895: Roentgen’s Discovery of X-Rays*. Accessed: 21.02.2022. URL: <https://www.aps.org/publications/apsnews/200111/history.cfm>.
- Armstrong, R. et al. (October 2016). “Beyond Darcy’s law: The role of phase topology and ganglion dynamics for two-fluid flow.” In: *Physical Review E* 94.4, page 043113. ISSN: 2470-0045. DOI: [10.1103/PhysRevE.94.043113](https://doi.org/10.1103/PhysRevE.94.043113).
- Barker, R. et al. (September 2018). “A review of iron carbonate (FeCO₃) formation in the oil and gas industry.” In: *Corrosion Science* 142, pages 312–341. ISSN: 0010938X. DOI: [10.1016/j.corsci.2018.07.021](https://doi.org/10.1016/j.corsci.2018.07.021).
- Batenburg, K. J. and J. Sijbers (2011). “DART: a practical reconstruction algorithm for discrete tomography.” In: *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* 20.9, pages 2542–2553. ISSN: 19410042. DOI: [10.1109/TIP.2011.2131661](https://doi.org/10.1109/TIP.2011.2131661).
- Beevil (2017). *Resolution in direct and indirect x-ray detectors*. Accessed: 30.01.2022. URL: https://commons.wikimedia.org/wiki/File:Resolution_in_direct_and_indirect_x-ray_detectors.svg.
- Beister, M., D. Kolditz, and W. Kalender (2012). “Iterative reconstruction methods in X-ray CT.” In: *Physica Medica* 28.2, pages 94–108. ISSN: 11201797. DOI: [10.1016/j.ejmp.2012.01.003](https://doi.org/10.1016/j.ejmp.2012.01.003).
- Ben-Zion, Y. (December 2008). “Collective behavior of earthquakes and faults: Continuum-discrete transitions, progressive evolutionary changes, and different dynamic regimes.” In: *Reviews of Geophysics* 46.4, RG4006. ISSN: 87551209. DOI: [10.1029/2008RG000260](https://doi.org/10.1029/2008RG000260).

- Berg, S. et al. (March 2013). “Real-time 3D imaging of Haines jumps in porous media flow.” In: *Proceedings of the National Academy of Sciences* 110.10, pages 3755–3759. ISSN: 0027-8424. DOI: [10.1073/pnas.1221373110](https://doi.org/10.1073/pnas.1221373110).
- Berger, M. et al. (January 2022). *XCOM: Photon Cross Section Database (version 1.5)*. Originally published as Berger, M.J. and Hubbell, J.H., “XCOM: Photon Cross Sections on a Personal Computer,” NBSIR 87-3597, National Bureau of Standards (former name of NIST), Gaithersburg, MD (1987); and as Berger, M.J. and Hubbell, J.H., “NIST X-ray and Gamma-ray Attenuation Coefficients and Cross Sections Database,” NIST Standard Reference Database 8, Version 2.0, National Institute of Standards and Technology, Gaithersburg, MD (1990). DOI: [10.18434/T48G6X](https://doi.org/10.18434/T48G6X). URL: <http://physics.nist.gov/xcom>.
- Bilderback, D., P. Elleaume, and E. Weckert (May 2005). “Review of third and next generation synchrotron light sources.” In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 38.9, S773–S797. ISSN: 09534075. DOI: [10.1088/0953-4075/38/9/022](https://doi.org/10.1088/0953-4075/38/9/022).
- Blunt, M. et al. (2013). “Pore-scale imaging and modelling.” In: *Advances in Water Resources* 51, pages 197–216. ISSN: 03091708. DOI: [10.1016/j.advwatres.2012.03.003](https://doi.org/10.1016/j.advwatres.2012.03.003).
- Boin, M. and A. Haibel (2006). “Compensation of ring artefacts in synchrotron tomographic images.” In: *Optics Express* 14.25, page 12071. ISSN: 1094-4087. DOI: [10.1364/oe.14.012071](https://doi.org/10.1364/oe.14.012071).
- Boone, J. and J. Seibert (1997). “An accurate method for computer-generating tungsten anode x-ray spectra from 30 to 140 kV.” In: *Medical Physics* 24.11, pages 1661–1670. DOI: [10.1118/1.597953](https://doi.org/10.1118/1.597953).
- BP (January 2020). *Energy charting tool – app*. Edited by BP p.l.c. Accessed: 17.01.2022. URL: <https://www.bp.com/en/global/corporate/energy-economics/energy-charting-tool-desktop.html>.
- Bührer, M. et al. (December 2021). “Deep learning based classification of dynamic processes in time-resolved X-ray tomographic microscopy.” In: *Scientific Reports* 11.1, page 24174. ISSN: 2045-2322. DOI: [10.1038/s41598-021-03546-8](https://doi.org/10.1038/s41598-021-03546-8).
- Bultreys, T., W. De Boever, and V. Cnudde (April 2016). “Imaging and image-based fluid transport modeling at the pore scale in geological materials: A practical introduction to the current state-of-the-art.” In: *Earth-Science Reviews* 155, pages 93–128. ISSN: 00128252. DOI: [10.1016/j.earscirev.2016.02.001](https://doi.org/10.1016/j.earscirev.2016.02.001).
- Buzug, T. (2008). *Computed tomography: From photon statistics to modern cone-beam CT*. Berlin, Heidelberg: Springer Berlin Heidelberg, pages 1–521. ISBN: 9783540394075. DOI: [10.1007/978-3-540-39408-2](https://doi.org/10.1007/978-3-540-39408-2).
- Chen, X. et al. (December 2020). “The dynamic evolution of compaction bands in highly porous carbonates: the role of local heterogeneity for nucleation and propagation.” In: *Progress in Earth and Planetary Science* 7.1, page 28. ISSN: 2197-4284. DOI: [10.1186/s40645-020-00344-0](https://doi.org/10.1186/s40645-020-00344-0).
- Cimmino, G. (1938). “Calcolo approssimato per le soluzioni dei sistemi di equazioni lineari.” In: *La Ricerca Scientifica* 9, pages 326–333.

- Cnudde, V. and M. Boone (2013). “High-resolution X-ray computed tomography in geosciences: A review of the current technology and applications.” In: *Earth-Science Reviews* 123, pages 1–17. ISSN: 00128252. DOI: [10.1016/j.earscirev.2013.04.003](https://doi.org/10.1016/j.earscirev.2013.04.003).
- Collette, A. et al. (October 2021). *h5py 3.5.0*. Version 3.5.0. DOI: [10.5281/zenodo.5585380](https://doi.org/10.5281/zenodo.5585380).
- Connolly, B. J. et al. (September 2006). “X-ray microtomography studies of localised corrosion and transitions to stress corrosion cracking.” In: *Materials Science and Technology* 22.9, pages 1076–1085. ISSN: 02670836. DOI: [10.1179/174328406X114199](https://doi.org/10.1179/174328406X114199).
- Coolth (2010). *X-ray tube (scheme)*. Accessed: 12.02.2022. URL: <https://commons.wikimedia.org/wiki/File:WaterCooledXrayTube.svg>.
- Croton, L. et al. (May 2019). “Ring artifact suppression in X-ray computed tomography using a simple, pixel-wise response correction.” In: *Optics Express* 27.10, page 14231. ISSN: 1094-4087. DOI: [10.1364/oe.27.014231](https://doi.org/10.1364/oe.27.014231).
- Danish Energy Agency (January 2017). *RProduction*. Technical report. Danish Energy Agency.
- (January 2021a). *About oil and gas*. Edited by Danish Energy Agency. Accessed: 25.01.2022. URL: <https://ens.dk/en/our-responsibilities/oil-gas/about-oil-and-gas>.
- (August 2021b). *Monthly and yearly production*. Edited by Danish Energy Agency. Accessed: 17.01.2022. URL: <https://ens.dk/en/our-services/oil-and-gas-related-data/monthly-and-yearly-production>.
- Dembinski, H. et al. (January 2022). *scikit-hep/iminuit: v2.9.0*. Version v2.9.0. DOI: [10.5281/zenodo.5831665](https://doi.org/10.5281/zenodo.5831665).
- Duan, X. et al. (October 2013). “Electronic noise in CT detectors: Impact on image noise and artifacts.” In: *American Journal of Roentgenology* 201.4, W626–W632. ISSN: 0361803X. DOI: [10.2214/AJR.12.10234](https://doi.org/10.2214/AJR.12.10234).
- Ebel, H. (July 1999). “X-ray Tube Spectra.” In: *X-Ray Spectrometry* 28.4, pages 255–266. ISSN: 00498246. DOI: [10.1002/\(SICI\)1097-4539\(199907/08\)28:4<255::AID-XRS347>3.0.CO;2-Y](https://doi.org/10.1002/(SICI)1097-4539(199907/08)28:4<255::AID-XRS347>3.0.CO;2-Y).
- Eftekhari, A. and R. Farajzadeh (April 2017). “Effect of Foam on Liquid Phase Mobility in Porous Media.” In: *Scientific Reports* 7.1, page 43870. ISSN: 2045-2322. DOI: [10.1038/srep43870](https://doi.org/10.1038/srep43870).
- Feeman, T. G. (2015a). “Back Projection.” In: pages 39–46. DOI: [10.1007/978-3-319-22665-1_3](https://doi.org/10.1007/978-3-319-22665-1_3).
- (2015b). “Filters and Convolution.” In: pages 77–95. DOI: [10.1007/978-3-319-22665-1_7](https://doi.org/10.1007/978-3-319-22665-1_7).
- (2015c). “The Radon Transform.” In: pages 13–37. DOI: [10.1007/978-3-319-22665-1_2](https://doi.org/10.1007/978-3-319-22665-1_2).
- (2015d). “Two Big Theorems.” In: pages 71–76. DOI: [10.1007/978-3-319-22665-1_7](https://doi.org/10.1007/978-3-319-22665-1_7).
- Feldkamp, L., L. Davis, and J. Kress (June 1984). “Practical cone-beam algorithm.” In: *Journal of the Optical Society of America A* 1.6, page 612. ISSN: 1084-7529. DOI: [10.1364/josaa.1.000612](https://doi.org/10.1364/josaa.1.000612).

- Ferrucci, M. et al. (2018). “Measurement of the X-ray computed tomography instrument geometry by minimization of reprojection errors—Implementation on simulated data.” In: *Precision Engineering* 54.March, pages 7–20. ISSN: 01416359. DOI: [10.1016/j.precisioneng.2018.03.012](https://doi.org/10.1016/j.precisioneng.2018.03.012).
- Gao, Y. et al. (2020). “Pore-scale dynamics and the multiphase Darcy law.” In: *Physical Review Fluids* 5.1, page 13801. ISSN: 2469990X. DOI: [10.1103/PhysRevFluids.5.013801](https://doi.org/10.1103/PhysRevFluids.5.013801).
- Gilbert, P. (1972). “Iterative methods for the three-dimensional reconstruction of an object from projections.” In: *Journal of Theoretical Biology* 36.1, pages 105–117. ISSN: 10958541. DOI: [10.1016/0022-5193\(72\)90180-4](https://doi.org/10.1016/0022-5193(72)90180-4).
- Godinho, J. and P. Withers (February 2018). “Time-lapse 3D imaging of calcite precipitation in a microporous column.” In: *Geochimica et Cosmochimica Acta* 222, pages 156–170. ISSN: 00167037. DOI: [10.1016/j.gca.2017.10.024](https://doi.org/10.1016/j.gca.2017.10.024).
- Gordon, R., R. Bender, and G. Herman (1970). “Algebraic Reconstruction Techniques (ART) for three-dimensional electron microscopy and X-ray photography.” In: *Journal of Theoretical Biology* 29.3, pages 471–481. ISSN: 0022-5193. DOI: [10.1016/0022-5193\(70\)90109-8](https://doi.org/10.1016/0022-5193(70)90109-8).
- Gregor, J. and T. Benson (July 2008). “Computational analysis and improvement of SIRT.” In: *IEEE Transactions on Medical Imaging* 27.7, pages 918–924. ISSN: 02780062. DOI: [10.1109/TMI.2008.923696](https://doi.org/10.1109/TMI.2008.923696).
- Gruner, S., M. Tate, and E. Eikenberry (August 2002). “Charge-coupled device area x-ray detectors.” In: *Review of Scientific Instruments* 73.8, page 2815. ISSN: 00346748. DOI: [10.1063/1.1488674](https://doi.org/10.1063/1.1488674).
- Haaksman, V. A. et al. (January 2017). “Characterization of feed channel spacer performance using geometries obtained by X-ray computed tomography.” In: *Journal of Membrane Science* 522, pages 124–139. ISSN: 18733123. DOI: [10.1016/j.memsci.2016.09.005](https://doi.org/10.1016/j.memsci.2016.09.005).
- Hansen, P. C., J. S. Jørgensen, and P. W. Rasmussen (2021). “Stopping Rules for Algebraic Iterative Reconstruction Methods in Computed Tomography.” In: *The 21st International Conference on Computational Science and its Applications: Application of Numerical Analysis to Imaging Science*.
- Harris, C. R. et al. (September 2020). “Array programming with NumPy.” In: *Nature* 585.7825, pages 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2).
- Hemberg, O. (July 2004). “Liquid-metal-jet anode x-ray tube.” In: *Optical Engineering* 43.7, page 1682. ISSN: 0091-3286. DOI: [10.1117/1.1737787](https://doi.org/10.1117/1.1737787).
- Hendriksen, A. A., M. Bühner, et al. (December 2021). “Deep denoising for multi-dimensional synchrotron X-ray tomography without high-quality reference data.” In: *Scientific Reports* 11.1, page 11895. ISSN: 20452322. DOI: [10.1038/s41598-021-91084-8](https://doi.org/10.1038/s41598-021-91084-8).
- Hendriksen, A. A., D. M. Pelt, and K. J. Batenburg (2020). “Noise2Inverse: Self-Supervised Deep Convolutional Denoising for Tomography.” In: *IEEE Transactions on Computational Imaging* 6, pages 1320–1335. ISSN: 2333-9403. DOI: [10.1109/TCI.2020.3019647](https://doi.org/10.1109/TCI.2020.3019647).

- Herbst, M. et al. (March 2019). “Misalignment compensation for ultra-high-resolution and fast CBCT acquisitions.” In: *Medical Imaging 2019: Physics of Medical Imaging*. Edited by H. Bosmans, G.-H. Chen, and T. Gilat Schmidt. March 2019. SPIE, page 57. ISBN: 9781510625433. DOI: 10.1117/12.2513276.
- Herman, G. and L. Meyer (1993). “Algebraic Reconstruction Techniques Can Be Made Computationally Efficient.” In: *IEEE Transactions on Medical Imaging* 12.3, pages 600–609. ISSN: 1558254X. DOI: 10.1109/42.241889.
- Hoheisel, M. (July 2006). “Review of medical imaging with emphasis on X-ray detectors.” In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 563.1, pages 215–224. ISSN: 01689002. DOI: 10.1016/j.nima.2006.01.123.
- Holbrow, C. H. et al. (2009a). “The Granularity of Light.” In: *Modern Introductory Physics*. New York, NY: Springer New York, pages 401–419. DOI: 10.1007/978-0-387-79080-0_13.
- (2009b). “X-Rays.” In: *Modern Introductory Physics*. New York, NY: Springer New York, pages 421–454. DOI: 10.1007/978-0-387-79080-0_14.
- Hounsfield, G. N. (December 1973). “Computerized transverse axial scanning (tomography): I. Description of system.” In: *British Journal of Radiology* 46.552, pages 1016–1022. ISSN: 00071285. DOI: 10.1259/0007-1285-46-552-1016.
- Huang, Y. et al. (2020). “Limited angle tomography for transmission X-ray microscopy using deep learning.” In: *Journal of Synchrotron Radiation* 27.2, pages 477–485. ISSN: 1600-5775. DOI: 10.1107/s160057752000017x.
- International Energy Agency (2018). *The Future of Petrochemical*. URL: <https://www.iea.org/reports/the-future-of-petrochemicals>.
- Jacobson, M. W., M. Ketcha, et al. (March 2017). “Geometric calibration using line fiducials for cone-beam CT with general, non-circular source-detector trajectories.” In: *Medical Imaging 2017: Physics of Medical Imaging*. Edited by T. G. Flohr, J. Y. Lo, and T. Gilat Schmidt. Volume 10132. March 2017, page 101320I. ISBN: 9781510607095. DOI: 10.1117/12.2255724.
- Jacobson, M. W., M. D. Ketcha, et al. (January 2018). “A line fiducial method for geometric calibration of cone-beam CT systems with diverse scan trajectories.” In: *Physics in Medicine & Biology* 63.2, page 025030. ISSN: 1361-6560. DOI: 10.1088/1361-6560/aa9910.
- James, F. and M. Roos (1975). “Minit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations.” In: *Comput. Phys. Commun.* 10, pages 343–367. DOI: 10.1016/0010-4655(75)90039-9.
- Jha, D. et al. (October 2014). “Adaptive center determination for effective suppression of ring artifacts in tomography images.” In: *Applied Physics Letters* 105.14, page 143107. ISSN: 00036951. DOI: 10.1063/1.4897441.
- Jone, R. C. (1959). “Quantum Efficiency of Detectors for Visible and Infrared Radiation.” In: edited by L. Marton and C. Marton. Volume 11. *Advances in Electronics and Electron Physics*. Academic Press, pages 87–183. DOI: 10.1016/S0065-2539(08)60994-4.

- Kaczmarz, S. (1937). “Angenäherte Auflösung von Systemen linearer Gleichungen.” In: *Bulletin International de l’Académie Polonaise des Sciences et des Lettres. Classe des Sciences Mathématiques et Naturelles. Série A, Sciences Mathématiques* 5, pages 355–357.
- Kang, E., J. Min, and J. C. Ye (October 2017). “A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction.” In: *Medical Physics* 44.10, e360–e375. ISSN: 00942405. DOI: [10.1002/mp.12344](https://doi.org/10.1002/mp.12344).
- Kazantsev, D., G. Van Eynhoven, et al. (May 2015). “Employing temporal self-similarity across the entire time domain in computed tomography reconstruction.” In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373.2043, pages 20140389–20140389. ISSN: 1364-503X. DOI: [10.1098/rsta.2014.0389](https://doi.org/10.1098/rsta.2014.0389).
- Kazantsev, D., E. Guo, et al. (March 2016). “Temporal sparsity exploiting nonlocal regularization for 4D computed tomography reconstruction.” In: *Journal of X-Ray Science and Technology* 24.2, pages 207–219. ISSN: 08953996. DOI: [10.3233/XST-160546](https://doi.org/10.3233/XST-160546).
- Kazantsev, D., W. M. Thompson, et al. (2015). “4D-CT reconstruction with unified spatial-temporal patch-based regularization.” In: *Inverse Problems & Imaging* 9.2, pages 447–467. ISSN: 1930-8345. DOI: [10.3934/ipi.2015.9.447](https://doi.org/10.3934/ipi.2015.9.447).
- Kingston, A., A. Sakellariou, A. Sheppard, et al. (August 2010). “An auto-focus method for generating sharp 3D tomographic images.” In: *Developments in X-Ray Tomography VII*. Edited by S. R. Stock. Volume 7804. September 2010, 7804J. ISBN: 9780819483003. DOI: [10.1117/12.860285](https://doi.org/10.1117/12.860285).
- Kingston, A., A. Sakellariou, T. Varslot, et al. (August 2011). “Reliable automatic alignment of tomographic projection data by passive auto-focus.” In: *Medical Physics* 38.9, pages 4934–4945. ISSN: 00942405. DOI: [10.1118/1.3609096](https://doi.org/10.1118/1.3609096).
- Kump, K. et al. (December 1998). “Digital X-ray detector technology.” In: *RBM - Revue Européenne de Technologie Biomedicale* 20.9, pages 221–226. ISSN: 02220776. DOI: [10.1016/S0222-0776\(99\)80006-6](https://doi.org/10.1016/S0222-0776(99)80006-6).
- Leo, W. R. (1994). “Scintillation Detectors.” In: *Techniques for Nuclear and Particle Physics Experiments: A How-to Approach*. Berlin, Heidelberg: Springer Berlin Heidelberg, pages 157–175. ISBN: 978-3-642-57920-2. DOI: [10.1007/978-3-642-57920-2_7](https://doi.org/10.1007/978-3-642-57920-2_7).
- Lesaint, J. et al. (November 2017). “Calibration for Circular Cone-Beam CT Based on Consistency Conditions.” In: *IEEE Transactions on Radiation and Plasma Medical Sciences* 1.6, pages 517–526. ISSN: 2469-7311. DOI: [10.1109/trpms.2017.2734844](https://doi.org/10.1109/trpms.2017.2734844).
- Li, G. et al. (December 2018). “A novel calibration method incorporating nonlinear optimization and ball-bearing markers for cone-beam CT with a parameterized trajectory.” In: *Medical Physics* 46.1, mp.13278. ISSN: 0094-2405. DOI: [10.1002/mp.13278](https://doi.org/10.1002/mp.13278).
- Li, J. et al. (May 2017). “Scale formation and control in oil and gas fields: A review.” In: *Journal of Dispersion Science and Technology* 38.5, pages 661–670. ISSN: 15322351. DOI: [10.1080/01932691.2016.1185953](https://doi.org/10.1080/01932691.2016.1185953).

- Lin, Q. et al. (May 2018). "Optimization of image quality and acquisition time for lab-based X-ray microtomography using an iterative reconstruction algorithm." In: *Advances in Water Resources* 115, pages 112–124. ISSN: 03091708. DOI: [10.1016/j.advwatres.2018.03.007](https://doi.org/10.1016/j.advwatres.2018.03.007).
- Løge, I. A. et al. (2022). "Crystal formation from high resolution 4D X-ray μ CT: Initiation, expansion, and densification." In preparation.
- Maire, E. and P. J. Withers (January 2014). "Quantitative X-ray tomography." In: *International Materials Reviews* 59.1, pages 1–43. ISSN: 0950-6608. DOI: [10.1179/1743280413Y.0000000023](https://doi.org/10.1179/1743280413Y.0000000023).
- Manrique, E. et al. (April 2010). "EOR: Current status and opportunities." In: *Proceedings - SPE Symposium on Improved Oil Recovery*. Volume 2. SPE, pages 1584–1604. ISBN: 9781617384158. DOI: [10.2118/130113-ms](https://doi.org/10.2118/130113-ms).
- Marone, F. and M. Stampanoni (November 2012). "Regridding reconstruction algorithm for real-time tomographic imaging." In: *Journal of Synchrotron Radiation* 19.6, pages 1029–1037. ISSN: 09090495. DOI: [10.1107/S0909049512032864](https://doi.org/10.1107/S0909049512032864).
- Masson-Delmotte, V. et al. (2021). "Summary for Policymakers." In: *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge: Cambridge University Press. Chapter 1.
- Meng, Y., H. Gong, and X. Yang (February 2013). "Online geometric calibration of cone-beam computed tomography for arbitrary imaging objects." In: *IEEE Transactions on Medical Imaging* 32.2, pages 278–288. ISSN: 02780062. DOI: [10.1109/TMI.2012.2224360](https://doi.org/10.1109/TMI.2012.2224360).
- Menke, H., M. Andrew, et al. (2016). "Reservoir condition imaging of reactive transport in heterogeneous carbonates using fast synchrotron tomography - Effect of initial pore structure and flow conditions." In: *Chemical Geology* 428, pages 15–26. ISSN: 00092541. DOI: [10.1016/j.chemgeo.2016.02.030](https://doi.org/10.1016/j.chemgeo.2016.02.030).
- Menke, H., B. Bijeljic, M. Andrew, et al. (April 2015). "Dynamic Three-Dimensional Pore-Scale Imaging of Reaction in a Carbonate at Reservoir Conditions." In: *Environmental Science & Technology* 49.7, pages 4407–4414. ISSN: 0013-936X. DOI: [10.1021/es505789f](https://doi.org/10.1021/es505789f).
- Menke, H., B. Bijeljic, and M. Blunt (2017). "Dynamic reservoir-condition microtomography of reactive transport in complex carbonates: Effect of initial pore structure and initial brine pH." In: *Geochimica et Cosmochimica Acta* 204, pages 267–285. ISSN: 00167037. DOI: [10.1016/j.gca.2017.01.053](https://doi.org/10.1016/j.gca.2017.01.053).
- Menke, H., C. Reynolds, et al. (March 2018). "4D multi-scale imaging of reactive flow in carbonates: Assessing the impact of heterogeneity on dissolution regimes using streamlines at multiple length scales." In: *Chemical Geology* 481, pages 27–37. ISSN: 00092541. DOI: [10.1016/j.chemgeo.2018.01.016](https://doi.org/10.1016/j.chemgeo.2018.01.016).
- Mohammadkhani, S., H. Shahverdi, and M. N. Esfahany (August 2018). "Impact of salinity and connate water on low salinity water injection in secondary and tertiary stages for enhanced oil recovery in carbonate oil reservoirs." In: *Journal of Geophysics and Engineering* 15.4, pages 1242–1254. ISSN: 1742-2132. DOI: [10.1088/1742-2140/aaae84](https://doi.org/10.1088/1742-2140/aaae84).

- Muders, J. and J. Hesser (February 2014). “Stable and Robust Geometric Self-Calibration for Cone-Beam CT Using Mutual Information.” In: *IEEE Transactions on Nuclear Science* 61.1, pages 202–217. ISSN: 0018-9499. DOI: [10.1109/TNS.2013.2293969](https://doi.org/10.1109/TNS.2013.2293969).
- Mundhenk, N. et al. (May 2013). “Corrosion and scaling as interrelated phenomena in an operating geothermal power plant.” In: *Corrosion Science* 70, pages 17–28. ISSN: 0010938X. DOI: [10.1016/j.corsci.2013.01.003](https://doi.org/10.1016/j.corsci.2013.01.003).
- Myers, G., M. Geleta, et al. (2015). “Bayesian approach to time-resolved tomography.” In: *Optics Express* 23.15, page 20062. ISSN: 1094-4087. DOI: [10.1364/OE.23.020062](https://doi.org/10.1364/OE.23.020062).
- Myers, G., A. Kingston, et al. (2011). “Dynamic tomography with a priori information.” In: *Applied Optics* 50.20, page 3685. ISSN: 0003-6935. DOI: [10.1364/ao.50.003685](https://doi.org/10.1364/ao.50.003685).
- Nielsen, S. M., I. Nesterov, and A. A. Shapiro (June 2016). “Microbial enhanced oil recovery—a modeling study of the potential of spore-forming bacteria.” In: *Computational Geosciences* 20.3, pages 567–580. ISSN: 15731499. DOI: [10.1007/s10596-015-9526-3](https://doi.org/10.1007/s10596-015-9526-3).
- Nikitin, V. V. et al. (September 2019). “Four-Dimensional Tomographic Reconstruction by Time Domain Decomposition.” In: *IEEE Transactions on Computational Imaging* 5.3, pages 409–419. ISSN: 2333-9403. DOI: [10.1109/TCI.2019.2898088](https://doi.org/10.1109/TCI.2019.2898088).
- Noo, F. et al. (November 2000). “Analytic method based on identification of ellipse parameters for scanner calibration in cone-beam tomography.” In: *Physics in Medicine and Biology* 45.11, pages 3489–3508. ISSN: 0031-9155. DOI: [10.1088/0031-9155/45/11/327](https://doi.org/10.1088/0031-9155/45/11/327).
- Nordsøfonden (January 2021). *Oil and gas in Denmark*. Edited by Nordsøfonden. Accessed: 25.01.2022. URL: <https://eng.nordsoefonden.dk/about/oil-and-gas-in-denmark/>.
- Okuta, R. et al. (2017). “CuPy: A NumPy-Compatible Library for NVIDIA GPU Calculations.” In: *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*.
- Oliveira, D. F. et al. (September 2019). “Characterization of scale deposition in oil pipelines through X-Ray Microfluorescence and X-Ray microtomography.” In: *Applied Radiation and Isotopes* 151, pages 247–255. ISSN: 18729800. DOI: [10.1016/j.apradiso.2019.06.019](https://doi.org/10.1016/j.apradiso.2019.06.019).
- Øren, P. et al. (August 2019). “In-situ pore-scale imaging and image-based modelling of capillary trapping for geological storage of CO₂.” In: *International Journal of Greenhouse Gas Control* 87.May, pages 34–43. ISSN: 17505836. DOI: [10.1016/j.ijggc.2019.04.017](https://doi.org/10.1016/j.ijggc.2019.04.017).
- Ouadah, S. et al. (April 2016). “Self-calibration of cone-beam CT geometry using 3D–2D image registration.” In: *Physics in Medicine and Biology* 61.7, pages 2613–2632. ISSN: 0031-9155. DOI: [10.1088/0031-9155/61/7/2613](https://doi.org/10.1088/0031-9155/61/7/2613).
- Palenstijn, W. J., K. J. Batenburg, and J. Sijbers (2011). “Performance improvements for iterative electron tomography reconstruction using graphics processing units

- (GPUs)." In: *Journal of Structural Biology* 176.2, pages 250–253. ISSN: 1047-8477. DOI: [10.1016/j.jsb.2011.07.017](https://doi.org/10.1016/j.jsb.2011.07.017).
- Panetta, D., N. Belcari, A. Del Guerra, A. Bartolomei, et al. (2012). "Analysis of image sharpness reproducibility on a novel engineered micro-CT scanner with variable geometry and embedded recalibration software." In: *Physica Medica* 28.2, pages 166–173. ISSN: 11201797. DOI: [10.1016/j.ejmp.2011.03.006](https://doi.org/10.1016/j.ejmp.2011.03.006).
- Panetta, D., N. Belcari, A. Del Guerra, and S. Moehrs (July 2008). "An optimization-based method for geometrical calibration in cone-beam CT without dedicated phantoms." In: *Physics in Medicine and Biology* 53.14, pages 3841–3861. ISSN: 0031-9155. DOI: [10.1088/0031-9155/53/14/009](https://doi.org/10.1088/0031-9155/53/14/009).
- Panetta, D. (2016). "Advances in X-ray detectors for clinical and preclinical Computed Tomography." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 809. Advances in detectors and applications for medicine, pages 2–12. ISSN: 0168-9002. DOI: [10.1016/j.nima.2015.10.034](https://doi.org/10.1016/j.nima.2015.10.034).
- Patel, V. et al. (December 2008). "Self-calibration of a cone-beam micro-CT system." In: *Medical Physics* 36.1, pages 48–58. ISSN: 00942405. DOI: [10.1118/1.3026615](https://doi.org/10.1118/1.3026615).
- Pierrehumbert, R. (2015). "Rayleigh Scattering." In: *Encyclopedia of Astrobiology*. Berlin, Heidelberg: Springer Berlin Heidelberg, pages 2154–2155. DOI: [10.1007/978-3-662-44185-5_1349](https://doi.org/10.1007/978-3-662-44185-5_1349).
- Radon, J. (December 1986). "on the Determination of Functions From Their Integral Values Along Certain Manifolds." In: *IEEE Transactions on Medical Imaging* MI-5.4, pages 170–176. ISSN: 02780062. DOI: [10.1109/tmi.1986.4307775](https://doi.org/10.1109/tmi.1986.4307775).
- Rasmussen, P. W. et al. (December 2021). "Improved dynamic imaging of multi-phase flow by constrained tomographic reconstruction." In: *Scientific Reports* 11.1, page 12501. ISSN: 20452322. DOI: [10.1038/s41598-021-91776-1](https://doi.org/10.1038/s41598-021-91776-1).
- Renard, F., B. Cordonnier, D. Dysthe, et al. (July 2016). "A deformation rig for synchrotron microtomography studies of geomaterials under conditions down to 10 km depth in the Earth." In: *Journal of Synchrotron Radiation* 23.4, pages 1030–1034. ISSN: 1600-5775. DOI: [10.1107/S1600577516008730](https://doi.org/10.1107/S1600577516008730).
- Renard, F., B. Cordonnier, M. Kobchenko, et al. (October 2017). "Microscale characterization of rupture nucleation unravels precursors to faulting in rocks." In: *Earth and Planetary Science Letters* 476, pages 69–78. ISSN: 0012821X. DOI: [10.1016/j.epsl.2017.08.002](https://doi.org/10.1016/j.epsl.2017.08.002).
- Renard, F., J. McBeck, et al. (August 2019). "Volumetric and shear processes in crystalline rock approaching faulting." In: *Proceedings of the National Academy of Sciences* 116.33, pages 16234–16239. ISSN: 0027-8424. DOI: [10.1073/pnas.1902994116](https://doi.org/10.1073/pnas.1902994116).
- Renard, F., J. Weiss, et al. (February 2018). "Critical Evolution of Damage Toward System-Size Failure in Crystalline Rock." In: *Journal of Geophysical Research: Solid Earth* 123.2, pages 1969–1986. ISSN: 21699313. DOI: [10.1002/2017JB014964](https://doi.org/10.1002/2017JB014964).
- Richmond, C. (2004). "Sir Godfrey Hounsfield." In: *BMJ* 329.7467, page 687. ISSN: 0959-8138. DOI: [10.1136/bmj.329.7467.687](https://doi.org/10.1136/bmj.329.7467.687).

- Rodríguez-Sánchez, Á. et al. (November 2020). “Review of the influence of noise in X-ray computed tomography measurement uncertainty.” In: *Precision Engineering* 66, pages 382–391. ISSN: 01416359. DOI: [10.1016/j.precisioneng.2020.08.004](https://doi.org/10.1016/j.precisioneng.2020.08.004).
- Röntgen, W. C. (1895). *The bones of a hand with a ring on one finger, viewed through x-ray. Photoprint from radiograph by W.K. Röntgen.* Wellcome Collection. URL: <https://wellcomecollection.org/works/wjc8ejn2>.
- (January 1896). “On a new kind of rays.” In: *Nature* 53.1369, pages 274–276. ISSN: 00280836. DOI: [10.1038/053274b0](https://doi.org/10.1038/053274b0).
- Sassi, K. M. and I. M. Mujtaba (July 2011). “Optimal design and operation of reverse osmosis desalination process with membrane fouling.” In: *Chemical Engineering Journal* 171.2, pages 582–593. ISSN: 13858947. DOI: [10.1016/j.cej.2011.04.034](https://doi.org/10.1016/j.cej.2011.04.034).
- Scanziani, A. et al. (2020). “Dynamics of enhanced gas trapping applied to CO₂ storage in the presence of oil using synchrotron X-ray micro tomography.” In: *Applied Energy* 259. July 2019, page 114136. ISSN: 03062619. DOI: [10.1016/j.apenergy.2019.114136](https://doi.org/10.1016/j.apenergy.2019.114136).
- Seyyedi, M. and M. Sohrabi (February 2017). “Pore-Scale Investigation of Crude Oil/CO₂ Compositional Effects on Oil Recovery by Carbonated Water Injection.” In: *Industrial & Engineering Chemistry Research* 56.6, pages 1671–1681. ISSN: 0888-5885. DOI: [10.1021/acs.iecr.6b04743](https://doi.org/10.1021/acs.iecr.6b04743).
- Seyyedi, M., M. Sohrabi, et al. (January 2018). “Quantification of oil recovery efficiency, CO₂ storage potential, and fluid-rock interactions by CWI in heterogeneous sandstone oil reservoirs.” In: *Journal of Molecular Liquids* 249, pages 779–788. ISSN: 01677322. DOI: [10.1016/j.molliq.2017.10.070](https://doi.org/10.1016/j.molliq.2017.10.070).
- Seyyedi, M., S. Tagliaferri, et al. (November 2018). “An integrated experimental approach to quantify the oil recovery potential of seawater and low-salinity seawater injection in North Sea chalk oil reservoirs.” In: *Fuel* 232, pages 267–278. ISSN: 00162361. DOI: [10.1016/j.fuel.2018.05.158](https://doi.org/10.1016/j.fuel.2018.05.158).
- Sijbers, J. and A. Postnov (July 2004). “Reduction of ring artefacts in high resolution micro-CT reconstructions.” In: *Physics in Medicine and Biology* 49.14, N247–N253. ISSN: 00319155. DOI: [10.1088/0031-9155/49/14/N06](https://doi.org/10.1088/0031-9155/49/14/N06).
- Singh, K., H. Menke, et al. (December 2017). “Dynamics of snap-off and pore-filling events during two-phase fluid flow in permeable media.” In: *Scientific Reports* 7.1, page 5192. DOI: [10.1038/s41598-017-05204-4](https://doi.org/10.1038/s41598-017-05204-4).
- Singh, K., H. Scholl, et al. (2017). “The Role of Local Instabilities in Fluid Invasion into Permeable Media.” In: *Scientific Reports* 7.1, pages 1–11. ISSN: 20452322. DOI: [10.1038/s41598-017-00191-y](https://doi.org/10.1038/s41598-017-00191-y).
- Six, N., J. De Beenhouwer, and J. Sijbers (November 2019). “poly-DART: A discrete algebraic reconstruction technique for polychromatic X-ray CT.” In: *Optics Express* 27.23, page 33670. ISSN: 1094-4087. DOI: [10.1364/oe.27.033670](https://doi.org/10.1364/oe.27.033670).
- Spurin, C. et al. (October 2019). “Intermittent fluid connectivity during two-phase flow in a heterogeneous carbonate rock.” In: *Physical Review E* 100.4, page 043103. ISSN: 2470-0045. DOI: [10.1103/PhysRevE.100.043103](https://doi.org/10.1103/PhysRevE.100.043103).

- Sternlof, K. R. et al. (July 2006). “Flow and transport effects of compaction bands in sandstone at scales relevant to aquifer and reservoir management.” In: *Water Resources Research* 42.7. ISSN: 00431397. DOI: [10.1029/2005WR004664](https://doi.org/10.1029/2005WR004664).
- Sun, Z. and J. C. Santamarina (August 2019). “Haines jumps: Pore scale mechanisms.” In: *Physical Review E* 100.2, page 023115. ISSN: 24700053. DOI: [10.1103/PhysRevE.100.023115](https://doi.org/10.1103/PhysRevE.100.023115).
- The Danish Hydrocarbon Research and Technology Centre (2022). *A Key Player in the Energy Transition - Annual Report 2019/2020*. URL: <https://www.oilgas.dtu.dk/english/Research/Publications/Annual-Reports>.
- The Danish Parliament (2021). *L 30 Forslag til lov om ændring af lov om anvendelse af Danmarks undergrund*. Edited by The Danish Parliament. URL: <https://www.ft.dk/samling/20211/lovforslag/L30/index.htm>.
- The HDF Group (1997). *Hierarchical Data Format, version 5*. URL: <https://www.hdfgroup.org/HDF5/>.
- TotalEnergies (2021). *TotalEnergies in Denmark*. Edited by TotalEnergies. Accessed: 25.01.2022. URL: <https://wiki.totalenergies.com/en/country/totalenergies-denmark>.
- Touir, R. et al. (December 2009). “Corrosion and scale processes and their inhibition in simulated cooling water systems by monosaccharides derivatives. Part I: EIS study.” In: *Desalination* 249.3, pages 922–928. ISSN: 00119164. DOI: [10.1016/j.desal.2009.06.068](https://doi.org/10.1016/j.desal.2009.06.068).
- Tsang, T. (1998). *Classical Electrodynamics*. UNITEXT for Physics. Cham: Springer International Publishing. ISBN: 978-3-319-91808-2. DOI: [10.1142/3388](https://doi.org/10.1142/3388).
- Tuy, H. K. (June 1983). “Inversion Formula for Cone-Beam Reconstruction.” In: *SIAM Journal on Applied Mathematics* 43, pages 546–552. ISSN: 00361399. DOI: [10.1137/0143035](https://doi.org/10.1137/0143035).
- Van Eyndhoven, G., K. J. Batenburg, D. Kazantsev, et al. (November 2015). “An Iterative CT Reconstruction Algorithm for Fast Fluid Flow Imaging.” In: *IEEE Transactions on Image Processing* 24.11, pages 4446–4458. ISSN: 1057-7149. DOI: [10.1109/TIP.2015.2466113](https://doi.org/10.1109/TIP.2015.2466113).
- Van Eyndhoven, G., K. J. Batenburg, and J. Sijbers (February 2014). “Region-based iterative reconstruction of structurally changing objects in CT.” In: *IEEE Transactions on Image Processing* 23.2, pages 909–919. ISSN: 10577149. DOI: [10.1109/TIP.2013.2297024](https://doi.org/10.1109/TIP.2013.2297024).
- Virtanen, P. et al. (2020). “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” In: *Nature Methods* 17, pages 261–272. DOI: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2).
- Vo, N. T., R. C. Atwood, and M. Drakopoulos (October 2018). “Superior techniques for eliminating ring artifacts in X-ray micro-tomography.” In: *Optics Express* 26.22, page 28396. ISSN: 1094-4087. DOI: [10.1364/oe.26.028396](https://doi.org/10.1364/oe.26.028396).
- Whiting, B. R. et al. (August 2006). “Properties of preprocessed sinogram data in x-ray computed tomography.” In: *Medical Physics* 33.9, pages 3290–3303. ISSN: 00942405. DOI: [10.1118/1.2230762](https://doi.org/10.1118/1.2230762).

- Wicklein, J. et al. (July 2012). “Image features for misalignment correction in medical flat-detector CT.” In: *Medical Physics* 39.8, pages 4918–4931. ISSN: 00942405. DOI: [10.1118/1.4736532](https://doi.org/10.1118/1.4736532).
- Wildenschild, D. and A. P. Sheppard (January 2013). “X-ray imaging and analysis techniques for quantifying pore-scale structure and processes in subsurface porous medium systems.” In: *Advances in Water Resources* 51, pages 217–246. ISSN: 03091708. DOI: [10.1016/j.advwatres.2012.07.018](https://doi.org/10.1016/j.advwatres.2012.07.018).
- Willeminck, M. et al. (November 2018). “Photon-counting CT: Technical principles and clinical prospects.” In: *Radiology* 289.2, pages 293–312. ISSN: 15271315. DOI: [10.1148/radiol.2018172656](https://doi.org/10.1148/radiol.2018172656).
- Würfl, T. et al. (2016). “Deep learning computed tomography.” In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Edited by S. Ourselin et al. Volume 9902 LNCS. Lecture Notes in Computer Science. Cham: Springer International Publishing, pages 432–440. ISBN: 9783319467252. DOI: [10.1007/978-3-319-46726-9_50](https://doi.org/10.1007/978-3-319-46726-9_50).
- Xiang, Q., J. Wang, and Y. Cai (August 2016). “A geometric calibration method for cone beam CT system.” In: *Eighth International Conference on Digital Image Processing (ICDIP 2016)*. Edited by C. M. Falco and X. Jiang. Volume 10033. August 2016, 100333G. ISBN: 9781510605039. DOI: [10.1117/12.2244579](https://doi.org/10.1117/12.2244579).
- Xu, Y. et al. (August 2017). “Simultaneous calibration phantom commission and geometry calibration in cone beam CT.” In: *Physics in Medicine & Biology* 62.17, N375–N390. ISSN: 1361-6560. DOI: [10.1088/1361-6560/aa77e5](https://doi.org/10.1088/1361-6560/aa77e5).
- Yang, K. et al. (May 2006). “A geometric calibration method for cone beam CT systems.” In: *Medical Physics* 33.6Part1, pages 1695–1706. ISSN: 00942405. DOI: [10.1118/1.2198187](https://doi.org/10.1118/1.2198187).
- Yang, X. et al. (March 2020). “Tomographic reconstruction with a generative adversarial network.” In: *Journal of Synchrotron Radiation* 27.2, pages 486–493. ISSN: 1600-5775. DOI: [10.1107/S1600577520000831](https://doi.org/10.1107/S1600577520000831).
- Yang, Y., S. Bruns, et al. (2018). “Retraction of the dissolution front in natural porous media.” In: *Sci. Rep.* 8, page 5693. DOI: [10.1038/s41598-018-23823-3](https://doi.org/10.1038/s41598-018-23823-3).
- Yang, Y., S. Hakim, et al. (June 2018). “Direct Observation of Coupled Geochemical and Geomechanical Impacts on Chalk Microstructure Evolution under Elevated CO₂ Pressure.” In: *ACS Earth and Space Chemistry* 2, pages 618–633. ISSN: 24723452. DOI: [10.1021/acsearthspacechem.8b00013](https://doi.org/10.1021/acsearthspacechem.8b00013).
- Yang, Y., Y. Li, et al. (April 2020). “Dynamic Pore-Scale Dissolution by CO₂-Saturated Brine in Carbonates: Impact of Homogeneous Versus Fractured Versus Vuggy Pore Structure.” In: *Water Resources Research* 56.4. ISSN: 19447973. DOI: [10.1029/2019WR026112](https://doi.org/10.1029/2019WR026112).
- Zhang, P., M. T. Tweheyo, and T. Austad (July 2007). “Wettability alteration and improved oil recovery by spontaneous imbibition of seawater into chalk: Impact of the potential determining ions Ca²⁺, Mg²⁺, and SO₄²⁻.” In: *Colloids and Surfaces A: Physicochemical and Engineering Aspects* 301.1-3, pages 199–208. ISSN: 09277757. DOI: [10.1016/j.colsurfa.2006.12.058](https://doi.org/10.1016/j.colsurfa.2006.12.058).

- Zhi, S. et al. (November 2021). “CycN-Net: A Convolutional Neural Network Specialized for 4D CBCT Images Refinement.” In: *IEEE Transactions on Medical Imaging* 40.11, pages 3054–3064. ISSN: 1558254X. DOI: [10.1109/TMI.2021.3081824](https://doi.org/10.1109/TMI.2021.3081824).
- Zhuge, X., W. J. Palenstijn, and K. J. Batenburg (January 2016). “TVR-DART: A more robust algorithm for discrete tomography from limited projection data with automated gray value estimation.” In: *IEEE Transactions on Image Processing* 25.1, pages 455–468. ISSN: 10577149. DOI: [10.1109/TIP.2015.2504869](https://doi.org/10.1109/TIP.2015.2504869).