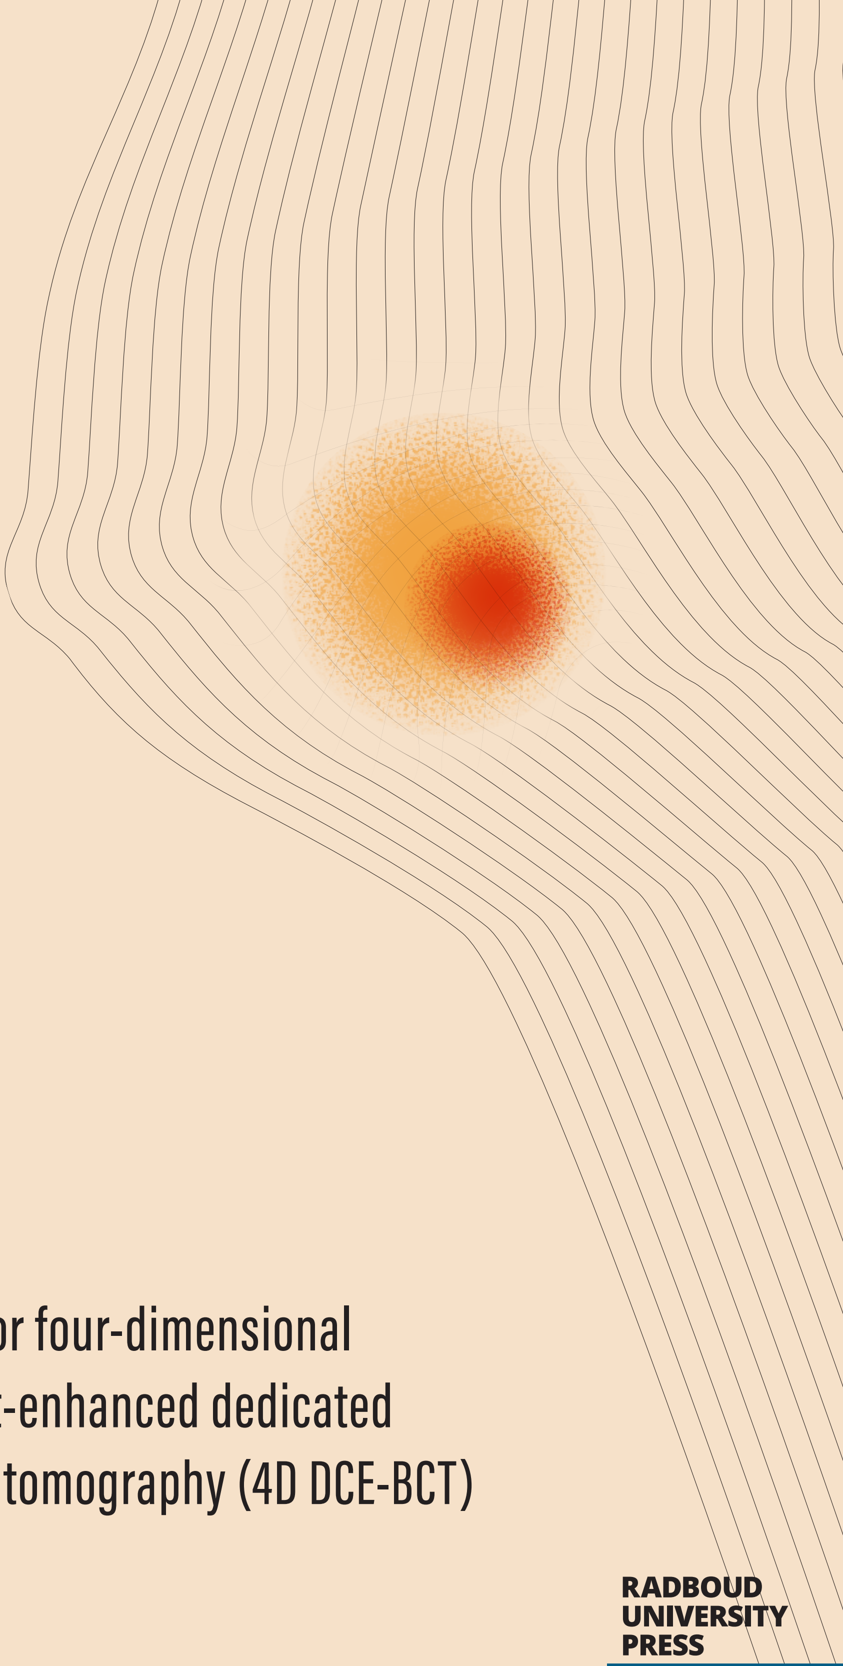

Reconstruction for four-dimensional dynamic contrast-enhanced dedicated breast computed tomography (4D DCE-BCT)

Mikhail Mikerov

RADBOUD
UNIVERSITY
PRESS

Radboud
Dissertation
Series

Reconstruction for four-dimensional dynamic contrast-enhanced dedicated breast computed tomography (4D DCE-BCT)

Reconstruction for four-dimensional dynamic contrast-enhanced dedicated breast computed tomography (4D DCE-BCT)
Mikhail Mikerov

Radboud Dissertation Series
ISSN: 2950-2772 (Online); 2950-2780 (Print)

Published by RADBOUD UNIVERSITY PRESS
Postbus 9100, 6500 HA Nijmegen, The Netherlands
www.radbouduniversitypress.nl

Design: Mikhail Mikerov
Cover: Proefschrift AIO | Guntra Laivacuma
Printing: DPN Rikken/Pumbo

ISBN: 9789465151731
DOI: 10.54195/9789465151731

RADBOUD
UNIVERSITY
PRESS

RECONSTRUCTION FOR FOUR-DIMENSIONAL DYNAMIC CONTRAST-ENHANCED DEDICATED BREAST COMPUTED TOMOGRAPHY (4D DCE-BCT)

Proefschrift ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen,
op gezag van de rector magnificus prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op

woensdag 25 februari 2026
om 12.30 uur precies

door

Mikhail Mikerov
geboren op 26 april 1993
te Perm, Rusland

Promotoren:

Prof. dr. I. Sechopoulos
Dr. R.M. Mann

Copromotor:

Dr. K.J.M. Michielsen

Manuscriptcomissie:

Prof. dr. ir. N.J.J Verdonschot
Prof. dr. G.J. Lord
Dr. C.R.L.P.N. Jeukens (Maastricht UMC+)

RECONSTRUCTION FOR FOUR-DIMENSIONAL DYNAMIC CONTRAST-ENHANCED DEDICATED BREAST COMPUTED TOMOGRAPHY (4D DCE-BCT)

Dissertation to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.M. Sanders,
according to the decision of the Doctorate Board
to be defended in public on

Wednesday, February 25, 2026
at 12.30 pm

by

Mikhail Mikerov
born on April 26, 1993
in Perm, Russia

PhD supervisors:

Prof. dr. I. Sechopoulos
Dr. R.M. Mann

PhD co-supervisor:

Dr. K.J.M. Michielsen

Manuscript Committee:

Prof. dr. ir. N.J.J. Verdonschot
Prof. dr. G.J. Lord
Dr. C.R.L.P.N. Jeukens (Maastricht UMC+)

CONTENTS

1

INTRODUCTION

1.1 BREAST CANCER TREATMENT

There is no doubt that breast cancer is the most diagnosed and one of the deadliest cancers among women. In 2020 alone, about 685 thousand women died because of it worldwide and more than 2.3 million new cases were diagnosed. Due to the world population becoming larger, older, and overall less healthy due to environmental factors, nutrition, and increasing exposure to stress, the number of new cases is predicted to increase to over 3 million per year and the number of deaths to exceed 1 million by 2040 [1].

Although still debated, the introduction of organized screening programs has led to a reduction of breast cancer-related deaths, because the cancers are detected earlier when the treatment can result in a better outcome [2]. For example, it is estimated that the breast cancer screening program in the UK prevented around 1300 deaths per year as of 2013 [3]. The introduction of breast cancer screening in the Netherlands accompanied with advances in treatment enabled the reduction of the breast cancer mortality among women older than 55 years by 42.3% in 2022 compared to 1989 [4].

If during screening the findings on the mammogram appear suspicious, the woman is recalled and will undergo a diagnostic procedure. In case the woman is diagnosed with breast cancer, different treatment options are available. The choice of treatment depends mostly on breast cancer type and stage as well as on availability of breast surgeons, and other factors. In some cases, women will first undergo neoadjuvant therapy (NAT) to reduce the tumor size. Neoadjuvant therapy can include chemotherapy, immune therapy, hormone therapy, radiotherapy, or any combination of the above. Sometimes, the neoadjuvant therapy can be very successful and lead to pathological complete response (pCR).

However, the achieved pCR rate is highly dependent on the molecular subtype of the tumor and treatment and can be anywhere between 8% and 60% [5, 6] or even exceed it for certain molecular subtypes. One of the possible reasons for the low success of neadjuvant therapy is the fact that this treatment is rarely personalized. Hence, the majority of women receive similar treatments and few adjustments are made before the treatment starts or during it, with no real verification to determine if the tumor is responding to it. The absence of personalization and adjustments is caused among other things by the lack of a suitable imaging modality that could be used to monitor treatment response. In the next section, the most commonly used imaging techniques in breast imaging treatment will be discussed.

1.2 IMAGING MODALITIES TO MONITOR BREAST CANCER TREATMENT

One of the most promising imaging modalities for breast cancer treatment monitoring is PET/CT. Seban *et al.* demonstrated that some PET biomarkers such as maximum standard uptake value and total metabolic tumor volume (SUVmax and TMTV) are independent factors for pCR prediction using multivariable logistic regression [7]. In early-stage triple-negative breast cancer cases that received neoadjuvant chemotherapy, the pCR rate was 85% for patients with high tumor SUVmax and low TMTV; 58% for patients with high tumor SUVmax and low TMTV; 6% for patients with low tumor SUVmax and high TMTV. Due to the major differences in the operational principle of PET and the developed imaging modality in this thesis, PET/CT will not be discussed further in detail in this work.

Reig *et al.* agree with the majority of the studies that dynamic contrast-enhanced MRI (DCE-MRI) is the most accurate modality to demonstrate response to therapy [8]. In their institution, DCE-MRI examination consists of a precontrast sequence followed by three axial T1-weighted postcontrast sequences. They have found that the changes in volume are more accurate in predicting response than other parameters that can be derived from images such as the longest diameter. Nevertheless, Chen *et al.* report that 6%-19% of cases are overestimated in MRI and 7%-28% are underestimated on MRI images [9]. Reig *et al.* list tumors with non-mass morphology or non-concentric shrinkage, and late enhancing foci as the main factors leading to underestimation of residual disease. On the other hand, fibrosis, i.e., post-treatment thickening of connective tissue, necrotic tumors that do not contain viable tumor cells but are still present, and residual benign masses are factors that can lead to overestimation of residual disease.

Fowler *et al.* agree that using the change of tumor as the main predictive quantity might lead to the lower accuracy of MRI or other modalities [10]. The authors indicate that the effect of therapy in terms of tumor size may not be immediate and some time is needed for the cells to die, which can then be easily detected by comparing the size of the tumor before and after therapy. However, since different areas of the tumor, which are either positively or negatively responding to the therapy, could perfuse blood differently, imaging of this difference could provide a valuable indication on the chances of positive treatment outcome. This kind of imaging is also known as functional imaging.

Functional imaging with DCE-MRI requires a higher number of images in the post-contrast sequence in order to calculate the perfusion parameters using both the images and mathematical modeling of blood flow in the breast. For example, the volume transfer constant K^{trans} between the extravascular extracellular space and plasma correlates with pathological response. Over the recent years more research was performed that confirmed the high predictive power of functional MRI imaging. Liang *et al.* could show in their study that a combination of parameters from DCE-MRI and apparent diffusion coefficient can achieve sensitivity and specificity above 80% [11]. Similar results were reported in a meta-analysis study by Chen *et al.* [9]. In a prospective study, Ramtohul *et al.* derived radiomics features from K^{trans}, volume fraction of extravascular and extracellular space, and maximum contrast agent uptake rate maps to achieve AUC of 0.8 for predicting pCR [12].

Although DCE-MRI provides functional information about the tumor, it has fundamental limitations that are very difficult to overcome. In general, temporal and spatial resolutions in DCE-MRI are coupled. Therefore, there is always a trade-off between the ability to depict smaller structures and sampling the perfusion curve adequately so as to not miss important phases. Moreover, motion artifacts become an important issue, which further decreases the accuracy of estimation of the perfusion curve. Since the acquisition takes longer in situations where an image with high spatial resolution is required, patient motion becomes more likely and has a greater influence on images with high spatial resolution. These limitations call for a different imaging modality, which can provide images with high spatio-temporal resolution.

1.3 Dynamic CT Perfusion imaging

Unlike MRI, X-ray imaging modalities can maximize both the spatial and temporal resolutions independently. The spatial resolution depends mainly on the geometric characteristics of the source and the number of imaging elements (pixels) on the detector. The flux output capability of the source, read-out speed of the detector, and the ability of the mechanical components to move the source and detector relative to the object being scanned determine the temporal resolution. Among all X-ray imaging modalities, computed tomography (CT) is the most commonly used when it comes to dynamic functional imaging (perfusion imaging).

One of the first ever published dynamic contrast-enhanced CT perfusion studies was presented already in 1979 [13], within a decade after the invention of the CT scanner by Godfrey Hounsfield. In this study the authors provided clear evidence that cerebral ischemia and infarction could be detected by quantifying the change in the attenuation coefficient as a function of time in the affected parts of the brain and could not be detected on static scans.

Since then, significant progress has been made in the field of dynamic CT perfusion imaging. The technical advances in CT technology have enabled imaging at much higher spatial and temporal resolutions. Improvements in both the hardware and software have led to decreased radiation dose to the patient during such exams. Finally, the parallel development of mathematical models of perfusion permitted calculation of quantitative maps, which are correlated with the clinical outcome of the studied pathologies. For example, mean transit time (MTT) describes the average time for blood to flow through vessels and increased values indicate the widening of blood vessels and decreased blood pressure. It is often used with cerebral blood flow (CBF) - another important quantitative map - to detect stroke. CBF is the amount of blood that flows per mass of tissue per minute. Low CBF identifies regions that suffered from infarct. One of the main requirements for the application of mathematical models that are used to calculate quantitative maps from perfusion sequences is stable performance in low-dose acquisition mode. Methods based on Bayesian modeling have been shown to be more robust in low-dose settings and are currently used in clinical practice [14].

Nowadays, dynamic CT perfusion imaging finds its application not only in identifying stroke, but also in diagnosing liver metastases, and cardiovascular diseases. The application of CT perfusion to monitor the effect of neoadjuvant therapy in the liver was demonstrated in the study led by Kim *et al.* [15]. Similar studies provided evidence that dynamic CT perfusion imaging at high spatio-temporal resolution is beneficial to assess treatment response compared to static scans [16]. CT perfusion was especially important when the treatment involved antiangiogenic therapies, i.e., therapies that have an indirect effect on the size of the tumor since the cytostatic effect is more prevalent than the cytotoxic effect. The treatment is called cytostatic when it affects the growth of tumor cells as opposed to cytotoxic that aims to kill tumor cells.

According to literature, to this day only a handful of CT perfusion imaging sequences of the breast have ever been accomplished. Hirasawa *et al.* could create perfusion CT images for 24 patients who were scanned in supine position using helical CT scanners limited to single-slice imaging [17]. The analysis of this data showed that it is possible to differentiate tumor types based on perfusion values. One of the few theoretical studies on functional CT imaging of the breast was performed on digital phantoms and was the foundation for this thesis. In that study, Caballo *et al.* demonstrated that dynamic contrast-enhanced imaging of the breast with a dedicated imaging modality, breast CT [18], would enable the differentiation of not only malignant from benign tumors, but also the assessment of treatment response on sub-tumor level. This dedicated CT scanner will be introduced in the next section.

1.4 BREAST CT

While full-body CT can be used to acquire images of the breast, it is not the optimal modality due to three main reasons. The first limitation is that the spatial resolution of full-body CT scanners is not high enough to accurately depict fine fibro-glandular structures and small cancers in the breast. The second reason for not using full-body CT for breast imaging is the relatively high dose administered to the patient since the X rays need to penetrate the whole chest without the image suffering from photon starvation artifacts. A related reason is the reduced soft tissue contrast that is caused by the higher energy x-ray beam that is necessary to achieve this. Finally, imaging in supine position will likely result in spatial misalignment of CT images at different time points because of patient motion due to breathing while imaging in prone position requires modifications of the patient table.

Dedicated breast CT is a relatively novel imaging modality that has been developed to acquire fully 3D images of the breast at high spatial resolution without exposing the chest region to direct radiation. This is achieved by placing the X-ray source and detector on a gantry that rotates in the horizontal plane around the breast in a pendant position. The typical positioning of a woman in the breast CT scanner is shown in figure 1.1. The woman lies prone on the table, with one breast pending through the opening in the table. The gantry with mounted source and detector in half-cone beam geometry rotates around the breast and acquires projection images that are then processed and used for reconstruction. The dimensions of the breast CT scanner vary. In the CT scanner used for the research in this thesis, the distance between the detector and source was 95 cm, while the center of rotation was located 60 cm away from the source. The detector was approximately 30 cm x 30 cm in size, enabling acquisition of breasts that are up to 19 cm long when they are placed exactly at the center of rotation.

One of the main differences between full-body and breast CT is that the latter uses lower-

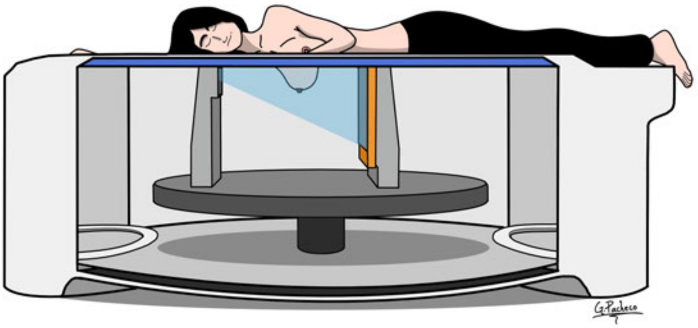

Figure 1.1: A woman being scanned with breast CT. The woman lies in a prone position with one breast pending through the opening in the patient table. The source and detector are positioned in half-cone beam geometry on a rotational platform.

energy X-ray photons. The maximum photon energy typically does not exceed 65 keV. It follows the same logic as in digital mammography and digital breast tomosynthesis (DBT) with the reasoning being that low-energy photons enable high contrast between adipose and fibro-glandular tissues. Unlike in the applications for which full-body CT is used, there are usually no highly attenuating structures in the breast. Thus, low-energy photons can still penetrate the breast and reach the detector. The absence of highly attenuating structures and the ability to perform CT scans without data truncation, i.e. the entire breast is projected onto the detector at all angles, allow for quantitative output in the form of linear attenuation coefficients at every location in the image. In theory, this quantitative output is only correlated to tissue properties and is independent of the breast's size, shape, and overall composition. However, physical effects such as scatter and beam hardening will reduce the quantitative accuracy if they are not appropriately accounted for.

The acquisition time of breast CT is affected not only by the readout times of the detector, but also by the size of the detector and mechanical components responsible for the rotation of the gantry. Some breast CT vendors incorporate small photon-counting detectors, requiring a spiral scan in order to perform the image acquisition of the whole breast. From the mechanical point of view, breast CT scanners are not as advanced as the full-body CT scanners produced by major vendors, thus, making CT acquisitions at comparable speeds is impossible as of today. Long acquisition times and the absence of the breast support or compression as is the case in digital mammography and DBT increase the likelihood of involuntary patient movement, which will result in motion artifacts, therefore reducing the diagnostic value and numerical accuracy of the reconstructed 3D images when not accounted for with a motion compensation method.

1.5 IMAGE RECONSTRUCTION

The acquired projections of the object can be used to reconstruct CT images with different image properties depending on the clinical task at hand by selecting an appropriate reconstruction method. The standard reconstruction method used by the vendor whose scanner was used in this project is filtered-backprojection in half-cone beam geometry, also known as the FDK algorithm [19], named after L.A Feldkamp, L.C Davis, and J.W. Kress who proposed it. This is a fully analytical method, which requires only a single backprojection operation of the previously prefiltered and weighted projections to obtain the distribution of linear attenuation coefficients at the mean energy of the X-ray spectrum in the breast. The filtering serves two purposes: removal of low frequency component introduced by oversampling near the center of the reconstructed volume and image smoothing or removal of high frequency components to reduce image noise. The main advantage of the FDK reconstruction is its computational speed. No other method can achieve equivalent image quality faster. Moreover, a high number of options is available to affect the image properties by changing the type of the frequency filter and its cut-off frequency. For example, a decision can be made to introduce more image smoothing to reduce image noise, but it will come at the cost of losing the ability to detect small structures. In an ideal situation, the inverse of the backprojection operation, the forward projection operation on the reconstructed volume, would yield the measured projections independent of the used filter.

Often, it is advantageous to be able to modify the image in the image domain directly, while still ensuring that the forward projection and backprojection are each others' inverse operations. This can be achieved by incorporating regularizations in the image domain into the reconstruction process; however, in order to be able to reconstruct the images in that way, the reconstruction must become iterative, i.e., the image is constantly modified by incorporating small updates until it converges fully.

One of the simplest examples is reconstruction of the images with additional total variation regularization to reduce high-frequency noise. However, this algorithm can sometimes lead to artificially-looking piece-wise constant output if performed with suboptimal settings. More advanced reconstruction methods take the statistical nature of the X-ray acquisition process into account and model the signal on the detector as samples drawn from the Poisson distribution. One such method, Maximum Likelihood in TRansmission (MLTR) [20], finds the most likely distribution of the linear attenuation coefficients in the imaged object to obtain the signal on the detector. Due to its specific modeling of the detector signal, MLTR is more likely to lead to better results than FDK with decreasing radiation dose. Figure 1.2 depicts how the same projection data can lead to a different appearance as a result of different algorithms being used to reconstruct the images.

In addition to modeling the signal on the detector, one could also model each part of the imaged object as a weighted combination of base materials or pseudo-materials. For this, both base materials and pseudo-materials are specified by energy-resolved linear attenuation curves, while the linear attenuation curves of pseudo-materials (μ_1, μ_2) can be mathematically described by using just a few parameters (a,b & c,d):

$$\mu(E) = \alpha\mu_1(a,b) + \beta\mu_2(c,d).$$

The introduction of such complex models is necessary to correctly account for the polychro-

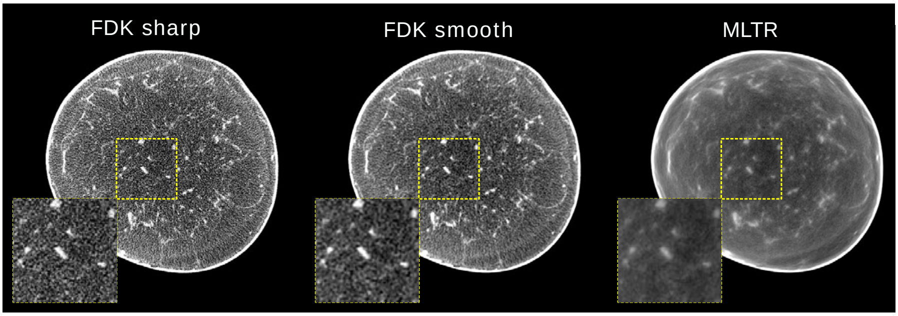

Figure 1.2: The same slice reconstructed with three different methods. Although all reconstructions are consistent with the projection data, their appearance is different.

matic nature of the X rays used in clinical CT. If the X rays are assumed to be monochromatic, it will lead to beam-hardening artifacts in the images. That is, the linear attenuation coefficient of the same tissues in the center of the image will always be lower than at the border.

A good combination of pseudo-materials to construct the linear attenuation coefficient in the low X-ray photon energy range typical for breast CT is the decomposition of the linear attenuation coefficient of the tissue into contributions due to two physical effects responsible for the attenuation of X rays in matter: Compton scatter and photoelectric absorption. In general, such models would require two independent sets of projections, e.g., acquired at different energies, to estimate the weighting coefficients α and β. However, in some situations, when it is possible to write the contributions due to those two physical effects as a mathematical function of a single parameter x, only one set of projections is required to reconstruct images without beam-hardening artifacts using a model-based method:

$$\mu(E) = \alpha(x)\,\mu_1(a,b) + \beta(x)\,\mu_2(c,d)\,.$$

The Iterative Maximum-likelihood Polychromatic Algorithm for CT (IMPACT) is an extension of MLTR method that achieves exactly that [21].

1.6 Automatic differentiation

Iterative model-based methods require the exact mathematical expression of the update step to be known in order to perform the optimization process. For example, MLTR and IMPACT maximize the posterior likelihood by constructing a monotonously increasing function and derive the update step this way. The major difficulty of this approach is that one is eventually limited in the complexity of the model that can be used. It is not always straightforward to derive a closed form of the update step. Moreover, each modification to the model requires a new derivation, notably limiting the number of changes that can be tested in a short time.

The goal of each optimization is the convergence of some cost metric as a function of input parameters of the model. In CT reconstruction, it is always possible to change one of the input parameters in the model and observe the influence of this change on the cost metric or, in other words, estimate the partial derivative. The easiest method to estimate

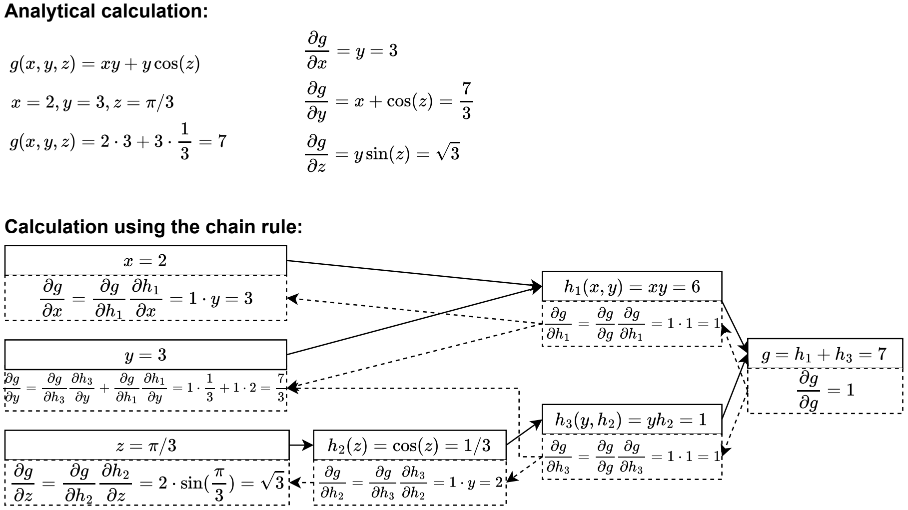

Figure 1.3: Symbolic and automatic differentiation lead to the same result. In the forward pass (solid lines) the value of g is calculated for given x, y, z. In the backward pass (dashed lines) the partial derivatives $\frac{\partial g}{\partial x}, \frac{\partial g}{\partial y}, \frac{\partial g}{\partial z}$ are calculated using the chain rule. At each link in the chain, the calculation of a partial derivative requires only the partial derivative from the previous step in the backward pass and analytical expression for the derivative in the current link.

the partial derivatives of any function is to introduce a very small change and observe how the cost metric reacts to it. This method is fully numerical and is called finite difference method of calculation of derivatives. If the partial derivatives of the cost metric in respect to all input parameters (the gradient) are known, gradient descent methods can be used to optimize the function to a local optimum.

Unfortunately, the finite difference method suffers from numerical inaccuracies and is very time consuming as each partial derivative needs to be calculated separately. A more efficient and accurate method is required to perform optimization with gradient descent of any differentiable function no matter how complicated it is.

Automatic differentiation is such an alternative. It is used to calculate analytical gradients of any composition of differentiable functions, since it can be shown that the composition of differentiable functions is differentiable. In order to perform automatic differentiation of a chained function, its computational graph needs to be constructed, which can then be reversed to propagate the gradients (see figure 1.3). In fact, efficient implementations of automatic differentiation accelerated the rise of deep learning and neural networks.

Strictly speaking, not any differentiable function can be minimized for the optimization result to be useful. Ideally, the function has a single point with zero gradient corresponding to the global minimum of this function. In practice, however, highly multidimensional functions usually contain local minima and saddle points where the gradient is also zero. In order to overcome local minima and saddle points, modern gradient descent methods incorporate momentum terms that allow to skip regions with very small gradients with inertia. It is similar to a massive tire rolling down a hill - it is very unlikely to stop in small potholes the faster it rolls.

Application of automatic differentiation in CT reconstruction opens many doors for very advanced regularization techniques and complicated models. A very promising direction is to approximate otherwise reference-based metrics with a neural network that can guide the reconstruction process. For example, Huang *et al.* taught a neural network to calculate Visual Image Fidelity, so that it can be used to reduce motion artifacts in an iterative motion compensation process when no ground truth is available [22]. Another application is the incorporation of the image domain resampling to reduce motion artifacts, as will be discussed later in this thesis.

1.7 Image reconstruction in 4D imaging

Methods such as FDK, MLTR, and IMPACT described earlier put serious requirements on the projection data. If these requirements are not fulfilled, the obtained reconstructions are not optimal and often cannot be used clinically. The primary requirement is that the projections must be densely sampled over an angular range sufficient to recover the image, i.e., covering at least 180° + the angle of the fan-beam. In other words, the angular spacing between the projections must be low. When this requirement is violated, characteristic streak artifacts will appear in the image. Although methods exist to minimize these artifacts, they all introduce changes to the images that are difficult to track down [9, 23–25]. Therefore, it is preferable to not use such techniques if highly accurate images are desired. The secondary requirement puts constraints on the noise in the projection domain. The lower the noise present in the projection domain, the smoother the images will appear in the image domain. Although there are regularization techniques to reduce the noise level in the image domain, they result in trade-off with bias that needs to be accounted for in applications requiring high numerical accuracy. In practice, those two requirements often lead to an increased radiation dose administered to the patient to ensure high quality of the images. Techniques such as tube current modulation can reduce the radiation dose by reducing photon fluence at certain angles. For example the photon fluence can be increased as function of the path length through the body that X rays need to propagate when acquiring projections from different angles.

Although it becomes increasingly more challenging to obtain single-phase images of comparable image quality with reduced dose, there is more freedom when dynamic imaging is performed. One of the biggest advancements in image processing in the last decade was the introduction of compressed sensing - a mathematical theory postulating that if some conditions are met, Nyquist's sampling theorem can be violated to obtain accurate signals from highly undersampled data. The most notable application of compressed sensing in CT was the development of the method for reconstruction of dynamic CT images from undersampled projection data sets, otherwise known as prior image constrained compressed sensing (PICCS) [26]. The base assumption of this method is that images in the post-contrast sequence differ from the pre-contrast image only at a few locations, such as in vessels or tumors, and that this change is not large. This is a valid assumption because only the perfused anatomies will change their attenuation value depending on the amount of blood, and at low concentrations of contrast agent this change is rather small. This assumption can be easily violated, though, if there is patient motion, since the change in the attenuation coefficient cannot be explained solely by uptake of contrast agent.

However, if the assumption of small changes only due to perfusion holds, just a small fraction of the pre-contrast image needs to be updated and it can be achieved using sparsely-sampled projection data. Moreover, the projection data can be acquired at lower dose compared to the pre-contrast image. As a result, both the sparse sampling and lower dose per projection lead to significant decreases in patient dose.

1.8 Noise reduction in dynamic CT

A probably even more straightforward approach to perform functional imaging at low dose compared to PICCS is to take the high noise level in the projection domain into account and use noise reduction techniques. Noise reduction techniques that are intrinsic to image reconstruction in addition to regularization in the image domain have been mentioned previously. On the one hand, frequency filters with lower cut-off frequencies can be used in FDK to obtain smoother images. On the other hand, methods like MLTR model the statistical nature of the X-ray attenuation and usually lead to images with lower noise levels.

Noise reduction methods can also be applied directly to reconstructed images. In general, the idea of such methods is to find similar voxels in the image and use them to calculate the average of these voxels. One of the most basic noise filters, the mean filter, calculates the mean value of all voxels in the spatial neighborhood of the voxel that is being filtered and replaces the voxel value with this new value. If the size of the neighborhood is large, it can lead to significant reduction of noise (see figure 1.4). Unfortunately, it also removes fine anatomical structures in the image, decreasing the spatial resolution. Such noise filters, which can be applied to a single image, always lead to worse spatial resolution.

Similar to the reconstruction of dynamic CT images, more options for noise reduction become available when working with sequential image data. Specifically, dynamic imaging removes the need to use voxels only from a local spatial neighborhood to calculate the average, since access to more information about each voxel other than its location is available. With dynamic imaging, the search for similar voxels can shift from the image domain into the time domain [27–31]. Therefore, the similarity of the voxels is determined based on

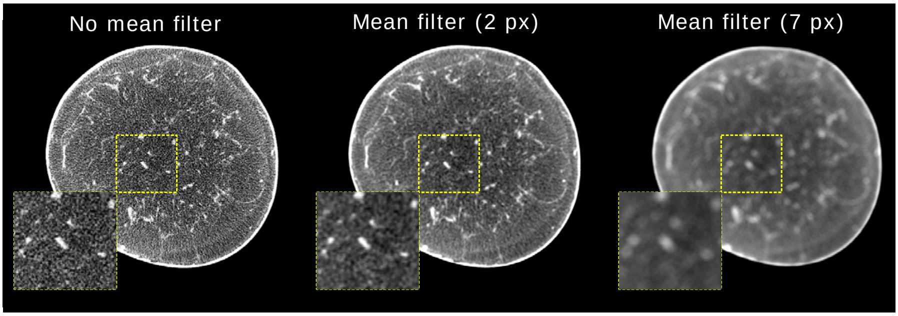

Figure 1.4: The same slice as in figure 1.2 reconstructed with a sharp FDK kernel before and after filtering with increasing strength. The noise reduction comes at the cost of worse spatial resolution.

the functional properties of corresponding anatomies, not based on their spatial proximity. For example, voxels corresponding to blood vessels exhibit completely different perfusion profiles than those of fatty tissue. A perfusion profile is the change of the linear attenuation coefficient in a specific voxel as a function of time. This has two consequences. First, more similar voxels can be found in the entire image than in a small local neighborhood based on their perfusion profiles. This leads to greater noise reduction since it is proportional to the square root of the number of similar voxels if they all are drawn from the same distribution. Second, there is no more averaging of voxels corresponding to different tissues if images are perfectly registered and aligned. Therefore, the spatial resolution is preserved. Again, the last requirement is often very easily violated in clinical practice when a patient moves during the scan sequence. In this case, patient movement needs to be counteracted by registering images at each phase to a common reference point, before such noise reduction methods are applied. One downside of these powerful methods that limits their application in the clinic is high computational cost. It is especially important when the patient needs to be diagnosed very quickly, for example when it is suspected to suffer from stroke. The reason for this is the working principle of these methods relying on brute force comparison of as many voxels as possible.

1.9 OUTLINE OF THIS THESIS

This thesis focuses on three aspects of image reconstruction needed for successful implementation of low-dose dynamic contrast-enhanced breast CT in the clinic. We present our approaches for accurate estimation of the iodine concentration in the breast at each time point in the post-contrast sequence, reduction of motion artifacts in single-phase breast CT acquisition, and advanced noise filtering when a series of independent very low-dose CT images is acquired after the administration of the contrast agent.

In **chapter 2**, we present our attempt to use the most accurate model of linear attenuation coefficients in model-based iterative reconstruction. We demonstrate in a simple simulation study that under strictly defined ideal conditions it could be beneficial to work with an undersampled system of equations to achieve numerically highly accurate reconstructions. Unfortunately, the results from this simplified theoretical study could not be directly transferred to real reconstruction problems. Hence, an alternative approach for the estimation of the iodine concentration in the breast needed to be developed. In **chapter 3**, we determine the optimal image acquisition and reconstruction parameters for PICCS in breast CT. Here, we focus on the numerical accuracy of the iodine quantification and absence of undersampling artifacts in the images. Next, we study the severity of motion artifacts that can be expected when performing very long scan sequences. In **chapter 4**, we examine DCE-MRI sequences for inter-frame motion and use this motion to develop a digital phantom to generate motion vector fields describing motion in DCE-MRI to guide us to a better understanding of the possible patient motion in breast CT. In **chapter 5**, we describe our non-rigid motion compensation algorithm for breast CT. We explain how we were able to reduce motion artifacts with a fully data-driven and model-free method. Finally, we verify its performance on clinical data from multiple institutions. In **chapter 6**, we accelerate a commercially available noise filter for dynamic contrast imaging by adapting it for execution on the GPU instead of the CPU and describe the changes to the algorithm that were needed to be made to achieve this. This last work lays the foundation for achieving very powerful noise reduction of 4D DCE-BCT as an alternative method to PICCS. In the **General Discussion**, we estimate the impact of this work on the application of breast CT for treatment planning and monitoring of treatment response in clinical practice.

REFERENCES

[1] M. Arnold, E. Morgan, H. Rumgay, A. Mafra, D. Singh, M. Laversanne, J. Vignat, J. R. Gralow, F. Cardoso, S. Siesling, and I. Soerjomataram. Current and future burden of breast cancer: Global statistics for 2020 and 2040. *The Breast*, 66:15–23, 2022.

[2] L. Tabar, M. Yen, B. Vitak, H. T. Chen, R. A. Smith, and S. W. Duffy. Mammography service screening and mortality in breast cancer patients: 20-year follow-up before and after introduction of screening. *The Lancet*, 361(9367):1405–1410, 2003.

[3] Independent UK Panel on Breast Cancer Screening. The benefits and harms of breast cancer screening: an independent review. *The Lancet*, 380(9855):1778–1786, 2012.

[4] Monitor dutch breast cancer screening programme 2022. Technical report, RIVM: Rijksinstituut voor Volksgezondheid en Milieu, 2023.

[5] N. Houssami, P. Macaskill, G. von Minckwitz, M. L. Marinovich, and E. Mamounas. Meta-analysis of the association of breast cancer subtype and pathologic complete response to neoadjuvant chemotherapy. *European Journal of Cancer*, 48(18):3342–3354, 2012.

[6] R. Wuerstlein and N. Harbeck. Neoadjuvant therapy for her2-positive breast cancer. *Reviews on Recent Clinical Trials*, 12(2):81–92, 2017.

[7] R. Seban, E. Arnaud, D. Loirat, L. Cabel, P. Cottu, L. Djerroudi, S. Hescot, P. Loap, C. Bonneau, F. Bidard, V. Huchet, N. Jehanno, A. Berenbaum, L. Champion, and I. Buvat. [18f]fdg pet/ct for predicting triple-negative breast cancer outcomes after neoadjuvant chemotherapy with or without pembrolizumab. *European Journal of Nuclear Medicine and Molecular Imaging*, 50(13):4024–4035, Nov 2023.

[8] B. Reig, A. A. Lewin, L. Du, L. Heacock, H. K. Toth, S. L. Heller, Y. Gao, and L. Moy. Breast mri for evaluation of response to neoadjuvant therapy. *RadioGraphics*, 41(3):665–679, 2021. PMID: 33939542.

[9] Y. Chen, L. Shi, Q. Feng, J. Yang, H. Shu, L. Luo, J. Coatrieux, and W. Chen. Artifact suppressed dictionary learning for low-dose ct image processing. *IEEE Transactions on Medical Imaging*, 33(12):2271–2292, 2014.

[10] A. M. Fowler, D. A. Mankoff, and B. N. Joe. Imaging neoadjuvant therapy response in breast cancer. *Radiology*, 285:358–375, 2017.

[11] X. Liang, X. Chen, Z. Yang, Y. Liao, M. Wang, Y. Li, W. Fan, Z. Dai, and Y. Zhang. Early prediction of pathological complete response to neoadjuvant chemotherapy combining dce-mri and apparent diffusion coefficient values in breast cancer. *BMC Cancer*, 22(1250), 2022.

[12] T. Ramtohul, V. Lepagney, C. Bonneau, M. Jin, E. Menet, J. Sauge, E. Laas, E. Romano, D. Bello-Roufai, F. Mechta-Grigoriou, A. Vincent Salomon, F. Bidard, A. Langer, C. Malhaire, L. Cabel, H.J. Brisse, and A. Tardivon. Use of pretreatment perfusion mri–based intratumoral heterogeneity to predict pathologic response of triple-negative breast cancer to neoadjuvant chemoimmunotherapy. *Radiology*, 312(3):e240575, 2024. PMID: 39225608.

[13] H. R. Heinz, P. Dubois, D. Osborne, B. Drayer, and W. Barett. Dynamic computed tomography study of the brain. *Journal of Computer Assisted Tomography*, 3(5):641–649, 1979.

[14] K. Mouridsen, K. Friston, N. Hjort, L. Gyldensted, L. Østergaard, and S. Kiebel. Bayesian estimation of cerebral perfusion using a physiological model of microvasculature. *NeuroImage*, 33(2):570–579, 2006.

[15] D. H. Kim, S. H. Kim, S. Im, S. Han, J. M. Goo, J. K. Willmann, E. S. Lee, J. S. Eo, J. C. Paeng, J. K. Han, and B. I. Choi. Intermodality comparison between 3d perfusion ct and 18f-fdg pet/ct imaging for predicting early tumor response in patients with liver metastasis after chemotherapy: Preliminary results of a prospective study. *European Journal of Radiology*, 81(11):3542–3550, 2012.

[16] S. H. Kim, A. Kamaya, and J. K. Willmann. Ct perfusion of the liver: Principles and applications in oncology. *Radiology*, 272(2):322–344, 2014. PMID: 25058132.

[17] H. Hirasawa, Y. Tsushima, S. Hirasawa, H. Takei, A. Taketomi-Takahasi, A. Takano, M. Amanuma, and K. Endo. Perfusion ct of breast carcinoma:: Arterial perfusion of nonscirrhous carcinoma was higher than that of scirrhous carcinoma. *Academic Radiology*, 14(5):547–552, 2007.

[18] M. Caballo, K. Michielsen, C. Fedon, and I. Sechopoulos. Towards 4d dedicated breast ct perfusion imaging of cancer: development and validation of computer simulated images. *Physics in Medicine and Biology*, 64:245004, 2019.

[19] L. A. Feldkamp, L. C. Davis, and J. W. Kress. Practical cone-beam algorithm. *J. Opt. Soc. Am. A*, 1(6):612–619, Jun 1984.

[20] K. Lang and R. Carson. Practical cone-beam algorithm. *J Comput Assist Tomogr.*, (8(2)):306–16, Apr 1984.

[21] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Transactions on Medical Imaging*, 20(10):999–1008, October 2001.

[22] H. Huang, J. H. Siewerdsen, W. Zbijewski, C. R. Weiss, M. Unberath, T. Ehtiati, and A. Sisniega. Reference-free learning-based similarity metric for motion compensation in cone-beam ct. *Physics in Medicine Biology*, 67(12):125020, jun 2022.

[23] G. Henrich. A simple computational method for reducing streak artifacts in ct images. *Computerized Tomography*, 4(1):67–71, 1980.

[24] E. Y. Sidky and X. Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine Biology*, 53(17):4777, aug 2008.

[25] D. Karimi and R. Ward. Reducing streak artifacts in computed tomography via sparse representation in coupled dictionaries. *Medical Physics*, 43(3):1473–1486, 2016.

[26] G. Chen, J. Tang, and S. Leng. Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets. *Medical physics*, 35(2):660–663, 2008.

[27] E. J. Smit and W. M. Prokop. Noise reduction in image data, WO 2017/194787 A1, 2017.

[28] A. M. Mendrik, E. Vonken, B. van Ginneken, H. W. de Jong, A. Riordan, T. van Seeters, E. J. Smit, M. A. Viergever, and M. Prokop. Tips bilateral noise reduction in 4d ct perfusion scans produces high-quality cerebral blood flow maps. *Physics in Medicine Biology*, 56(13):3857, jun 2011.

[29] F. Pisana, T. Henzler, S. Schönberg, E. Klotz, B. Schmidt, and M. Kachelrieß. Low dose CT perfusion using k-means clustering. In Despina Kontos and Thomas G. Flohr, editors, *Medical Imaging 2016: Physics of Medical Imaging*, volume 9783, page 97833M. International Society for Optics and Photonics, SPIE, 2016.

[30] F. Pisana, T. Henzler, S. Schönberg, E. Klotz, B. Schmidt, and M. Kachelrieß. Noise reduction and functional maps image quality improvement in dynamic ct perfusion using a new k-means clustering guided bilateral filter (kmgb). *Medical Physics*, 44(7):3464–3482, 2017.

[31] Z. Li, L. Yu, S. Leng, E. E. Williamson, A. L. Kotsenas, D. R. DeLone, A. Manduca, and C. H. McCollough. A robust noise reduction technique for time resolved ct. *Medical Physics*, 43(1):347–359, 2016.

2

ARE MORE ACCURATE MODELS ALWAYS USEFUL?

So I hit my head up against the wall
Over and over and over and over again, and again
'Cause I don't wanna be like them
I hit my head up against the wall
Over and over and over and over again
And again and again

Nothing But Thieves

Accurate models of the x-ray attenuation process are required for quantitative estimation of iodine concentration with model-based reconstruction methods. The choice of model is influenced not only by the accuracy sought but also by the increasing complexity when more free parameters need to be reconstructed. The applicability of three attenuation models was investigated in a single pixel problem using either two or three monochromatic beams near the K-edge energy of iodine. We found that an empirical model with five components, proposed by Midgley, leads to the lowest error when modeling iodine free materials and small error in estimating iodine concentration (0.1% and 3.39%), whereas the decomposition into contributions due to photoelectric effect and incoherent scatter results in more accurate estimation of the iodine concentration (0.72%) but has larger error (8.9%) when reconstructing iodine free materials. Decomposition into base materials shows the worst results on both objectives (8.9% and 62%).

This chapter appeared as
🗎 *M. Mikerov, K. Michielsen, J. G. Nagy, and I. Sechopoulos.*
Effect of attenuation model on iodine quantification in contrast-enhanced breast CT
7th International Conference on Image Formation in X-Ray Computed Tomography [1].

2.1 Introduction

Tumor characterization through quantitative functional imaging may allow for better treatment decisions in patients with breast cancer [2]. Dynamic contrast-enhanced breast CT is a new imaging modality being developed with aim to provide such functional information at good spatial and temporal resolutions. However, to maximize its clinical potential, accurate estimation of iodine concentration in the breast CT images is crucial. Current knowledge based on body CT imaging and computer simulations of contrast-enhanced breast CT indicate that the iodine concentration in the areas of interest, especially the tumor, can be expected to be in the range of 0.5 to 3.5 mg I per mL blood [3]. Coupled with sparse projections of typical breast CT systems and low photon energies that are required to increase the contrast (typically around 30 keV), the estimation of iodine concentration is a challenging task. In our implementation of contrast-enhanced breast CT, to save acquisition time and dose to the patient, individual projections are acquired only once with one of the x-ray spectra being used. This makes the use of decomposition methods in the projection domain not applicable.

Model-based methods are well suited to solve this reconstruction problem since they can use all available information about the acquisition, such as system geometry, utilized spectra, and physics models of attenuation processes. The latter determines, among other things, the number of free parameters that need to be estimated. Accurate modeling of the attenuation process is more difficult at low photon energies, where, in addition to photoelectric effect and incoherent scatter, coherent scatter also plays a role. Various parameterization schemes with different amounts of free parameters have been proposed to model attenuation.

In this work, we present the results of two experiments in which we examine the performance of three different parameterization schemes for energy dependent linear attenuation coefficients in a single pixel reconstruction problem. We focus on biological materials in the breast at x-ray energies below 49 keV and on accuracy of constrast quantification after adding iodine in the attenuation models, so we can determine which parameterization is most suitable to extend for our application, and include in our reconstruction method for quantitative breast CT imaging.

2.2 Methods

To avoid confounding influences, we focus on estimation of the energy dependent linear attenuation coefficient between 10 keV and 49 keV in a single pixel with monochromatic beams in dual and triple energy systems. Therefore, we are not solving the geometric aspect of the CT reconstruction problem, but are rather showing the adequacy of possible models of the energy dependency of linear attenuation.

2.2.1 Forward model

All values in the projection domain are obtained using the Beer-Lambert law:

$$p_E = I_0 \exp(-L \cdot \mu_E), \tag{2.1}$$

where μ_E is the linear attenuation coefficient at energy E, I_0 is the signal before attenuation, and L is the intersection length of the ray with the pixel of interest. The values of I_0 and L

were set to 1 for all experiments.

2.2.2 SOLUTION OF LINEAR SYSTEMS

Limiting the estimation of the linear attenuation coefficient to a single pixel allows us to solve the linear system $A\mathbf{x} = \mathbf{b}$ using two analytical methods. The first is non-negative least squares, which is applied when the linear system has either full rank or is overdetermined [4]. We solve underdetermined systems of equations using the conjugate gradient method [4], an iterative method that requires a good initial guess. It can be shown that for underdetemined systems of equations, conjugate gradient methods applied to $A^T A\mathbf{x} = A^T\mathbf{b}$, such as CGLS and LSQR, will converge to the minimum norm solution $A^T(AA^T)^{-1}\mathbf{b}$, making their application feasible [4].

2.2.3 MODELS OF ATTENUATION

We consider three different models for the energy dependent attenuation. The first model consists of a decomposition into base materials. In breast imaging, decomposition into adipose, fibro-glandular, and iodine components is the most evident choice. The system of equations for the single pixel problem in dual energy systems then takes the following form:

$$-\begin{pmatrix} \mu_a^L & \mu_g^L & \mu_i^L \\ \mu_a^H & \mu_g^H & \mu_i^H \\ -1 & -1 & -1 \end{pmatrix} \begin{pmatrix} f_a \\ f_g \\ f_i \end{pmatrix} = \begin{pmatrix} \ln p_L \\ \ln p_H \\ 1 \end{pmatrix} \tag{2.2}$$

The subscripts L and H refer to low and high energies. The last row puts a constraint on the otherwise underdetermined system of equations by enforcing conservation of volume, possibly causing undesirable behavior when the selected base materials are suboptimal to represent all expected tissues. When more than two spectra are used to acquire the images, additional rows can be added, making this system overdetermined.

The second model makes use of the physical processes underlying the attenuation of x rays. It decomposes the attenuation profile into contributions due to photoelectric effect and incoherent scatter. The energy dependency of the photoelectric effect is usually modeled using the power law

$$\Phi(E) = 1/E^3, \tag{2.3}$$

where E is the x-ray energy, whereas the Klein-Nishina equation is employed to describe incoherent (Compton) scatter

$$\begin{aligned} \Theta(E) = & \frac{(1+\alpha)}{\alpha^2}\left(\frac{2(1+\alpha)}{1+2\alpha} - \frac{\ln(1+2\alpha)}{\alpha}\right) \\ & + \frac{\ln(1+2\alpha)}{2\alpha} - \frac{1+3\alpha}{(1+2\alpha)^2}, \end{aligned} \tag{2.4}$$

where $\alpha = E/511$ keV. However, this parameterization cannot model K-edges. Thus, in order to use this model to estimate the concentration of contrast agents, the attenuation

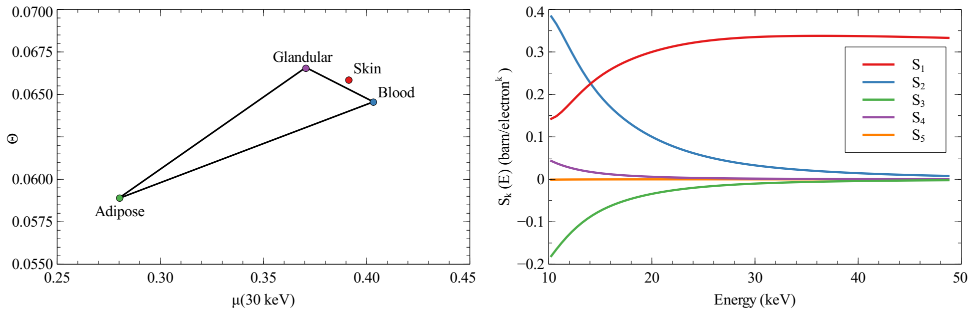

(a) (b)

Figure 2.1: (a) Fitted contribution due to incoherent scatter as function of linear attenuation coefficient at 30 keV for materials present in the breast. (b) The S-parameters were calculated using compositions of materials found in the breast.

characteristics of iodine must be included in the model. Consequently, the linear system of equations becomes:

$$-\begin{pmatrix} \phi^L & \Theta^L & \mu^L_{iodine} \\ \phi^M & \Theta^M & \mu^M_{iodine} \\ \phi^H & \Theta^H & \mu^H_{iodine} \end{pmatrix} \begin{pmatrix} f_\phi \\ f_\Theta \\ f_i \end{pmatrix} = \begin{pmatrix} \ln p_L \\ \ln p_M \\ \ln p_H \end{pmatrix}, \tag{2.5}$$

meaning that measurements at three different energies are required for the linear system to have full rank.

This model can be used in dual-energy setups following the method that was proposed by Depypere et al [5, 6]. If one can write ϕ and Θ as piece-wise linear functions of some parameter, e.g., the attenuation coefficient at a given energy, only two different measurements are needed. Hence, two parameters — the attenuation coefficient at a fixed energy and the iodine volume fraction — can be reconstructed using dual-energy setups. The difficulty of applying this method for breast imaging is that it is not possible to write the attenuation coefficient as a piece-wise linear function, as indicated in figure 2.1a. Blood, skin, adipose and glandular tissues have very similar properties. Since any possible combination of materials in the triangle adipose-glandular-blood is theoretically possible, a systematic error is introduced as soon as the background tissue does not lie on the chosen line.

The third parameterization we included in our experiment was proposed by Midgley [7]. It decomposes the linear attenuation coefficient into energy-dependent S-parameters that are weighted by composition-dependent a-parameters:

$$\mu_E = \sum_{k=1}^{5} a_k S_k. \tag{2.6}$$

The number of S-parameters is not fixed and more parameters will lead to more accurate models. However, five S-parameters are sufficient to keep the maximum error below 2% in the energy range up to 50 keV [7]. Moreover, following that $a_{k+1} \geq a_k$, μ_E monotonically decreases with energy. The S-parameters for biological tissues in breast imaging are shown in figure 2.1b. They were calculated by taking into account that a-parameters are functions of atomic number and electron density. Thus, the S-parameters are obtained by solving a least-squares problem [8], in our instance for elements that are found in breast tissues, specifically H, B, C, N, O, Na, Al, Si, P, S, Cl, Ar, K, Ca and Fe. As with the previous model, this parameterization scheme cannot account for K-edges. Thus, we extended this model to include the attenuation coefficient of iodine

$$\mu_E = \sum_{k=1}^{5} a_k S_k (1-f) + f \mu_{iodine}. \tag{2.7}$$

Assuming that the iodine volume fraction f is very small, the resulting system of equations in a dual-energy setup can be written as

$$-\begin{pmatrix} S_1^L & S_2^L & S_3^L & S_4^L & S_5^L & \mu_{iodine}^L \\ S_1^H & S_2^H & S_3^H & S_4^H & S_5^H & \mu_{iodine}^H \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \\ a_5 \\ f \end{pmatrix} = \begin{pmatrix} \ln p_L \\ \ln p_H \end{pmatrix}. \tag{2.8}$$

2.2.4 EXPERIMENTS

We examine the effect of the attenuation model on the accuracy of skin and iodine reconstruction in dual- and triple-energy acquisitions. We use skin as a good example of a background material that needs to be modeled correctly in order to avoid systematic errors that will otherwise propagate into the estimated iodine concentration in model-based reconstruction. Since model-based reconstruction tries to minimize the mismatch in the projection domain, incorrect linear attenuation coefficients in background material must be compensated by adjusting iodine concentration, which leads to lower accuracy.

The tissue background for the experiment including iodine contrast consists of a mixture of 50% adipose tissue, 50% fibro-glandular tissue; the volume fraction of blood is 30%. Varying amounts of iodine (0–20 mg/mL) were then added to the blood fraction. Our contrast agent is modeled as pure iodine. We assume that the attenuation of the contrast agent suspension medium is equal to that of blood. All materials were modeled with corresponding tissue substitutes [9–11]. and the attenuation profiles of elements were obtained from XrayDB [12].

The energies of the low and high energy spectra in the dual-energy setup were set to 30 keV and 34 keV, and the third energy for the triple-energy setup was set to 38 keV. Water was used as initialization for Midgley's parameterization. No initialization was needed for the other models since the resulting systems of equations were not underdetermined. The decomposition into photoelectric effect and incoherent scatter was considered only in the triple-energy setting.

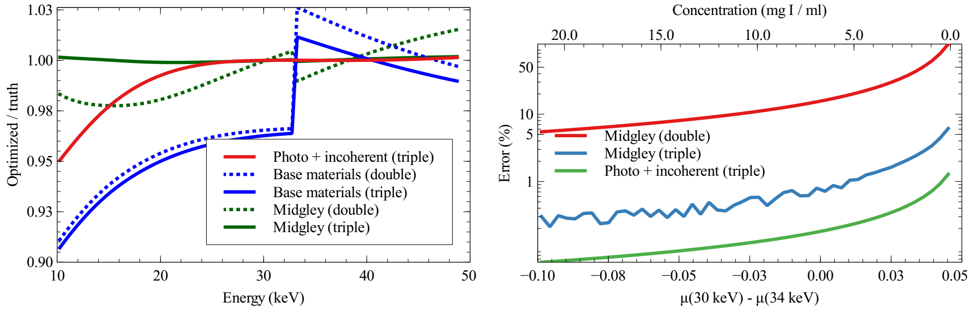

(a) (b)

Figure 2.2: (a) Proportional error between reconstructed and true linear attenuation coefficient of skin when using different models of attenuation and number of spectra. (b) Error in estimated iodine concentration as function of the difference between the linear attenuation of the mixture below and above the K-edge of iodine and concentration in blood. The negative region indicates rise in the attenuation due to K-edge. The background tissue was 50% adipose and 50% fibro-glandular tissues; volume fraction of blood was 30%.

2.3 Results

2.3.1 Reconstruction of skin

Figure 2.2a shows the relative error of the reconstructed attenuation coefficient of skin with three methods. The largest deviation (8.9% at 10 keV and $1.44{\cdot}10^{-4}$ volume fraction of iodine) results from the model that decomposes the background into adipose and glandular tissues. The most accurate result is achieved with Midgley's model in a triple-energy setup (0.1% at 10 keV and $-2.27{\cdot}10^{-6}$ volume fraction of iodine). Complete results are shown in table 2.1.

Table 2.1: Mean and maximum absolute relative errors on attenuation between 10 keV and 49 keV, and reconstructed iodine volume fraction for skin.

attenuation model	mean error	maximum error	iodine volume fraction
PE + Compton (double)	0.71%	4.95%	0.00
Base materials (double)	3.50%	8.88%	$1.44{\cdot}10^{-4}$
Base materials (triple)	3.50%	9.28%	$1.14{\cdot}10^{-4}$
Midgley (double)	1.13%	2.25%	$-3.80{\cdot}10^{-5}$
Midgley (triple)	0.07%	0.18%	$-2.27{\cdot}10^{-6}$

2.3.2 Estimation of iodine concentration

Figure 2.2b shows the error in iodine concentration estimation as a function of the difference between linear attenuation at 30 keV and 34 keV, resulting from the iodine K-edge when reconstructed using Midgley's model in dual- and triple-energy setups. The corresponding concentration of iodine is shown on top. The error is 62% and 3.39% at 0.97 mg I per mL of blood with double-energy and triple-energy setups, respectively. An even lower error (0.72%) is achieved with the decomposition into photoelectric effect and incoherent scatter with monochromatic radiation.

2.4 Discussion

The aim of this study was to find the most accurate parameterization scheme for model-based reconstructions methods for breast CT with a focus on quantitative accuracy of the estimation of iodine concentration. Unsurprisingly, the base material decomposition into iodine, adipose and glandular tissues is the least accurate method to reconstruct skin. The error is substantial as it clearly shows that the iodine concentration in skin will not be zero. The accuracy improves only slightly when using a triple-energy system.

Decomposition into contributions due to photoelectric effect and incoherent scatter on the other hand, is very accurate above 25 keV. However, its accuracy drops noticeably in the lower energy range. This can be explained by the contribution due to coherent scatter in this low-energy range, which is absent from this model.

Finally, Midgley's parameterization in the triple-energy setup leads to the most accurate results for skin. The parameterization with three energies achieves better results than the one with two. In particular, triple energy optimization has negligible iodine signal in skin even though the linear system remains underdetermined. This is most likely explained by our initialization which assigns a good approximation of the relative contributions of the scatter processes. Parameters not well constrained in the problem would then not end up far from a reasonable value. Further examination of how the initialization with water holds up for calcifications will give an indication on the limitations of using a simple homogeneous initialization in patient images.

Our experiment shows that the decomposition into contributions due to photoloelectric effect and incoherent scatter is equally good as Midgley's parameterization at energies above 25 keV. However, the result shown in figure 2.2a represents the ideal case. Introduction of real polychromatic spectra will complicate the situation. Besides, one of the spectra must have mean average energy below the K-edge of iodine to capture it, which leads to nonzero fluence in the low energy range. Taking into account that the attenuation coefficients are highest in this range, it is favorable to avoid systematic errors that are introduced due to less accurate models when using model-based reconstruction methods. Nevertheless, the mono-energetic simplification is the main limitation of this study that needs to be addressed.

The triple-energy system leads to better results in the relevant concentration range in our application, i.e., below 3.5 mg I per mL blood. Notably, it is not enough to just have beam energies on either side of the K-edge. Accurate quantification requires the K-edge that causes an increase in attenuation at relevant energies, otherwise, the accuracy quickly decreases when the K-edge is not causing such increase as can be observed in figure 2.2b. Since any allowed combinations of S-parameters are monotonically decreasing functions, the K-edge does not need to be modeled in order to connect the points on the linear attenuation profile at relevant energies. The continuous decrease in accuracy can be explained by different slopes of iodine before and after the K-edge, which means that there is still some information about iodine content present. The situation changes when the measurement at the third energy is added. Now, the only way to account for rise in attenuation and different slopes below and above the K-edge is to add the correct fraction of iodine.

2.5 CONCLUSION AND OUTLOOK

In this simplified numerical study we have shown that the attenuation model influences accuracy of iodine reconstruction in contrast-enhanced breast CT at energies below 49 keV. Such models could lead to underdetermined linear systems. Nevertheless, an accurate solution can be found if triple-energy systems and appropriate initialization are used. In our future research, we will examine if these conclusions remain valid for polychromatic spectra before incorporating the preferred attenuation model in our model-based reconstruction for breast CT.

ACKNOWLEDGMENT

This study was supported by ERC grant 864929.

REFERENCES

[1] M. Mikerov, K. Michielsen, J. G. Nagy, and I. Sechopoulos. Effect of attenuation model on iodine quantification in contrast-enhanced breast ct. In Joseph Webster Stayman, editor, *7th International Conference on Image Formation in X-Ray Computed Tomography*, Proceedings of SPIE - The International Society for Optical Engineering, United States, 2022. SPIE.

[2] J. Wu, G. Cao, X. Sun, J. Lee, D. L. Rubin, S. Napel, A. W. Kurian, B. L. Daniel, and R. Li. Intratumoral spatial heterogeneity at perfusion mr imaging predicts recurrence-free survival in locally advanced breast cancer treated with neoadjuvant chemotherapy. *Radiology*, 288(1):26–35, 2018. PMID: 29714680.

[3] M. Caballo, R. Mann, and I. Sechopoulos. Patient-based 4D digital breast phantom for perfusion contrast-enhanced breast CT imaging. *Medical Physics*, 45(10):4448–4460, October 2018.

[4] Å. Björck. *Numerical Methods for Least Squares Problems.* SIAM, 1996.

[5] M. Depypere, J. Nuyts, N. van Gastel, G. Carmeliet, F. Maes, and P. Suetens. An iterative dual energy CT reconstruction method for a K-edge contrast material. page 79610M, Lake Buena Vista, Florida, March 2011.

[6] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Transactions on Medical Imaging*, 20(10):999–1008, October 2001.

[7] S. M. Midgley. A parameterization scheme for the x-ray linear attenuation coefficient and energy absorption coefficient. *Physics in Medicine and Biology*, 49(2):307–325, January 2004.

[8] S. M. Midgley. A method for estimating radiation interaction coefficients for tissues from single energy CT. *Physics in Medicine and Biology*, 59(23):7479–7499, December 2014.

[9] ICRP, 2009. Adult Reference Computational Phantoms. ICRP Publication 110. Ann. ICRP 39 (2).

[10] D. R. White, J. Booz, R. V. Griffith, J. J. Spokas, and I. J. Wilson. Report 44. *Journal of the International Commission on Radiation Units and Measurements*, os23(1):NP–NP, 04 2016.

[11] J. W. Byng, J. G. Mainprize, and M. J. Yaffe. X-ray characterization of breast phantom materials. *Physics in Medicine and Biology*, 43(5):1367–1377, may 1998.

[12] Xraydb. `https://github.com/xraypy/XrayDB`.

3

USING HIGH-QUALITY DATA TO OBTAIN RELIABLE INFORMATION FROM LOWER-QUALITY DATA

Background: *Four-dimensional dynamic contrast-enhanced breast CT (4D DCE-bCT) offers promising high-resolution spatial and temporal imaging capabilities for the characterization and monitoring of breast tumors. However, the optimal combination of parameters for iodine quantification in image space remains to be determined.*

Purpose: *This study aims to optimize a dedicated bCT system to perform long dynamic contrast-enhanced scans with high spatio-temporal resolution while maintaining a reasonable radiation dose.*

Methods: *Our protocol includes the acquisition of a high-quality prior image that is reconstructed with a polychromatic iterative algorithm (IMPACT). The acquisition of the post-contrast sequence is continuous but sparse and these images are reconstructed using prior image constrained compressed sensing (PICCS). A four-step optimization process is performed using images of a physical phantom. First, the optimal tube current is selected by taking the noise level into account. Second, the optimal number of angles is selected based on the absence of streak artifacts. Third, the number of iterations in IMPACT is specified at the lowest value that achieves the highest spatial resolution. Finally, the number of iterations in PICCS is determined based on the quantitative accuracy of a range of iodine concentrations.*

Results: *When a high-quality prior image is available, the imaging of post-contrast images can be performed using just 40 projection angles with a tube current of 32 mA. The noise level in the post-contrast images is inherited from the prior image and no streak artifacts are visible. Mean difference between the linear attenuation coefficients of samples containing iodine reconstructed with IMPACT using all 360 projections and PICCS using 40 projections is 0.0004* mm^{-1} *at most. The spatial resolution of images reconstructed with PICCS is lower than* *the one of IMPACT images and is concentration dependent. The cut-off frequency at 10% MTF drops from 1.55 mm^{-1} in the prior image to 0.9 mm^{-1} when the target with the largest concentration is evaluated. The total mean glandular dose of the protocol does not exceed 22.5 mGy.*

Conclusions: *This study found the optimal acquisition and reconstruction parameters for a low-dose dynamic contrast-enhanced bCT protocol. The numerical accuracy of the proposed protocol was ensured by performing a physical phantom study.*

3.1 INTRODUCTION

Around 16% of all cancer deaths among women are due to breast cancer, resulting in an estimated 685 thousand deaths in 2020 worldwide. This problem is only growing, with breast cancer being predicted to take more than 1 million lives in 2040 [2]. Screening programs have already enabled early detection and diagnosis of breast cancer, thereby increasing positive treatment outcomes [3]. After diagnosis, some women first undergo neoadjuvant therapy to reduce tumor size, followed by breast conserving therapy. Cases where neoadjuvant therapy leads to pathologically complete response are correlated with both outcome and survival rates improvements [4–8]. In these cases it might even be possible to eliminate surgical treatment [9, 10].

Currently, neoadjuvant treatment is less efficient because it is rarely personalized due to limited availability of suitable imaging modalities. However, dynamic perfusion imaging may provide the necessary functional information to personalize treatment. This raises the importance of accurate quantification of the contrast agent concentration throughout the tumor over time, regardless of the breast shape, size, or composition.

Dynamic contrast–enhanced magnetic resonance imaging (DCE–MRI) and dedicated breast CT (bCT) are both suitable modalities for functional imaging. DCE–MRI with gadolinium contrast agents is already an established modality in breast imaging [11], while the feasibility of DCE imaging in bCT has only been shown on digital phantoms [12]. However, bCT offers several advantages over MRI, such as better uniformity, higher spatial and temporal resolution, and the ability to depict calcifications and to provide quantitative results. Iodine is a well suited contrast agent for bCT because of its K–edge energy of 33.2 keV, which enables the use of x–ray spectra with lower average x–ray energy to increase contrast between adipose and fibro–glandular tissues.

The most straightforward way to obtain the iodine uptake and wash-out curves is temporal subtraction imaging. The idea of this method is that when two images are subtracted from each other, the difference is solely explained by the change in the iodine concentration between the two acquisitions. Although one can use regular setups for anatomic imaging with already existing reconstruction methods, patient exposure becomes the limiting factor for this approach. Specifically, all images in the sequence must be

---

This chapter appeared as
🗎 *M. Mikerov, J.J. Pautasso, L. Goris, K. Michielsen, and I. Sechopoulos.*
4D Dynamic contrast-enhanced breast CT: Phantom-based reconstruction parameter optimization for iodine quantification
Medical Physics 2025 [1].

acquired at high dose with densely-sampled projections to minimize the artifacts and noise in the difference images. High dose limits the applicability of this method for sequences with high temporal resolution, as it negatively affects both the patient and can lead to overheating of the x-ray tube.

Alternatively, multi-energy acquisition protocols could be employed, where, depending on the setup, iodine concentration estimation can be achieved without patient motion artifacts. However, quantitatively accurate reconstruction of iodine distribution in multi-energy bCT is challenging due to several factors, such as sparse and spatially misaligned projections in half-cone beam geometry in the cone-angle direction. For example, the number of projections during one whole revolution of bCT typically ranges from 225 to 500 [13]. This number is further split between each spectrum if the multi–energy mode of operation is achieved by spectrum switching.

Thus, any reconstruction that requires material decomposition in the projection domain is not possible without interpolation over a large angular range of missing data [14]. On the other hand, methods based on decomposition in the image domain will need to handle any artifacts that are introduced due to the reconstruction of sparse data. Furthermore, methods that perform direct decomposition are limited to a maximum of two base materials in the most commonly used double–energy setups, with a third base material with K–edge possible given sufficient spectral separation around the K–edge energy. Although, this limitation is tolerated when three base materials are different, e.g., bone, soft tissue, and iodine, it is not sufficiently accurate for breast imaging, where it becomes important to distinguish between similar tissues—blood, adipose, and glandular—especially when they are combined with low concentrations of iodine.

In the future, multi-binned photon-counting detectors could enable the use of advanced models of attenuation to provide accurate quantitative results. However, to the best of our knowledge, photon-counting CT currently uses no more than two bins at the time of writing this work, and is limited to use in spiral bCT due to the high cost of these detectors.

Prior image constrained compressed sensing (PICCS) sets itself apart from the above-mentioned methods as a well suited technique for tracking of the contrast agent in bCT [15]. Briefly, this method uses a high quality prior image to then allow reconstruction of subsequent images from severely undersampled data of the same anatomy under the assumption that only minor changes are present that can be captured through compressed sensing. Although PICCS faces challenges with patient motion when applied in long dynamic scan sequences, it enables low-dose imaging with high temporal resolution and accurate quantification of the contrast agent using a single spectrum.

In this work we aim to determine the optimal data acquisition settings and reconstruction parameters for accurate iodine quantification in 4D dynamic contrast-enhanced bCT (4D DCE-bCT) using PICCS. We performed an extensive physical phantom study to quantify the influence of dose, number of projections, and number of iterations in iterative methods on the accuracy of iodine quantification.

3.2 Materials and Methods

3.2.1 DCE-bCT imaging protocol

A bCT system (Vera, Koning Corp, Norcross, GA, USA) was modified to perform multiple scan sequences, each up to 100 seconds in duration. The system is equipped with a pulsed x-ray tube and a flat-panel detector (FPD) mounted on a rotating gantry. The x-ray tube (M-1581, Varian Medical Systems, Salt Lake City, UT) can generate spectra with voltages between 49 kV and 65 kV, tube currents between 32 mA and 160 mA, and pulse durations between 5 ms and 8 ms in the current realization of the bCT. The FPD, a CMOS-based Xineos 3030HS (Teledyne Dalsa, Waterloo, ON, Canada), possesses a 0.152 mm pixel pitch and a 0.700 mm thick CsI:Tl columnar scintillator. The geometry of the system places the source 950 mm from the detector and 600 mm from the isocenter. The FPD can be shifted laterally to increase the field of view, but this option is not used in our setup since the detector is large enough for the vast majority of breasts encountered in our clinic and this allows higher temporal resolution since complete angular sampling is reached after a half rotation plus fan-angle instead of a full rotation.

The scan protocol we want to optimize is intended for patients with a known lesion to be characterized, follows the general idea of the protocol introduced by Prionas *et al.* [16]. The main difference is that instead of four post-contrast images in the original protocol, we aim to have continuous acquisitions over 300 seconds. The protocol starts with a pre-contrast scan of the contralateral breast. Afterward, the patient is repositioned and the pre-contrast scan of the ipsilateral breast is made. The long post-contrast imaging sequence of the ipsilateral breast begins immediately after contrast injection to ensure the wash-in is captured. This sequence consists of two or three 100 seconds long individual scans during which the projections are acquired continuously while the gantry rotates at 0.1 Hz, resulting in 10 full revolutions and 20 independently-reconstructed post-contrast images per 100-second scan, with half rotation overlap between each time-point. Finally, the patient is repositioned once more to perform a late-phase contrast-enhanced scan of the contralateral breast.

For the imaging of the contralateral breast we will use temporal subtraction with registration as developed by Gazi *et al* [17]. The reconstruction of the ipsilateral dynamic image sequence is described in section 3.2.2.

3.2.2 Reconstruction of 4D sequence

We reconstruct post-contrast images using the Prior Image Constrained Compressed Sensing (PICCS) method. This algorithm consists of two steps. To begin with, a high-quality artifact-free prior image needs to be reconstructed. This image is acquired before the administration of the contrast agent. It is always preferable to acquire this image at the highest possible exposure with densely sampled projections and without any patient motion, to maximize the chances that the prior image is not affected by artifacts, has good spatial resolution, and is ideally close to being piece-wise constant.

Following that, each post-contrast image in the sequence is reconstructed from severely undersampled projection data acquired at lower exposure levels using the high-quality prior to guide the reconstruction. Under the assumption that only a small fraction of voxels in the post-contrast image is different from the corresponding voxels in the high-quality prior image, their attenuation values can be estimated even from undersampled data by enforcing data sparsity.

Reconstruction of the prior image

We use the iterative maximum-likelihood polychromatic algorithm for CT (IMPACT) to obtain the prior image [18]. IMPACT is a single-energy polychromatic reconstruction method, which enables reconstruction of the image without cupping and metal artifact if the list of base materials is chosen appropriately for the imaged anatomy. IMPACT performs the tissue decomposition into a weighted sum of the energy-dependent contributions due to photoelectric effect $\Phi(e)$ given by Braggs-Pierce power law and Compton scatter $\Theta(e)$ given by the Klein-Nishina equation:

$$\mu_j(e) = \phi_j \cdot \Phi(e) + \theta_j \cdot \Theta(e). \tag{3.1}$$

At first glance, this model has two unknowns for every voxel j, ϕ_j and θ_j, which must be estimated, meaning that two different spectra are needed to solve the reconstruction problem. However, in practice both ϕ_j and θ_j are functions of a common parameter. In the work of De Man *et al.* this was chosen to be linear attenuation at a reference energy of 70 keV, which is appropriate for clinical CT while we use the linear attenuation at 30 keV, which is approximately the average energy in the breast CT x-ray spectrum. Fig 3.1 shows the representation of the contributions of both effects as a piece-wise linear function of linear attenuation at 30 keV for typical breast tissues. This data is also available in Table 3.1. Except for the choice of reference energy and base materials, our implementation matches the update step provided in the work of De Man *et al.* exactly. For convenience, this expression can be found in the appendix.

Table 3.1: Decomposition of the linear attenuation coefficient according to equation 3.1 for tissues present in breast. Iron is included to ensure numerical stability of the reconstruction.

	θ (1/cm)	ϕ (1/cm)	$\mu_{30\text{keV}}$ (1/cm)
air	1.8993e-04	8.1078e-05	0.0004
adipose tissue [19]	0.1651	0.0382	0.2907
fibro-glandular tissue [19]	0.1860	0.0630	0.3748
skin [20]	0.1863	0.0709	0.3959
calcification	1.1581	1.5463	5.0015
iron	16.8700	19.3133	64.3895

Reconstruction of post-contrast images

PICCS reconstruction can be written as the following minimization problem:

$$\underset{X}{\operatorname{argmin}}\, \alpha\|\nabla\left(X - X_p\right)\|_{l_1} + (1-\alpha)\|\nabla X\|_{l_1}, \text{ s.t. } \|\mathbf{A}X - Y\| \le \epsilon. \tag{3.2}$$

Here, X and X_p are the post-contrast and prior images, respectively; $\mathbf{A}$ is the forward operator, Y is the projection data. Discrete gradient operator ∇ is used as a sparsification

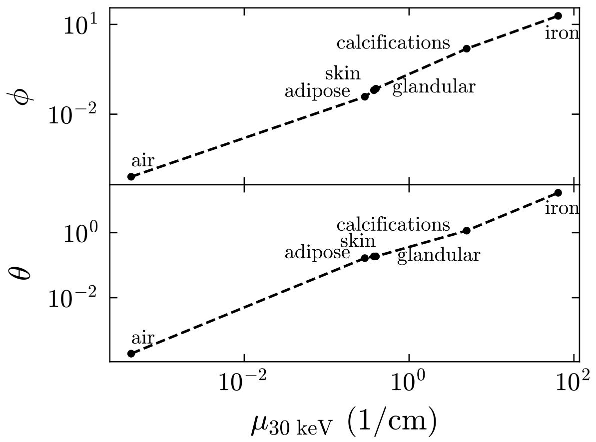

Figure 3.1: Contributions of photoelectric effect and Compton scatter to the total linear attenuation as functions of linear attenuation coefficient at 30 keV. Each combination of two materials is given by a linear equation.

transformation. Hyperparameter α can be seen as a control parameter for the smoothness of the image X.

The discrete gradient operator ∇ is chosen so that its application on the difference image sets all unchanged voxel values close to zero and keeps only a small fraction of large values. It is also applied to reduce the influence of a bad prior by enforcing image smoothness. We set $\alpha = 0.92$ to favor accuracy over image smoothness [21]. In other words, the post-contrast images can have a higher noise level and be affected by sparse reconstruction artifacts, but are expected to provide more accurate quantitative values.

After initialization with the prior image, the reconstruction process is performed iteratively, comprising two alternating main steps in each iteration. First, a single iteration of the IMPACT method is performed to keep the current image estimate consistent with the projection data. Second, a single step is taken to minimize the loss in eq. (3.2) using gradient descent with backtracking line search [21]. The pseudo-code of our implementation of the joint optimization procedure can be found in the appendix.

3.2.3 Breast phantom and iodinated contrast sample preparation

Iodine dilutions were prepared from a 300 mg/mL Iomeron (Bracco Imaging, Konstanz, Germany) stock solution, with concentrations ranging from 10 mg/mL to a 0 mg/mL control to mimic a scenario of a 100 kg patient with 5 liter blood volume. In such a case, a 150 mL injection of Iomeron, with a 300 mg/mL concentration, would provide an iodine blood concentration close to 9.0 mgI/mL. This led us to establish 10 mgI/mL as the upper limit for consistency. Each sample volume was adjusted to 50 mL, achieving dilutions from 0.5 to 10 mgI/mL. To ensure accurate and consistent sample volumes, water was initially added to reach 50 mL, then an equivalent volume was replaced with Iomeron stock, enabling precise dilutions and effective mixing. To ensure reproducible phantom positioning, a custom 3D-printed support was developed to hold the 50 mL tube within a conically shaped

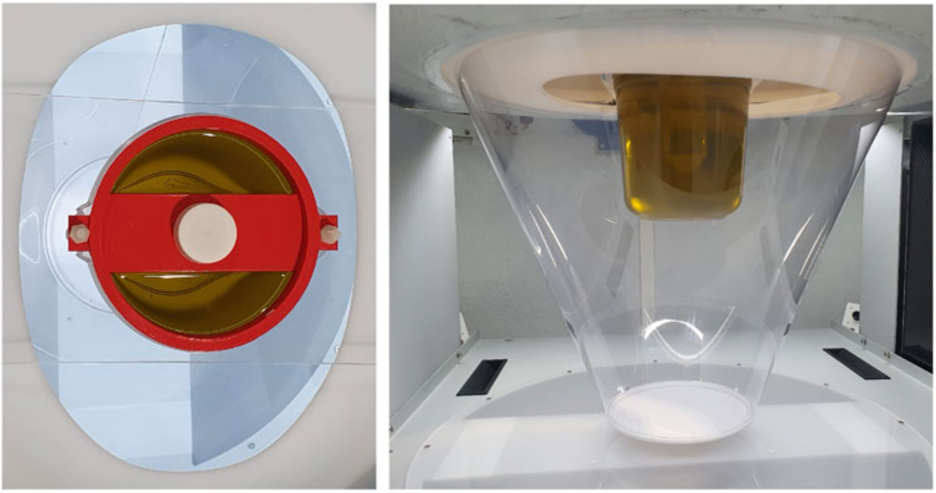

Figure 3.2: Phantom for imaging, including a 50 mL tube with iodine dilution, secured by a 3D-printed holder, all submerged in olive oil.

phantom, which measures 150 mm in height, transitioning from a 120 mm diameter at the top to an 110 mm diameter at the base. This phantom, filled with olive oil, serves to mimic the radiodensity characteristics of adipose tissue (Fig. 3.2). Due to the absence of motion, a homogeneous adipose-equivalent phantom is appropriate to study image reconstruction with PICCS since the background image is inherited from high-quality prior image [21, 22]. Thus, a phantom with a more complex background that includes features representing fibro-glandular tissue is not needed for this optimization process.

3.2.4 Finding the best acquisition and reconstruction parameters

The imaging protocol and reconstruction methods described in the previous sections have several parameters that need to be carefully selected to obtain good image quality and maximize quantitative accuracy. On the image acquisition side, the x-ray spectrum was selected to provide optimal iodine detection in projection domain in our previous work [23] while the tube current and the number of projections per rotation remain to be chosen.

On the image reconstruction side only the number of iterations to reconstruct the prior image and the number of iterations in PICCS need to be set, since in preliminary results it was confirmed, as expected, that tube current (above a minimum) had no impact on the reconstruction quality. These four parameters could be selected sequentially using the following scheme:

1. **Determine the tube current.** We reconstructed the phantom with the 6 mgI/mL target acquired with 60 projections over 360 degrees and at four different tube currents (32/50/64/80 mA) using PICCS. Mean values and standard deviations were calculated inside and outside the target in the ROIs displayed in Fig. 3.3 across

20 slices. This approach was based on the premise that demonstrating outcome consistency across settings was crucial, rather than the precision of individual values. Hence, we chose sparse acquisition with 60 projections due to the absence of streak artifacts and predetermined the lowest exposure setting that could be used if outcomes were consistently similar.

2. **Determine the optimal number of angles**. After the optimal exposure per projection was selected, the minimum required number of angles was determined using a method described in 3.2.6. The phantom with the highest concentration (10 mgI/mL) was used, since from our previous analysis we knew that PICCS performs worse when the difference between the images is large.
3. **Determine the number of iterations in IMPACT**. When reconstructing the prior image, a higher number of iterations will lead to better spatial resolution, but will also increase the noise level if the reconstruction is not regularized. However, noise regularization results in added bias and resolution loss. Because high quantitative accuracy and high resolution are important in this task, noise regularization was not used. The difference in the spatial resolution between reconstructions was given by the modulation transfer function (MTF) in the image domain, which was estimated using the method proposed by Richard *et al* [24]. The number of iterations that leads to the highest cutoff-frequency at 10% MTF while maintaining low noise level was chosen for the final optimization.
4. **Determine the number of iterations in PICCS.** The same analysis was performed to estimate the optimal number of PICCS iterations. However in this analysis, we focused on the convergence of the average value within the target, rather than on the noise level.

Had the sequential optimization not worked as described above, we would have performed a multidimensional sampling of the space of all parameters.

3.2.5 Scatter correction

Since the deep learning-based scatter correction method we developed is only valid for patient data [25], we implemented a dedicated correction for the phantom projection data and note here that this correction is not expected to perform differently from our approach for patient data. The concentration of iodine was the only variable modified; the shape and the location of the phantom remained unchanged. Following reconstruction of a single scan and segmentation of the phantom, primary and scatter projection images were produced using a validated Monte Carlo simulation algorithm [26]. These simulations emulated the bCT's geometry and settings. A 65 kV spectrum with a 0.257 mm Cu filter was modeled, and 10^9 photons were tracked. From the simulation, a set of 360 projections images (1° steps) was generated. Scatter correction was achieved by subtracting these MC-simulated scatter projections after normalization by the right magnitude from the phantom scans.

3.2.6 Quantification of the streak artifacts

We developed a quantitative method to estimate the optimal number of projections making use of the radial symmetry of the phantom. The method starts by aligning the geometrical

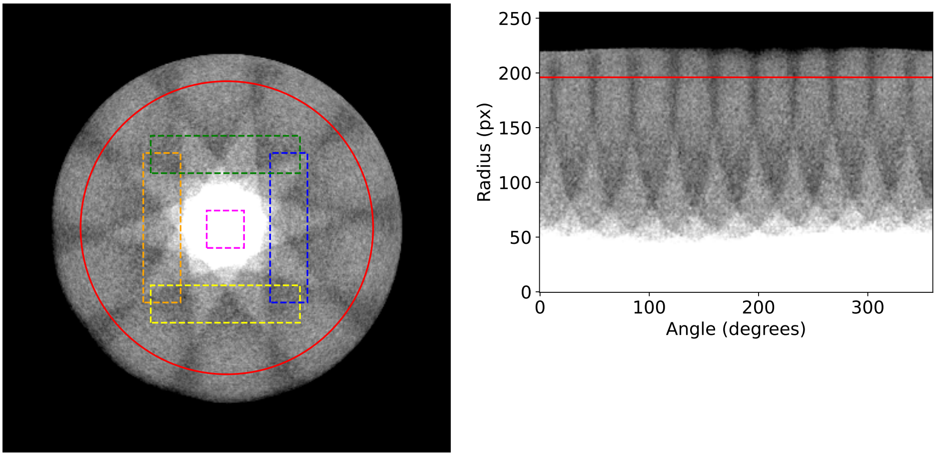

Figure 3.3: Representation of a single slice of the 10 mgI/mL phantom acquired with 10 angles and reconstructed with PICCS in Cartesian and polar coordinates. The data for frequency domain representation was collected along the red line. Ten radial lines attributed to the streak artifacts are clearly visible in both coordinate systems. Rectangular ROIs used to calculate the mean and standard deviation are overlaid on the Cartesian representation.

center of the phantom in the center of the 2D matrix. After that, coordinate transformation is applied to represent each 2D slice in polar coordinates system. Then, a line close to the outer border of the phantom is selected and averaged over 40 slices in the z-direction. Finally, we represent this averaged line in the frequency domain, in which each peak corresponds to the frequency of the streak artifacts. We select the angle of projections as a trade-off between the exposure and visibility of the artifacts. Fig. 3.3 illustrates the analysis of the phantom with the 10 mgI/mL target.

3.2.7 Air kerma measurement and dose calculation

To investigate the magnitudes of glandular doses resulting from the acquisition of an image sequence using a 4D DCE-bCT modality, the geometry of the bCT system described in Section 3.2.1 was replicated using a previously validated Monte Carlo-based algorithm [27]. This algorithm measured radiation captured by a pencil beam (100 mm × 10 mm) situated at the isocenter, with the simulation tracking 10^6 photons. The validation process involved comparing the simulated exposure distribution to empirical measurements obtained at the isocenter using a pencil 100 × 10 mm ion chamber and a Radcal Accu-Gold+ Touch Pro dosimeter.

Following validation, the simulation incorporated breasts of varying sizes, establishing chest wall diameters (CWD) at 10, 12, 14, 16, and 18 cm. The chest-to-nipple distances (CND) were set at 0.5, 0.75, and 1.0 times these diameters, yielding a total of 15 size variations [28].

Subsequently, the normalized mean glandular dose in mGy per mGy air kerma (DgN) from bCT was calculated, considering a glandular composition of 14.3% by volume [29]. To determine the mean glandular dose delivered to breast, the DgN coefficient is multiplied by the measured air kerma (mGy) at the isocenter for the defined 4D DCE-bCT image sequence.

3.3 Results

3.3.1 Optimal exposure level

Fig. 3.4a shows the mean values and standard deviations in five regions of the phantom: inside the sample with iodine and four regions just outside the iodine tube. As expected, the mean value inside the tube in the prior image is lower than in the post-contrast images acquired at various tube currents because it contains no iodine. It is evident that the tube current has no effect on neither the mean value nor standard deviation. Thus, we continue our further analysis only with the lowest tube current.

3.3.2 Determine optimal number of angles

The frequency-domain representation of the mean value over the angles shown in Fig. 3.4b indicates that 40 projections are sufficient to not cause major streaking artifacts. When the number of projections is decreased to 36, the peak at 36° becomes significantly larger than the peak at 40° leading to more visible streak artifacts as depicted in Fig. 3.5. Therefore, we continue our further analysis only with 40 projections.

3.3.3 Optimal number of iterations to reconstruct the prior image

Fig. 3.6a shows the radially-averaged edge profile of the tube filled with water in the image domain. We clearly see that the sharpness increases with higher number of iterations. Values near the center of the tube have larger spread due to the oversampling nature of the method. The standard deviation becomes larger with higher number of iterations as the noise increases. Therefore, only the range from 10 mm to 18 mm was used to calculate the MTF that is shown in Fig. 3.6b. The MTF plot supports the conclusion that better spatial resolution is achieved when the number of iterations is higher. Furthermore, in order to provide the decision based on a numerical value, we plotted the standard deviation outside the phantom and cutoff-frequency, at which MTF falls to 10%, as a function of the number of iterations. From Fig. 3.7 we conclude that 60 IMPACT iterations will be sufficient, as the benefits of having lower noise outweigh slightly better spatial resolution. Fig. 3.8 provides visual confirmation that higher number of iterations leads to better spatial resolution but also increases the noise level.

3.3.4 Optimal number of iterations in PICCS reconstruction

Fig. 3.9a visualizes the radially-averaged edge profiles of the tube filled with a 10 mg/mL iodinated solution in the image domain reconstructed with different number of PICCS iterations. It is apparent that a higher number of PICCS iterations is required for the mean value to converge. Moreover, a higher number of iterations also leads to a sharper depiction of the edge. However, the PICCS reconstruction is unable to reach the same spatial resolution as the IMPACT reconstruction of the prior image. Fig. 3.9b shows the image-domain MTF of the prior image and the phantom with the highest concentration

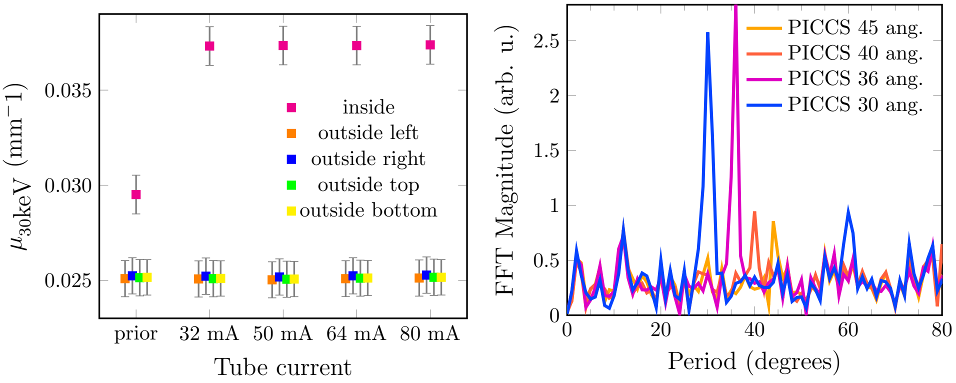

(a) (b)

Figure 3.4: (a) Mean attenuation values and standard deviations in five regions of the phantom with 6 mgI/mL insertion. (b) Frequency-domain representation of the streak artifacts due to undersampled data acquisition in the image-domain.

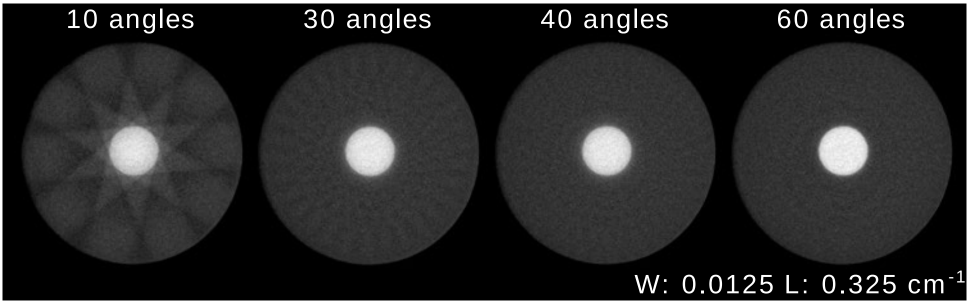

Figure 3.5: PICCS reconstruction of the phantom with 10 mgI/mL acquired with different number of angles.

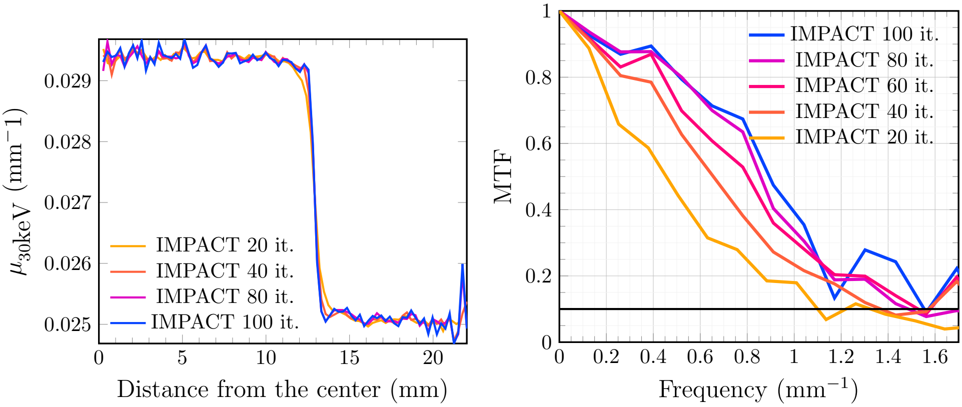

(a) (b)

Figure 3.6: (a) Average edge-profile of the tube with water in the image domain reconstructed for various number of iterations. (b) Image-domain MTF of the reconstruction of the prior image for various number of iterations.

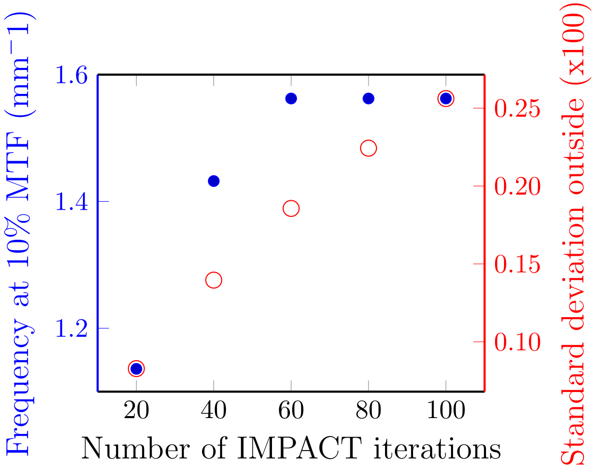

Figure 3.7: Spatial resolution and noise level as a function of the number of iterations when reconstructing the prior image.

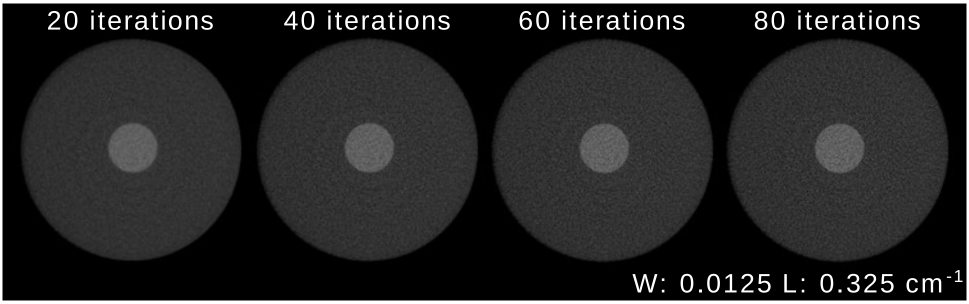

Figure 3.8: Reconstruction of the prior image with increasing number of iterations.

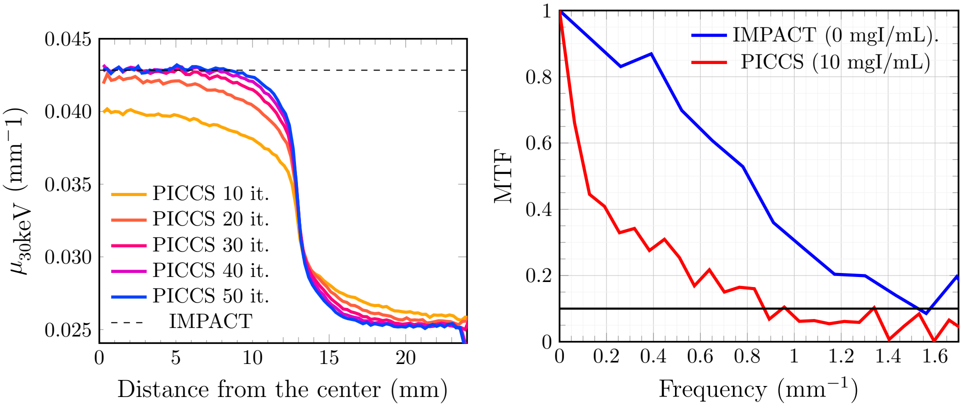

(a) (b)

Figure 3.9: (a) Average edge-profile of the tube with 10 mgI/mL in the image-domain reconstructed with various number of PICCS iterations. (b) Image-domain MTF of the prior image and image of the phantom with 10 mgI/mL target.

(10 mgI/mL) - the most challenging case for PICCS. We observe significant decrease in the ability to resolve small spatial features.

3.3.5 Calibration curves

Fig. 3.10a and 3.10b show that the maximum attenuation difference between values obtained with PICCS reconstruction with just 40 angles and IMPACT reconstruction with 360 angles does not exceed 0.0004 mm^{-1}. With the optimized settings, we estimated the calibration curves. The calibration curve using full reconstruction with IMPACT is

$$\mu_{30\text{keV}} = 0.001303 \cdot c + 0.029581, \tag{3.3}$$

While the same calibration curve obtained using PICCS reconstructions with 40 angles is

$$\mu_{30\text{keV}} = 0.001331 \cdot c + 0.029553, \tag{3.4}$$

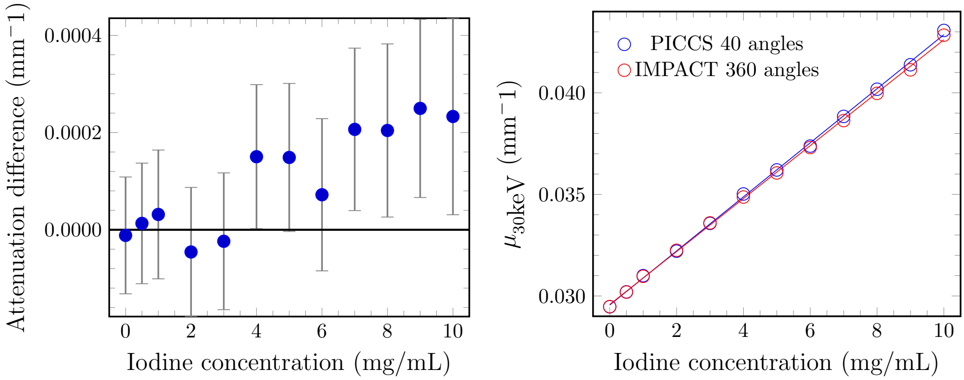

(a) (b)

Figure 3.10: (a) Relative difference between the linear attenuation coefficients reconstructed with IMPACT using all 360 projections and PICCS using 40 projections. (b) Calibration curves using IMPACT reconstructions of data sets with 360 projections and PICCS reconstructions of data sets with 40 projections.

where c is the concentration of iodine. The obtained parameters for the linear fit were compared using ANCOVA and both slopes (p=0.068) and intercepts (p=0.057) were found to not be different between both reconstructions.

3.3.6 OBTAINED AIR KERMA AND MEAN GLANDULAR DOSE

Air kerma measured at the isocenter showed that 360 projections at 65 kV and 80 mA with a 0.257 mm Cu filter resulted in 12.2 mGy, while 40 projections at 32 mA resulted in 0.56 mGy. In Table 3.2, the normalized mean glandular dose (DgN) coefficients for breasts of varying sizes are presented, with a mean value of 0.762 (ranging from a minimum of 0.667 to a maximum of 0.851).

Table 3.2: Normalized mean glandular dose (DgN) coefficients in mGy per mGy of air kerma for breasts of varying sizes, considering a glandularity of 14.3% by volume. Air kerma at the isocenter of the bCT system (65 kV/0.257 mm Cu) was taken as reference. "CWD" refers to chest wall diameters, and "CND" refers to chest-to-nipple distances.

CWD (cm)	CND (cm)	DgN (mGy/mGy)
10	5	0.763
10	7.5	0.818
10	10	0.851
12	6	0.744
12	9	0.797
12	12	0.827
14	7	0.712
14	10.5	0.772
14	14	0.801
16	8	0.686
16	12	0.746
16	16	0.775
18	9	0.667
18	13.5	0.720
18	18	0.745

3.4 DISCUSSION

The goal of this study was to perform the optimization of reconstruction parameters to define a protocol for DCE bCT. Since the aim is to continuously acquire post-contrast images over a minimum timespan of two minutes, it was important to reduce patient dose and prevent tube overheating. PICCS reconstruction was deemed suitable for this task as it allows image reconstruction from severely undersampled data acquired at low tube currents. However, the quantitative accuracy of the method needed to be ensured.

To begin with, no dependency of the mean value inside the tube and noise level on the tube current could be found. Image noise properties in the post-contast images reconstructed with PICCS are inherited from the prior image. Thus, the lowest tube current is selected for this protocol.

The minimum and maximum tube currents were determined by the bCT's settings and its 7 kW power capacity, respectively.

Following this, we recommend using 40 projection angles for post-contrast images, even though a lower number of angles can further reduce the dose while still maintaining sufficient image quality for diagnostic or localization purposes. Nevertheless, it will come at the cost of more prominent streak artifacts, which will negatively influence quantitative values.

The air kerma measured at the isocenter showed that 360 projections at 65 kV and 80 mA with a 0.257 mm Cu filter resulted in 12.20 mGy, while 40 projections at 32 mA resulted in 0.56 mGy. From Table 3.1, if an average breast size of 14 cm CWD and 10.5 cm CND is considered, then the mean glandular dose delivered to the breast is equivalent to 9.41

mGy for the pre-contrast phase and 0.43 mGy for the single image in the post-contrast case. During the entire post-contrast phase of 300 seconds, the total mean glandular dose delivered to the patient will be 12.96 mGy.

Next, we wanted to ensure that the degree of enhancement can be quantified, since making longitudinal comparisons throughout neoadjuvant treatment will be beneficial for personalizing and adapting treatment to response. Due to the large variability of breast sizes, shapes, and inner structures, FDK-based methods will suffer from cupping artifacts due to insufficient beam hardening correction. To enable longitudinal comparisons, the images were reconstructed with the iterative IMPACT method. While IMPACT reconstruction combined with scatter correction eliminates beam hardening artifacts, its spatial resolution and noise levels depend on the number of iterations. We determined that 60 iterations offer an optimal balance between spatial resolution and image noise.

Finally, we studied the impact of the number of PICCS iterations on the accuracy of the post-contrast images. In general, the higher the number of PICCS iterations, the more accurate and sharper the images will be. The analysis showed that the spatial resolution of the images reconstructed with PICCS will not reach the spatial resolution of the prior image even when the number of iterations is high. However, this analysis was performed with a very high concentration of 10 mgI/mL as the worst case scenario, while the iodine concentration in clinical applications is not expected to exceed 6 mgI/mL [30]. In general, the numerical accuracy of images reconstructed with PICCS will decrease with increasing concentration of the contrast agent, due to the difference constraint with the prior image and with smaller size of the perfused anatomy [31], the latter in our opinion likely due to partial volume effects.

Future work will be dedicated, among other things, to incorporating noise regularization methods in IMPACT and to improving its convergence rate. Additionally, the lower spatial resolution in the post-contrast images compared to prior image, which we identify as the main limitation of the proposed method, must be addressed. Moreover, the influence of intra-scan motion during the acquisition of the prior image and inter-scan motion between different frames needs to be studied as it could not be replicated with the presented phantom. Our approach for reducing inter-scan motion is rigid registration of the post-contrast projections to the corresponding projections of the prior image. This approach should significantly reduce the registration artifacts; in the best-case scenario, only non-rigid motion will remain. The artifacts due to intra-scan motion during the acquisition of the prior image can be reduced using our previously developed non-rigid motion compensation method [32]. An important point to consider is that even though it is able to produce artifact-free images when motion is moderate, there is no control over the final position of the breast image and its geometrical shape. Thus, it could lead to a very different representation of the prior and post-contrast images, limiting the accuracy of the method. Finally, validation on a dynamic phantom with known ground truth [33] is required before the protocol can be applied to scan patients. In terms of phantom design, addition of solid structures representing non-perfused fibro-glandular tissue is necessary to confirm that the background image is fully inherited from high-quality prior image without causing artifacts when using PICCS, but the optimization performed here would most probably not be affected.

3.5 Conclusion

We propose a protocol for DCE bCT tailored to provide accurate quantitative values at low dose. With this protocol, the pre-contrast image and a 300 seconds long sequence of post-contrast images can be acquired at 22.37 mGy.

Acknowledgments

This study was supported by ERC grant 864929.

Conflict of interest statement

Ioannis Sechopoulos is a Scientific Advisory Board member of Koning Corp.

Appendix

Joint optimization

In listing 1, we provide pseudo-code for the joint optimization method combining IMPACT and PICCS methods. In this implementation the total variation of volume I is calculated as

$$\mathrm{TV}(I) = \sum_{i,j,k} \sqrt{(I_{i,j,k} - I_{i-1,j,k})^2 + (I_{i,j,k} - I_{i,j-1,k})^2 + (I_{i,j,k} - I_{i,j,k-1})^2 + \varepsilon},$$

with indices wrapping around and $\varepsilon = 10^{-8}$, and the total gradient as

$$\begin{aligned}\mathrm{grad}_{TV}(I) = {} & (I_{i,j,k} - I_{i-1,j,k}) + (I_{i,j,k} - I_{i,j-1,k}) + (I_{i,j,k} - I_{i,j,k-1}) \\ & + (I_{i+1,j,k} - I_{i,j,k}) + (I_{i+1,j,k} - I_{i+1,j-1,k}) + (I_{i+1,j,k} - I_{i+1,j,k-1}) \\ & + (I_{i,j+1,k} - I_{i-1,j+1,k}) + (I_{i,j+1,k} - I_{i,j,k}) + (I_{i,j+1,k} - I_{i,j+1,k-1}) \\ & + (I_{i,j,k+1} - I_{i-1,j,k+1}) + (I_{i,j,k+1} - I_{i,j-1,k+1}) + (I_{i,j,k-1} - I_{i,j,k}).\end{aligned}$$

IMPACT update step

Let us assume that the image to be reconstructed has J voxels, while projections have I pixels. Then, given the measured projections y_i, forward projector operator $\mathrm{FP} = \sum_{j=1}^{J} l_{ij}$ and corresponding backprojector operator $\mathrm{BP} = \sum_{i=1}^{I} l_{ij}$, the update step for the IMPACT method can be written as follows:

$$\mu_j^{n+1} = \mu_j^n + \frac{\phi_j' \cdot \mathrm{BP}\left(\kappa_i Y_i^{\Phi}\right) + \theta_j' \cdot \mathrm{BP}\left(\kappa_i Y_i^{\Theta}\right)}{\phi_j' \cdot \mathrm{BP}(M_i) + \theta_j' \cdot \mathrm{BP}(N_i)},$$

where

$$\phi_j' = \left(\frac{d\phi}{d\mu}\right)|_{\mu_j^n}$$

$$\theta_j' = \left(\frac{d\theta}{d\mu}\right)|_{\mu_j^n}$$

$$\kappa_i = 1 - \frac{y_i}{\hat{y}_i}$$

Algorithm 1 Joint optimization

```
Require: Fully sampled prior projections: $P_f$
Require: Sparsely sampled projections: $P_s$
Require: Backtracking parameters: $c \leftarrow 0.1$, $\beta \leftarrow 0.5$, $\gamma \leftarrow 10$
$I^0_{prior} \leftarrow$ initial guess (e.g. 0)
for i = 1 to N do
    $I^i_{prior} \leftarrow \mathrm{IMPACT}(P_f, I^{i-1}_{prior})$
end for
$I^0_{piccs} \leftarrow I^N_{prior}$
for i = 1 to M do
    $I^i_{piccs} \leftarrow \mathrm{IMPACT}(P_f, I^{i-1}_{piccs})$
    $G_{TV} \leftarrow \alpha\ \mathrm{grad}_{TV}(I^i_{piccs} - I^N_{prior}) + (1-\alpha)\mathrm{grad}_{TV}(I^i_{piccs})$
    $TV_i \leftarrow \alpha \mathrm{TV}(I^i_{piccs} - I^N_{prior}) + (1-\alpha)\mathrm{TV}(I^i_{piccs})$
    $TV_\gamma \leftarrow \alpha \mathrm{TV}(I^i_{piccs} - I^N_{prior} - \gamma G_{TV}) + (1-\alpha)\mathrm{TV}(I^i_{piccs} - \gamma G_{TV})$
    while $TV_\gamma > TV_i - c\gamma \|G_{TV}\|_2$ do
        $\gamma \leftarrow \beta\gamma$
        $TV_\gamma \leftarrow \sum_j \alpha \nabla(I^i_{piccs} - I^N_{prior} - \gamma G_{TV}) + (1-\alpha)\nabla(I^i_{piccs} - \gamma G_{TV})$
    end while
    $I^i_{piccs} \leftarrow I^i_{piccs} - \gamma G_{TV}$
    $\gamma \leftarrow 10\gamma$
end for
```

$$M_i = \mathrm{FP}\left(\phi_j'\right) \cdot \left(Y_i^{\Phi\Phi}\kappa_i + \frac{y_i\left(Y_i^{\Phi}\right)^2}{\hat{y}_i^2}\right) + \mathrm{FP}\left(\theta_j'\right) \cdot \left(Y_i^{\Phi\Theta}\kappa_i + \frac{y_i Y_i^{\Phi} Y_i^{\Theta}}{\hat{y}_i^2}\right)$$

$$N_i = \mathrm{FP}\left(\theta_j'\right) \cdot \left(Y_i^{\Theta\Theta}\kappa_i + \frac{y_i\left(Y_i^{\Theta}\right)^2}{\hat{y}_i^2}\right) + \mathrm{FP}\left(\phi_j'\right) \cdot \left(Y_i^{\Phi\Theta}\kappa_i + \frac{y_i Y_i^{\Phi} Y_i^{\Theta}}{\hat{y}_i^2}\right),$$

and polychromatic sinograms are given as:

$$Y_i^{\Phi} = \sum_{e=1}^{E} \Phi(e)\, \hat{y}_{ie}$$

$$Y_i^{\Theta} = \sum_{e=1}^{E} \Theta(e)\, \hat{y}_{ie}$$

$$Y_i^{\Phi\Phi} = \sum_{e=1}^{E} \Phi(e)\Phi(e)\, \hat{y}_{ie}$$

$$Y_i^{\Theta\Theta} = \sum_{e=1}^{E} \Theta(e)\Theta(e)\, \hat{y}_{ie}$$

$$Y_i^{\Phi\Theta} = \sum_{e=1}^{E} \Phi(e)\Theta(e)\,\hat{y}_{ie}.$$

Finally, the polychromatic forward model is given by

$$\hat{y}_i = \sum_e \hat{y}_{ie} = \sum_e b_{ie} \exp\left(-\Phi(e)\mathrm{FP}\left(\phi\left(\mu_j^n\right)\right) - \Theta(e)\mathrm{FP}\left(\theta\left(\mu_j^n\right)\right)\right),$$

where b_{ie} is the energy fluence for each pixel i at energy bin e in absence of any materials between the x-ray source and the detector.

References

[1] M. Mikerov, J. J. Pautasso, L. Goris, K. Michielsen, and I. Sechopoulos. 4d dynamic contrast-enhanced breast ct: Phantom-based reconstruction parameter optimization for iodine quantification. *Medical Physics*, 52(4):2212–2223, 2025.

[2] M. Arnold, E. Morgan, H. Rumgay, A. Mafra, D. Singh, M. Laversanne, J. Vignat, J. R. Gralow, F. Cardoso, S. Siesling, and I. Soerjomataram. Current and future burden of breast cancer: Global statistics for 2020 and 2040. *The Breast*, 66:15–23, 2022.

[3] A. N. Wilkinson, J. Billette, L. F. Ellison, M. A. Killip, N. Islam, and J. M. Seely. The impact of organised screening programs on breast cancer stage at diagnosis for canadian women aged 40–49 and 50–59. *Current Oncology*, 29(8):5627–5643, 2022.

[4] K. R. Broglio, M. Quintana, M. Foster, M. Olinger, A. McGlothlin, S. M. Berry, J. Boileau, C. Brezden-Masley, S. Chia, S. Dent, K. Gelmon, A. Paterson, D. Rayson, and D. A. Berry. Association of Pathologic Complete Response to Neoadjuvant Therapy in HER2-Positive Breast Cancer With Long-Term Outcomes: A Meta-Analysis. *JAMA Oncology*, 2(6):751, June 2016.

[5] D. J. P. van Uden, M. C. van Maaren, P. Bult, L. J. A. Strobbe, J. J. M. van der Hoeven, C. F. J. M. Blanken-Peeters, S. Siesling, and J. H. W. de Wilt. Pathologic complete response and overall survival in breast cancer subtypes in stage III inflammatory breast cancer. *Breast Cancer Research and Treatment*, 176(1):217–226, July 2019.

[6] M. Golshan, S. Loibl, S. M. Wong, J. B. Huober, J. O'Shaughnessy, H. S. Rugo, N. Wolmark, M. D. McKee, D. Maag, D. M. Sullivan, O. Metzger-Filho, G. Von Minckwitz, C. E. Geyer, W. M. Sikov, and M. Untch. Breast Conservation After Neoadjuvant Chemotherapy for Triple-Negative Breast Cancer: Surgical Results From the BrighTNess Randomized Clinical Trial. *JAMA Surgery*, 155(3):e195410, March 2020.

[7] P. Cortazar, L. Zhang, M. Untch, K. Mehta, J. P. Costantino, N. Wolmark, H. Bonnefoi, D. Cameron, L. Gianni, P. Valagussa, S. M. Swain, T. Prowell, S. Loibl, D. L. Wickerham, J. Bogaerts, J. Baselga, C. Perou, G. Blumenthal, J. Blohmer, E. P. Mamounas, J. Bergh, V. Semiglazov, R. Justice, H. Eidtmann, S. Paik, M. Piccart, R. Sridhara, P. A. Fasching, L. Slaets, S. Tang, B. Gerber, C. E. Geyer, R. Pazdur, N. Ditsch, P. Rastogi, W. Eiermann, and G. von Minckwitz. Pathological complete response and long-term clinical benefit in breast cancer: the ctneobc pooled analysis. *The Lancet*, 384(9938):164–172, 2014.

[8] H. M. Kuerer, L. A. Newman, A. U. Buzdar, K. K. Hunt, K. Dhingra, T. A. Buchholz, S. M. Binkley, F. C. Ames, B. W. Feig, M. I. Ross, G. N. Hortobagyi, and S. E. Singletary. Residual metastatic axillary lymph nodes following neoadjuvant chemotherapy predict disease-free survival in patients with locally advanced breast cancer. *The American Journal of Surgery*, 176(6):502–509, 1998.

[9] A. B. Tadros, W. T. Yang, S. Krishnamurthy, G. M. Rauch, B. D. Smith, V. Valero, D. M. Black, A. Lucci, A. S. Caudle, S. M. DeSnyder, M. Teshome, C. H. Barcenas, M. Miggins, B. E. Adrada, T. Moseley, R. F. Hwang, K. K. Hunt, and H. M. Kuerer. Identification of Patients With Documented Pathologic Complete Response in the

Breast After Neoadjuvant Chemotherapy for Omission of Axillary Surgery. *JAMA Surgery*, 152(7):665, July 2017.

[10] J. H. Volders, V. L. Negenborn, P. E. Spronk, N. M. A. Krekel, L. J. Schoonmade, S. Meijer, I. T. Rubio, and M. P. van den Tol. Breast-conserving surgery following neoadjuvant therapy-a systematic review on surgical outcomes. *Breast Cancer Research and Treatment*, 168(1):1–12, February 2018.

[11] L. W. Turnbull. Dynamic contrast-enhanced MRI in the diagnosis and management of breast cancer. *NMR in Biomedicine*, 22(1):28–39, January 2009.

[12] M. Caballo, C. Fedon, L. Brombal, R. Mann, R. Longo, and I. Sechopoulos. Development of 3D patient-based super-resolution digital breast phantoms using machine learning. *Physics in Medicine & Biology*, 63(22):225017, November 2018.

[13] A. Sarno, G. Mettivier, and P. Russo. Dedicated breast computed tomography: Basic aspects. *Medical Physics*, 42(6Part1):2786–2804, 2015.

[14] R. E. Alvarez and A. Macovski. Energy-selective reconstructions in X-ray computerised tomography. *Physics in Medicine and Biology*, 21(5):733–744, September 1976.

[15] G. Chen, J. Tang, and S. Leng. Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets. *Medical physics*, 35(2):660–663, 2008.

[16] N. D. Prionas, K. K. Lindfors, S. Ray, S. Huang, L. A. Beckett, W. L. Monsky, and J. M. Boone. Contrast-enhanced dedicated breast ct: Initial clinical experience. *Radiology*, 256(3):714–723, 2010. PMID: 20720067.

[17] P. M. Gazi, S. Aminololama-Shakeri, K. Y., and J. M. Boone. Temporal subtraction contrast-enhanced dedicated breast ct. *Physics in Medicine Biology*, 61(17):6322, aug 2016.

[18] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. An iterative maximum-likelihood polychromatic algorithm for CT. *IEEE Transactions on Medical Imaging*, 20(10):999–1008, October 2001.

[19] D. R. White, J. Booz, R. V. Griffith, J. J. Spokas, and I. J. Wilson. Report 44. *Journal of the International Commission on Radiation Units and Measurements*, os23(1):NP–NP, 04 2016.

[20] ICRP, 2009. Adult Reference Computational Phantoms. ICRP Publication 110. Ann. ICRP 39 (2).

[21] P. T. Lauzier, J. Tang, and G. Chen. Prior image constrained compressed sensing: Implementation and performance evaluation. *Medical Physics*, 39(1):66–80, 2012.

[22] Z. Qi and G. Chen. Performance studies of four-dimensional cone beam computed tomography. *Physics in Medicine Biology*, 56(20):6709, sep 2011.

[23] J. J. Pautasso, K. Michielsen, and I. Sechopoulos. Technical note: Characterization, validation, and spectral optimization of a dedicated breast ct system for contrast-enhanced imaging. *Med Phys*, April 2024.

[24] S. Richard, D. B. Husarik, G. Yadava, S. N. Murphy, and E. Samei. Towards task-based assessment of ct performance: System and object mtf across different reconstruction algorithms. *Medical Physics*, 39(7Part1):4115–4122, 2012.

[25] J. J. Pautasso, M. Caballo, M. Mikerov, J. M. Boone, K. Michielsen, and I. Sechopoulos. Deep learning for x-ray scatter correction in dedicated breast ct. *Medical Physics*, 50(4):2022–2036, 2023.

[26] I. Sechopoulos, S. Vedantham, S. Suryanarayanan, C. J. D'Orsi, and A. Karellas. Monte carlo and phantom study of the radiation dose to the body from dedicated CT of the breast. *Radiology*, 247(1):98–105, April 2008.

[27] I. Sechopoulos, S. S. Feng, and C. J. D'Orsi. Dosimetric characterization of a dedicated breast computed tomography clinical prototype. *Med Phys*, 37(8):4110–4120, 2010. [published correction appears in Med Phys. 2012 Apr;39(4):2314].

[28] J. M. Boone, N. Shah, and T. R. Nelson. A comprehensive analysis of dgn(ct) coefficients for pendant-geometry cone-beam breast computed tomography. *Med Phys*, 31(2):226–235, 2004.

[29] M. J. Yaffe, J. M. Boone, N. Packard, and et al. The myth of the 50-50 breast. *Med Phys*, 36(12):5437–5443, 2009.

[30] L. Volterrani, F. Gentili, A. Fausto, V. Pelini, T. Megha, F. Sardanelli, and M. A. Mazzei. Dual-energy ct for locoregional staging of breast cancer: Preliminary results. *American Journal of Roentgenology*, 214(3):707–714, 2020. PMID: 31939699.

[31] B. E. Nett, R. Brauweiler, W. Kalender, H. Rowley, and G. Chen. Perfusion measurements by micro-ct using prior image constrained compressed sensing (piccs): initial phantom results. *Physics in Medicine Biology*, 55(8):2333, apr 2010.

[32] M. Mikerov, K. Michielsen, N. Moriakov, and I. Sechopoulos. Non-rigid motion compensation for breast ct. In *2022 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–2, 2022.

[33] L. C. Goris, M. Mikerov, J. J. Pautasso, and I. Sechopoulos. In-line spectroscopy for iodine quantification in dynamic contrast enhanced dedicated breast ct. n/a(n/a). Submitted to *Medical Physics*.

4

USING DIFFERENT MODALITIES TO CREATE ARTIFICIAL DATA

Extending dedicated breast CT to dynamic contrast-enhanced breast CT will allow functional data by following contrast distribution over multiple time points. As a result of longer scans, unwanted patient motion becomes more likely and motion correction is necessary to avoid artifacts. A crucial part of the development of a motion compensation method is its validation on data with known ground truth. Thus, anthropomorphic phantoms with realistic motion patterns are needed.

We present a method to combine motion vector fields, extracted from dynamic contrast-enhanced breast magnetic resonance imaging (DCE-MRI), with digital breast phantoms. DCE-MRI is an ideal source for these data due to its frequent clinical use and a similar prone patient positioning compared to breast CT. Our algorithm consists of three steps. First, the inter-scan motion vector fields are obtained by registration of consecutive images in the DCE-MRI sequence. In the second stage, the digital breast phantoms are aligned on MRI scans using an affine transformation. Finally, the obtained motion vector fields are transformed and applied to the phantoms after parameterization such that the total motion is distributed smoothly in time.

The applied motion is evaluated in reconstructions of the simulated breast CT acquisition by qualitative comparison to clinical cases with intra-scan motion. We show that phantoms with simulated motion exhibit the same artifacts as in clinical data such as smooth transitions at the tissue interfaces and ghosting of fine structures.

This chapter appeared as
🗎 *M. Mikerov, K. Michielsen, N. Moriakov, and I. Sechopoulos.*
Adding patient motion from DCE-MRI to anthropomorphic phantoms for dedicated breast CT
SPIE Medical Imaging 2022 [1].

4.1 DESCRIPTION OF PURPOSE

Dynamic contrast-enhanced dedicated breast CT (DCE-BCT) is expected to provide functional information at high spatio-temporal resolution [2]. However, with increased duration of the scan sequence as compared to acquisition of a single anatomical image, unwanted patient motion becomes more likely. In image acquisitions with comparable patient positioning during dynamic contrast-enhanced magnetic resonance imaging (DCE-MRI), the expected displacement magnitude for a given voxel was estimated to be 0.59 ± 0.23 mm between the pre-enhanced and first post-enhanced volumes [3]. Motion of this order of magnitude (almost twice the voxel size of breast CT) will result in motion artifacts in breast CT that will lead to loss of spatial resolution and decrease of the accuracy of the estimation of contrast agent concentration. To avoid these problems, a motion compensation algorithm needs to be applied, which in turn requires data with known ground truth for validation.

In this work, we present our method to create such data for validation by modifying existing voxelized digital breast phantoms [4] by incorporating realistic motion patterns. Since DCE-BCT is still in the development phase and no data is available, we obtain motion from existing DCE-MRI sequences by extracting the motion vector fields (MVFs) using image registration (4.2.2). After that, we estimate the transformation needed for the phantoms to match the anatomy of the breast imaged with MRI and, finally, apply extracted MVFs to them (4.2.3).

4.2 METHODS

4.2.1 PHANTOM AND PATIENT DATA

Due to limited breast DCE-BCT data, we decided to obtain MVFs from inter-scan motion in DCE-MRI, where patients are scanned in a similar prone position, and which is frequently used in clinical practice. To get high quality data, we selected image series acquired using the compressed sensing volume interpolated breath-hold examination (CS-VIBE) sequence. This choice was motivated by the absence of view sharing (and resulting temporal blur) and the short acquisition time for individual images in this sequence (4.5 s), which is close to the time of a whole revolution of our intended breast CT system (6 s). Eleven patient cases from an ongoing ethics board-approved study examining the CS-VIBE sequence were included in this work.

The anthropomorphic digital phantoms to which the DCE-MRI motion is applied are based on segmented clinical breast CT scans described in our previous work [4, 5]. In summary, each voxelized phantom contains voxels labeled as skin, adipose and glandular tissues. Moreover, phantom resolution is improved to 0.091 mm from the initial clinical resolution of 0.273 mm using a neural network [5].

To evaluate the realism of our method, we compared the artifacts present in reconstructed phantom images to BCT patient cases with significant intra-scan motion obtained with a dedicated breast CT clinical system (Koning Corp., Norcross, GA). The images were reconstructed using filtered back projection with isotropic voxel size of 0.273 mm with the vendor-specific cupping artifact correction.

4.2.2 MOTION VECTOR FIELDS ESTIMATION FROM MRI SCANS

While MRI and breast CT acquisitions have the same anatomical positioning, the MRI volume is larger and contains both breasts. To match these with the breast CT data, we manually determined a bounding box to crop the right breast in all MRI images prior to the registration of subsequent images in the sequence. Next, N4 bias field correction was applied to remove low-frequency intensity non-uniformity [6].

We used BSpline image registration to extract the MVFs from the DCE-MRI sequences. This method is intensity based, and as such can be influenced by changing intensity due to the varying presence of a contrast agent. We observed that if the time separation between the images was large, not only the sought physical motion but also apparent motion due to contrast agent was captured. However, we assumed that this effect can be neglected if the registration is limited to subsequent frames, which are separated by only 4.5 s. A Bspline transformation $T_{bspline}$ was obtained for all volume pairs using 3D Slicer [7] in an automated workflow controlled by a Python script. Only MVFs with average motion of at least half the MRI voxel size (0.4 mm) are considered for further processing due to the decreasing accuracy of registration at that scale.

4.2.3 APPLICATION OF THE MOTION FIELDS

In order to transfer the motion from the MRI source to the phantom data, both volumes need to be aligned. However, this step can tolerate lower accuracy since the Bspline-based MVF extends smoothly outside the breast volume. Because of the difficulty of aligning the source and target breasts with potentially large variations in size and anatomical structure, this step was carried out manually in 3D Slicer to obtain an affine transformation matrix T_{affine}. Then, the transformations representing patient motion and phantom alignment are combined as follows:

$$T_{combined} = T_{affine}^{-1} \circ T_{bspline} \circ T_{affine}. \tag{4.1}$$

When applied to a phantom, this transformation scales and moves it to the location of the segmented MRI breast, applies the measured patient motion, and finally, restores the scaling and moves the modified phantom back to the initial location.

For simulation of a BCT acquisition, we also require the transformation for each intermediary point between the start and end of the motion. To obtain this, we parameterize the MVFs as follows:

$$d(x,y,z,t) = f(t) \cdot d(x,y,z,t=0) \text{ for } t = [0,1]. \tag{4.2}$$

The upper limit of the function's domain corresponds to the time required for the acquisition of the complete set of projections (n=300 in our case). Many possible options for $f(t)$ exist; for the results shown here, we selected two scenarios. The first case is a steady continuous motion during the whole acquisition. The second motion pattern is modeled by a sigmoid function where most motion is concentrated in a short span of 30 frames approximately halfway through the scan. Results for both scenarios are shown in figure 4.1.

4.2.4 ACQUISITION SIMULATION AND RECONSTRUCTION

To simulate a breast CT scan of the moving phantom, 300 volumes were created, each corresponding to a different time point t of the parameterization. In this process, discrete

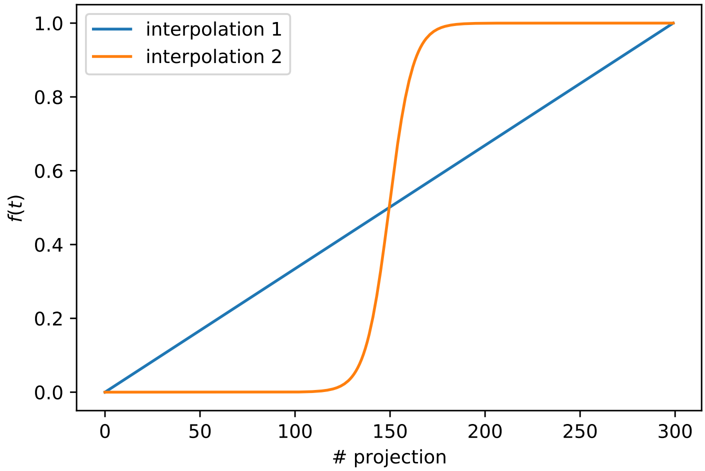

Figure 4.1: Two example functions that were used to create intermediate MVFs. The blue line describes continuous motion. The orange curve represents a case of sudden repositioning halfway during the acquisition of the breast.

tissue compositions in the original phantom were replaced by interpolated tissue mixtures at tissue edges. For example, if the adipose-glandular interface ended up at the center of a voxel after applying the MVF, that voxel would be assigned a tissue composition of 50% glandularity. This way smoother transitions between tissue types during the reconstruction were achieved.

The forward projection of each of the 300 volumes was simulated at its specified x-ray source position with a monochromatic beam at 30 keV, after which Poisson noise was added. The linear attenuation coefficients for skin, glandular, and adipose tissues were set to 0.395, 0.374, and 0.282 cm^{-1}, respectively. The photon count in the air was set to $3x10^5$. The phantom projections were reconstructed using the FDK algorithm with a Shepp-Logan filter. The voxel size was set to 0.273 mm, matching the patient reconstruction. Hence, the resolution in the reconstructions is three times lower than that of the simulated phantom. No cupping artifact correction was required since the simulation does not include scatter or beam hardening.

4.3 RESULTS

Out of 11 patient cases, with a total of 231 evaluated image pairs, two MVFs with sufficiently large motion were identified. Figure 4.2 depicts the displacement of each voxel relative to the initial volume of the glandular tissue in a transformed phantom. It was created by transforming the voxelized segmentations of glandular tissues (at initial position and after motion) into meshes using marching cubes and subtraction of those meshes [8]. Color variation indicates that the motion is not uniform. In the region where most displacement can be observed, the magnitude is in the range from 0.23 mm to 0.70 mm. This result corresponds well with result from literature [3].

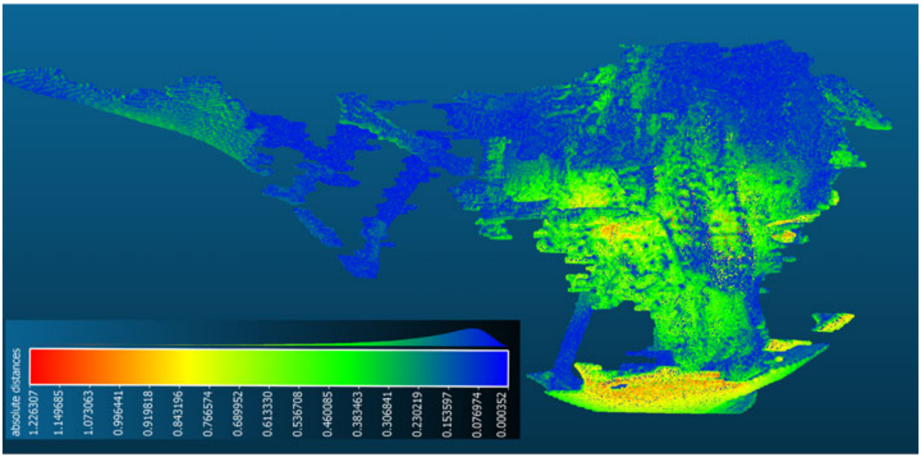

Figure 4.2: Segmented glandular tissue with colormap representing displacement magnitude.

Figure 4.3 shows artifacts present in a patient case and in our simulation. The window widths and levels were selected to emphasize different motion artifacts that can be found in the patient case. Particularly, the display range in the left column highlights streak artifacts in the air. The top row displays the reconstruction of an example patient case with significant intra-scan motion. Some specific features due to this motion are, for example: 4.3a) the prominent background streak artifacts which originate when the skin-air interface changes its position during acquisition; 4.3b) appearance of ghosting of fine structures and skinline; and 4.3c) shadowing artifact at the border of the glandular tissue. The reconstructed phantom with sigmoid motion interpolation is shown in the middle row of figure 4.3 (4.3d, 4.3e, and 4.3f). All three artifacts observed in the patient image, can be identified here as well. Taking into account the parameterization of motion, these effects can be explained by inconsistencies in the projection data obtained from the initial and transformed positions. The artifacts are less prominent when the total motion is distributed equally over all projections with linear interpolation (4.3g, 4.3h, and 4.3i). Especially ghosting of fine structures is absent from the reconstructions. This can be explained by the fact that there is rather continuous motion that does not cause projections that are acquired at two or more different positions.

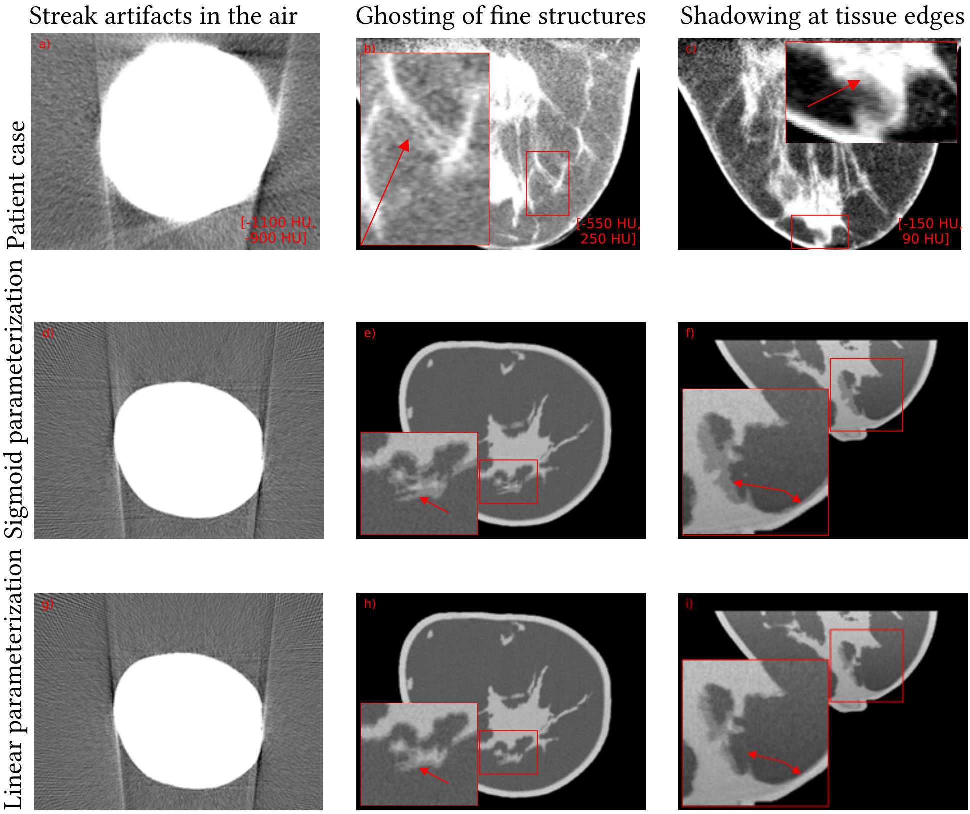

Figure 4.3: Reconstructions of the patient case (top) and phantom with added motion (middle and bottom). The middle row shows motion interpolation with sigmoid function, while the bottom row represents the linear case (cfr figure 4.1). The display range in the top row is unique for each column and is shown in images. The images in the middle and bottom rows were not transformed to HU and have arbitrary scale to visualize the artifacts.

4.4 Discussion

In this work we have presented a method to create digital phantoms for breast CT that include patient motion and could be used to validate motion correction algorithms. For these phantoms to be useful in this regard, they need to include anatomical structures, which are inherited from the original static phantoms, and the phantom deformation caused by the applied motion needs to be realistic and result in similar artifacts as seen in clinical cases, which is demonstrated in figure 4.3. While we only found two occurrences of a sufficiently large motion in the first 11 examined cases, it indicates that the approach works. More motion patterns can be created as the dataset grows.

A substantial limitation of this method is that the whole process is laborious and quite inefficient computationally. Thus, its application is limited to final testing of motion correction methods and not, for example, to create large datasets for machine learning. As such, the limited number of available DCE-MRI sequences that can be used as input for our method is less problematic.

Attempting to apply the MFV derived from a specific breast to a phantom with a potentially very different size and different interior structure poses some challenge. In this work, the affine transformation between the two was set by hand, but at least partial automation is preferred before attempting to process further data. A specific consequence of matching motion from different sized breasts that we should note is that the magnitude of the original MVF also scales with the affine transformation. And while this means that the applied motion does not exactly match the measurement, the actual motion pattern remains realistic even after scaling its magnitude.

This study was supported by ERC grant 864929.

References

[1] M. Mikerov, K. Michelsen, N. Moriakov, and I. Sechopoulos. Adding patient motion from dce-mri to anthropomorphic phantoms for dedicated breast ct. In *Medical Imaging 2022*, Progress in Biomedical Optics and Imaging - Proceedings of SPIE, United States, April 2022. SPIE.

[2] M. Caballo, K. Michielsen, C. Fedon, and I. Sechopoulos. Towards 4D dedicated breast CT perfusion imaging of cancer: development and validation of computer simulated images. *Physics in Medicine & Biology*, 64(24):245004, December 2019.

[3] A. Melbourne, J. Hipwell, M. Modat, T. Mertzanidou, H. Huisman, S. Ourselin, and D. J. Hawkes. The effect of motion correction on pharmacokinetic parameter estimation in dynamic-contrast-enhanced MRI. *Physics in Medicine and Biology*, 56(24):7693–7708, November 2011. Publisher: IOP Publishing.

[4] M. Caballo, R. Mann, and I. Sechopoulos. Patient-based 4D digital breast phantom for perfusion contrast-enhanced breast CT imaging. *Medical Physics*, 45(10):4448–4460, October 2018.

[5] M. Caballo, C. Fedon, L. Brombal, R. Mann, R. Longo, and I. Sechopoulos. Development of 3D patient-based super-resolution digital breast phantoms using machine learning. *Physics in Medicine & Biology*, 63(22):225017, November 2018.

[6] N.J. Tustison, B.B. Avants, P.A. Cook, Yuanjie Z., A. Egan, P.A. Yushkevich, and J.C Gee. N4ITK: Improved N3 Bias Correction. *IEEE Transactions on Medical Imaging*, 29(6):1310–1320, June 2010.

[7] R. Kikinis, S.D. Pieper, and K.G. Vosburgh. *3D Slicer: A Platform for Subject-Specific Image Analysis, Visualization, and Clinical Support*, pages 277–289. Springer New York, New York, NY, 2014.

[8] W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, August 1987.

5

DO WE NEED TO KNOW THE TRUTH?

Objective: *The goal of this work is to develop a data-driven empirical motion-artifact reduction algorithm for non-rigid motion in dedicated breast CT.*

Methods: *Breast CT is a novel imaging modality that offers fully 3D images at good spatial resolution without breast compression and tissue overlap. However, the slow rotation speed of the gantry in such systems increases the likelihood of motion artifacts. Because of the breast anatomy, motion-artifact reduction techniques need to be able to handle artifacts induced by non-rigid motion, which cannot be modeled due to variable motion patterns and the breasts' inner structure, shape, and size. In this work, we present an iterative data-driven empirical algorithm to reduce motion artifacts in breast CT. The highlight of our method is the ability to perform transformations in the image domain using b-spline fields that are defined for each angle and can be efficiently updated with gradient descent and automatic differentiation.*

Result: *We test the method using a simulation study, on physical phantoms, and clinical cases, and show that it can significantly reduce the appearance of motion artifacts.*

Conclusion and Significance: *This work introduces a fully data-driven empirical motion-artifact reduction capable of identifying and minimizing motion artifacts without an underlying model of motion.*

This chapter appeared as
M. Mikerov, K. Michielsen, N. Moriakov, J. J. Pautasso, S. A. M. Tunissen, A. M. Hernandez, J. M. Boone and I. Sechopoulos.
Empirical motion-artifact reduction for non-rigid motion in dedicated breast CT
TBME 2025 [1].

5.1 INTRODUCTION

Despite early detection through screening and improvements in treatment, breast cancer still takes over half a million lives worldwide every year [2, 3]. Breast cancer screening typically involves acquiring mammographic or digital breast tomosynthesis images. Both imaging modalities not only require breast compression, which some patients experience as painful, but also suffer from tissue overlap, which can lead to missed cancers and false positives in dense breasts. Dedicated breast CT is a further development capable of full 3D imaging that can produce high-quality images of the breast at good spatial resolution without tissue overlap [4]. However, it can not yet match the spatial resolution of mammography or breast tomosynthesis.

To improve the visibility of microcalcifications and other small details, breast CT developers tend to reduce detector pixel size as it improves spatial resolution of the scanner [5, 6]. The reduction in detector pixel size results in longer scan times due to a commensurate increase in readout time of the detector and a reduction in rotation speed to minimize blurring the x-ray source due to the finite pulse width. Patient motion mitigation strategies such as breath-hold become impractical for many patients, due to the rather long acquisition times of a single image, and thus artifacts and blurring due to involuntary patient motion become the factors limiting spatial resolution.

To examine the need for motion-artifact reduction in dedicated breast CT, we previously estimated patient motion in breast imaging modalities with prone patient positioning by registering subsequent frames of dynamic contrast-enhanced (DCE) MRI breast acquisitions [7]. Of the 11 patient cases examined, with a total of 231 evaluated image pairs, we found two cases containing motion, with an estimated amplitude in the range from 0.23 mm to 0.70 mm, which was in good agreement with a similar study by Melbourne *et al.*, who reported motion of 0.55 mm ± 0.21 mm [8]. It is important to note that the patient positioning in DCE-MRI differs slightly from that in breast CT. Although both imaging modalities require prone positioning, only a single breast is suspended in the field of view in breast CT while in MRI both breasts are in the field of view.

Brombal *et al.* estimated patient motion in synchrotron-based breast CT to be in the range from 0.4 mm to 1 mm in the vertical and 0.2 mm to 0.5 mm in the horizontal directions using optical tracking with markers placed on the skin [9]. Interestingly, in their study healthy volunteers exhibited periodic motion patterns in both directions, with a frequency of 9 to 18 cycles/minute, despite our expectation that prone positioning of the patient would mitigate motion of the breast due to breathing. It is worth mentioning that in synchrotron-based breast CT, as opposed to conventional dedicated breast CT, the patient rotates rather than the gantry, which can lead to slightly higher motion amplitude.

Despite the availability of many motion compensation algorithms applicable for cone-beam breast CT [10], to the best of our knowledge, none are used in clinical practice. Several factors make their development difficult. To begin with, the breast is non-rigidly deformable, even when freely hanging in the scanner, requiring modeling of such motion. This limits the usefulness of optical markers since they do not provide explicit information on deformations inside the breast. Next, even though Brombal *et al.* demonstrated that periodic motion is present [9], the motion we observed in our evaluation of DCE-MRI cases presented abrupt motion of large amplitude, which originally motivated us to develop a motion-artifact reduction method. The occurrence of this abrupt motion makes the use of periodic models of motion inapplicable. In addition, limited availability of existing breast CT data with or without motion makes the use of machine learning-based algorithms for motion compensation not suitable at this time.

The uniqueness of each breast and its motion during a scan leaves us with data-driven methods as the most appropriate choice since they estimate motion from projection data only and are not limited in what motion they can model. For a more in depth overview on different data-driven methods, we refer the reader to the recent review by Kyme and Fulton [11]. One class of such methods involves autofocus methods. It is known that patient motion reduces sharpness of the reconstructed image. Thus, it is possible to estimate the motion in such way that its removal improves an image quality metric in the image domain. Capostagno *et al.* optimized the gradient entropy in the image domain by dividing the whole reconstructed volume into smaller cubes that could move rigidly [12]. Moreover, they assumed that motion is smooth during the scan and used temporal b-splines to enforce smoothness of each motion parameter. In their follow-up paper, they used Visual Information Fidelity, a reference-based similarity metric that, when optimized, removes motion artifacts from images as perceived by human observers [13].

Another class of data-driven methods are joint motion estimation and update methods. Sun *et al.* developed an iterative method in which they find in the motion estimation phase a rigid translation of the projected reconstruction from the previous iteration in such a way that it minimizes the mismatch in the projection domain for each projection [14]. Afterwards, the inverse transformation is incorporated into the projection matrix and an image update takes place.

Both classes of methods can correct the whole volume or its parts for rigid motion only. Moreover, both use non-gradient based estimation of motion parameters (covariance matrix adaptation evolution strategy in [12, 13] or numerical finite difference method of gradient estimation in [14]) that require long computation times to evaluate all possible combinations. Consequently, such optimization methods cannot be applied to correct for non-rigid motion since the number of parameters that need to be optimized increases drastically. The alignment of a scan sequence with motion to a reference image simplifies the problem by removing the need to update the reference after each motion estimate, as shown by Flach *et al.* who present a non-rigid registration algorithm based on projection domain misalignment [15]. In their recent work on autofocus methods, Huang *et al.* [16] implemented gradient-based optimization of the motion parameters, which should improve the computational performance; however, they are still dividing the whole volume into large cubes that can only move rigidly.

In this work, we propose a data-driven empirical motion-artifact reduction algorithm in which we minimize the data mismatch in the projection domain by non-rigidly and globally deforming the reconstructed volume. To minimize the already high computational requirements associated with non-rigid motion-artifact reduction techniques, we avoided adding any constraints or physics models with additional cost. Our method builds upon the diffeomorphic demons registration algorithm [17] by moving the cost function to the projection domain, and uses the automatic differentiation [18, 19] functionality of modern deep learning frameworks to efficiently optimize b-spline fields with gradient descent. We demonstrate that our method reduces image artifacts both on digital and physical phantoms, and for patient data.

It is important to point out that the goal is thus to reduce or remove the appearance of motion artifacts and not to estimate the motion that occurred during the scan. This is analogous to how empirical beam hardening corrections aim to avoid the appearance of cupping artifacts without needing explicit knowledge of x-ray spectral properties of the scan.

5.2 ALGORITHM PROPOSED

We will start with a high level overview of our empirical motion-artifact reduction algorithm in section 5.2.1 before describing specific aspects in more detail. First, section 5.2.2 explains estimation of parameters needed to deform the reconstructed volume. Then, sections 5.2.3, 5.2.4, and 5.2.5 clarify how we perform a differentiable transformation in the image domain. Finally, we provide a theoretical analysis of the runtime in section 5.2.6.

5.2.1 DATA-DRIVEN EMPIRICAL MOTION-ARTIFACT REDUCTION ALGORITHM

In general, our method follows the scheme of classical joint motion estimation and image update methods. However, instead of estimating the motion vector field that was applied to the object in the physical space, we estimate b-spline fields that parameterize the deformation of the volume to match the forward projected data to the measured data. The reader is referred to the algorithm pseudocode 2 for a schematic representation of our method.

Each iteration of the method consists of two stages: an image update followed by estimation of b-spline fields. The b-spline field $d_{n,k}$ for angle k at iteration n is defined on a sparse sampling grid, is initialized with zeros, and can then be interpolated using cubic b-splines to create the corresponding dense representation $D_{n,k}$, as further described in section 5.2.5. Both $d_{n,k}$ and $D_{n,k}$ have three channels to specify the deformation in xyz-direction. In the image update phase, we use the inverse of the current image transformation $D_{n,k}$ to transform the backprojection $\mathrm{BP}(P_k)$ of a single projection P_k weighted with cosine weights and filtered with a ramp filter in the image domain. This inverse is guaranteed to exist since we require the transformation to be diffeomorphic, as shown in section 5.2.4. Finally, we sum all transformed backprojections J_k to obtain the reconstructed volume I_n. Since we are working with single-phase data, the projection data must be acquired over at least 180 degrees plus fanbeam angle. In the second phase, we perform the following operation for each projection angle k: we first densify the sparse b-spline field $d_{n,k}$, and use it to transform the current reconstruction estimate I_n before performing a forward projection FP_k of the transformed volume at this angle. We then calculate the L1 loss $\mathcal{L}$ between the resulting forward projection $p_{n,k}$ and the measured projection P_k at this angle and backpropagate the error to update the sparse field $d_{n,k}$. Practical aspects of the optimization process are further discussed in section 5.2.2.

These steps are repeated until the loss is no longer decreasing by a substantial amount as defined by stopping threshold ϵ. In this work, the iterative process was interrupted if the average L1 loss did not decrease by more than 1% compared to the L1 loss at the previous iteration. Alternatively, the algorithm can be run for a fixed number of iterations N, or a combination of these approaches. After the algorithm has finished, the final b-spline field

Algorithm 2 Empirical motion-artifact reduction

```
Require: Log-converted projections $P$
Require: Stopping threshold $\epsilon$
$n \leftarrow 0$
repeat

    Image update
    for all angles $k$ do
        $D_{n,k} \leftarrow \text{Densify}(d_{n,k})$
        $D_{n,k}^{-1} \leftarrow \text{Inverse}(D_{n,k})$
        $J_k \leftarrow \text{BP}(P_k) \circ D_{n,k}^{-1}$
    end for
    $I_n \leftarrow \sum_k J_k$

    Estimation of b-spline fields
    for all angles $k$ do
        $D_{n,k} \leftarrow \text{Densify}(d_{n,k})$
        $p_{n,k} \leftarrow \text{FP}_k(I_n \circ D_{n,k})$
        $\mathcal{L}_{n+1,k} \leftarrow \|P_k - p_{n,k}\|_1$
        $\nabla d \leftarrow \frac{\partial \mathcal{L}_{n+1,k}}{\partial d_{n,k}}$
        $d_{n+1,k} \leftarrow \text{ADAM}\left(d_{n,k}, \nabla d\right)$
    end for
    $n \leftarrow n + 1$
    $\mathcal{L}_{n+1} \leftarrow \sum_k \mathcal{L}_{n+1,k}$
until $(\mathcal{L}_n - \mathcal{L}_{n+1}) < \epsilon$
return d
```

$d_{N,k}$ can be used for one last image update followed by a deformation to an arbitrary angle k. This last step ensures improved volume preservation after motion-artifact reduction. While the transformations are diffeomorphic and thus invertible, their combination at all angles is not guaranteed to preserve volume. However, at each angle the projected volume is constrained by the measurement. The last reconstruction can be performed using a different (smoother) FDK filter, as done in this work. Alternatively, any other reconstruction method that accepts the estimated b-spline fields as input can be used to obtain the final result.

We note here that the filtered backprojection BP and forward projection FP ray tracing operators do not model the physics of the acquisition process fully (e.g., no resolution loss or spectral information is considered), and therefore the calculated loss is a combination of data discrepancy due to patient motion and imperfect modeling. To minimize the effects of these discrepancies, after log-conversion the measured projection data should ideally be processed to correct for the presence of scatter and beam hardening. Masking out the background signal in the projection data also helps by eliminating the effect of streak artifacts in the air in the image domain that could hinder algorithm performance.

5.2.2 Estimation of b-spline fields

The objective of the second phase in the algorithm 2 is to minimize the mismatch between the forward projection of the reconstructed image after deformation and the measured projection at each angle, by optimizing the sparse parameters d. We use gradient descent to solve this minimization problem. The gradient to update the b-spline field $d_{n,k}$ is given by the partial derivatives of the objective function:

$$\frac{\partial \mathcal{L}}{\partial d_{n,k}} = \frac{\partial}{\partial d_{n,k}} \| P_k - \mathrm{FP}_k(I_n \circ D_{n,k}) \|_1. \tag{5.1}$$

A similar approach of optimizing parameters by minimizing the data fidelity function in the projection domain was previously used by Wang *et al.* [20]. However, our approach differs in two key ways: we do not require multi-phase projection data for the optimization, and we also use automatic differentiation, a computational technique used for numerical optimization, to compute necessary gradients. Automatic differentiation applies the chain rule to calculate the derivatives. Therefore, as long as any operation in the chain can be composed of functions for which their analytical derivative is known, such operation will also be differentiable. In strict mathematical sense, the derivative of the L1 loss is not defined at 0, however, it is often assumed 0, as it would mean that no change is required to further reduce the loss value. Our choice of L1 loss as the cost function as opposed to, e.g., MSE loss, is driven by the fact that the derivative of L1 loss with respect to each parameter is always 1 or -1 and does not decrease with decreasing difference between current estimated data and measured data, ensuring that even small mismatches contribute to parameter update but large differences due to possible artifacts do not dominate the loss value. We use ADAM optimizer to adaptively and independently reduce the step size for each element in the b-spline field as it converges to a stationary value [21]. The used ADAM optimizer is provided by PyTorch and has the following parameters: learning rate $\gamma = 0.0001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, weight decay $w = 0$ [19].

The requirement that each operation needs to be differentiable for the automatic differentiation to work means that the forward projection operation (FP) that projects a 3D voxelized volume into a 2D pixelized projection requires special attention. Despite the fact that it can be formulated as a chain of simple algebraic operations, this would be highly inefficient. In our previous work, we have shown that if the forward projector operator FP is the Hermitian adjoint to the backprojector operator BP, the calculation of the derivative of FP requires applying BP to 2D images of the same size as projections filled with gradients that are known by following the chain rule from latter operations in the forward pass. This allows us to use the existing projector implementations with only minor modifications [22].

Similar to any other gradient-based optimization method in high-dimensional space where limited data leads to an underdetermined setup, our approach is not guaranteed to converge to the global optimum. The final solution is a result of local optimization that reduces the data mismatch in the projection domain but is not necessarily the best solution.

5.2.3 VOLUME TRANSFORMATION IN THE IMAGE DOMAIN

A crucial part of our method is the ability to transform 3D volumes in a non-rigid manner in such a way that this transformation is reversible and differentiable using accurate analytical derivatives. The transformation by sampling of a 3D volume I requires a 3D grid D of the same size as I, where each component $-1 \leq \vec{\delta_i} \leq 1$ indicates the relative position in three directions of the voxel in the original volume I that is needed to obtain the value of the matching voxel in the new volume. Jaderberg *et al.* showed that this sampling operation can be efficiently implemented using textures on a GPU with bicubic interpolation [23]. They also proved that this sampling operation is differentiable, and therefore, it is possible to backpropagate the error to update each $\vec{\delta_i}$ to reduce the loss function.

Nevertheless, this operation comes with a downside due to the additional bicubic interpolation in the image domain, which combined with the interpolation due to upsampling of the b-spline fields leads to some resolution loss. In this work, we decided that chaining two simple, already-available components in favor of implementing a highly complex operator that combines the deformation with the projection operators was a sound trade-off.

5.2.4 DIFFEOMORPHIC TRANSFORMATION

In general, the sampling grid D can be constructed or modified in such a way that the transformation is not reversible. We employ the scaling and squaring method from the diffeomorphic demons registration algorithm to construct the sampling grid as a composition of deformations with smaller spatial changes, approximating the exponential mapping that ensures that the resulting composite deformation is reversible and differentiable [17, 24, 25].

5.2.5 SPARSE AND DENSE REPRESENTATIONS OF THE IMAGE TRANSFORMATION

Looking at the overview of the method presented in subsection 5.2.1, we require a 3D grid $D_{n,k}$ for each projection k and each of the xyz-components of $D_{n,k}$ is as large as the reconstructed volume. It is computationally not practical to store all this information for a clinical dataset at a useful resolution. Hence, we need a way to use smaller sparse field $d_{n,k}$

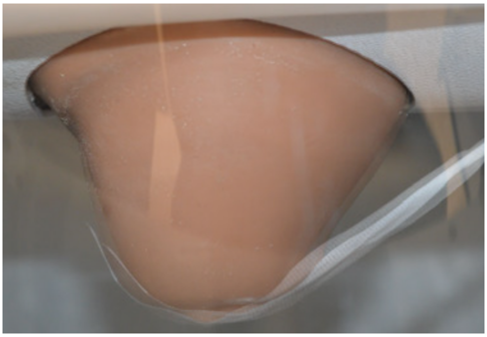

Figure 5.1: Breast phantom for ultrasound, mammography, and MRI, CIRS Model 073, positioned in the gantry of the breast CT scanner with attached sticky tape. The transparent part of the photo depicts the approximate deformation of the phantom.

that can be used to calculate the dense field $D_{n,k}$ on the fly. This operation also needs to be differentiable using analytical differentiation so that we can update $d_{n,k}$ using gradient descent. One way to perform this operation is to use 3D convolution using a 3D cubic b-spline kernel B. Such a (3D) cubic b-spline kernel can be seen as a (3D) convolution of a rectangular function with itself three times [26]. Given a sparse field $d_{n,k}$ with a 3D cubic b-spline kernel B, the resulting dense field is a convolution operation:

$$D_{n,k} = d_{n,k} \star B. \tag{5.2}$$

We use the Airlab library for densification of the field and to ensure that the transformation is diffeomorphic. This implementation allows error backpropagation and therefore can be used for automatic calculation of the gradient [27]. Even though the enforced diffeomorphism and sparse representation add constraints on possible image transformations, we still consider our motion-artifact reduction fully model free and data driven. These constraints are forming a regularization that reduce the likelihood of overfitting on noise rather than excluding anatomically plausible image transformations that can be approximated using these assumptions.

5.2.6 Runtime analysis

We can provide a theoretical analysis of expected runtime using algorithm 2. For each angle k at each iteration n of the method, one FP and two BP (FDK and backpropagation) operations are performed, similar to what is required for a maximum likelihood iterative reconstruction algorithm. Additionally, the transformation that includes densification of sparse b-spline fields is combined with application of scaling and squaring algorithm as well as image resampling (denoted by '$\circ$') is computed twice. Therefore, the total time to process each angle is as follows:

$$\tau_k = \tau_{FP} + 2\tau_{BP} + 2\tau_{trafo} + \tau_{trafo}^{backpr} + \epsilon, \tag{5.3}$$

where ϵ is the time required to perform the remaining mainly algebraic and tensor operations, such as addition of reconstructed images and loading correct geometry configurations.

Runtimes τ_{FP} and τ_{BP} are dependent on the implementation of the projector operators used and system geometry; however, since the FP and two BP operations are performed just for a single projection angle, their combined runtime is expected to be much lower than τ_{trafo}. The densification operation is implemented using 3D convolution of the sparse field with a b-spline kernel and its runtime depends on how sparse the fields are chosen and volume size, with denser sampling decreasing the cost due to the smaller convolution kernel this requires. In general, this time depends on the number of the of steps in the scaling and squaring algorithm. The time τ_{trafo}^{backpr} needed to backpropagate the error back to the sparse b-spline field is closely related to it. We can omit inverting the image transformation from the calculation, since it is just an additional multiplication with -1 in the scaling and squaring method that ensures diffeomorphism in the densification step.

5.3 Materials and Methods

We demonstrate the performance of our method using a simulation study with a digital breast phantom and by applying it to real scans of physical phantoms and clinical cases.

5.3.1 Simulation study

Motion simulation

Realistic linear attenuation coefficients were assigned to skin, adipose, and fibro-glandular tissues in a digital breast phantom [28]. The non-rigid deformation without rotation and translation at each angle k of the phantom was given by the following affine matrix:

$$A_k = \begin{pmatrix} 1 & t_k & t_k & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{5.4}$$

The corresponding maximum transformation was achieved at projection angles 75 and 150. The object kept its original shape for the first and last 75 projection angles. The shape of the object between projection angles 75 and 150 varied between both extreme states according to the function in Fig. 5.2. The digital phantom was deformed individually for each of the 300 projection angles using SimpleITK with the linear interpolation option [29].

Monochromatic simulations were performed at an energy of 30 keV using the geometry of the Koning 1st generation breast CT scanner described in more detail in the following section. The projections were corrupted by Poisson noise with mean $6 \cdot 10^5$ in the air, similar to the noise level observed in clinical scans on the simulated system.

Motion-artifact reduction

The sparse b-spline field was initialized with zeros and had the following dimensions at each projection angle: 3 x 22 x 29 x 29. Here, the first dimension corresponds to the X, Y, and Z components of the deformation parameters. Ramp filter was used to reconstruct all images, i.e., the initial reconstruction affected by motion artifacts and resulting images at each iteration.

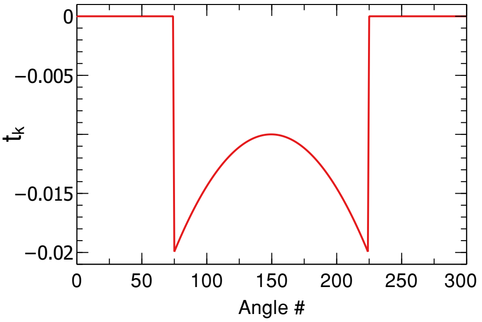

Figure 5.2: Parameterization of deformation given by (5.4). The object has its original shape in the first and last quarters of the scan. The maximum deformation is achieved at projection angles 75 and 150.

We tracked the L1 loss in the projection domain as function of iteration number to verify that the empirical stopping criterion defined in section 5.2.1 is an appropriate choice for this method. The sum of all image intensity values was tracked in both the static position and the deformed position at the first angle to examine volume preservation after motion-artifact reduction. Moreover we investigated whether the resulting b-spline image transformation is correlated to the motion field applied to the object in the object space by plotting the values of a single parameter as a function of projection angle.

Finally, we used the resulting parameters in an iterative reconstruction method, specifically in maximum-likelihood in transmission reconstruction (ML-TR) [30], to show that the estimated b-spline fields can be used with any reconstruction method.

5.3.2 PHYSICAL PHANTOMS AND CLINICAL CASES

The data acquired on three different breast CT systems consists of six datasets: two measurements of physical phantoms with known types of motion and four patient scans without specific information on motion patterns. Table 5.1 lists important parameters of each breast CT system; table 5.2 provides information about each measurement.

Both phantom images were acquired using the Koning 2nd generation breast CT scanner. The first phantom, a *dry orange*, is as the name suggests a solid dried out orange with distinct fiber structures. It was mounted on coupled translational stages that moved in horizontal and vertical directions separately in a loop. The dry orange was attached with tape that caused the orange to oscillate in the air leading to motion that was not quantifiable. While this motion was rigid, it could not be modeled by only using translations. The motion artifacts were reduced using the fully non-rigid empirical motion-artifact reduction method.

The second physical phantom was a *multi-modal breast phantom* (CIRS Model 073 [31]) compatible with ultrasound, mammography, and MRI. This phantom material is easily deformable, so in order to simulate non-rigid motion, it was placed into the gantry with a long piece of tape attached near the nipple as shown in Fig. 5.1. The tape was pulled to deform the phantom before the start of the acquisition, and was then released approximately halfway through the scan, leading to a measurement with non-rigid motion. In addition to visual inspection, a homogeneous area of this phantom was used to evaluate the noise texture before and after motion-artifact reduction by computing the noise power

Table 5.1: Overview of the used breast CT scanners

System	Spectrum	# projections	magnification factor	detector pixel size	additional pixel binning
Koning 1. gen	49 kV with Al (1.5 mm)	300	1.42	$(0.388\ \text{mm})^2$	
Koning 2. gen	49 kV with Al (1.5 mm)	225	1.58	$(0.152\ \text{mm})^2$	2x2
"Doheny"	60 kV with Gd (0.2 mm)	492	1.38	$(0.150\ \text{mm})^2$	2x2

Table 5.2: Overview of the acquired datasets and reconstruction parameters

ID	Scanner	Estimated motion amplitude	Air Kerma at the center of rotation (mGy)	Projection smoothing	Cut-off frequency	Voxel size
Dry orange	Koning 2. gen		17.6		$0.4\ f_{Nyq}$	$(0.25\ \text{mm})^3$
CIRS 073	Koning 2. gen		9.6	2D Gaussian $\sigma = 2$	$0.4\ f_{Nyq}$	$(0.35\ \text{mm})^3$
RUMC_P1	Koning 1. gen	1.6 mm	12.7		f_{Nyq}	$(0.273\ \text{mm})^3$
RUMC_P2	Koning 1. gen	4.0 mm	12.7		f_{Nyq}	$(0.273\ \text{mm})^3$
UCD_P1	"Doheny"		8.2	2D Gaussian $\sigma = 2$	f_{Nyq}	$(0.3\ \text{mm})^3$
UCD_P2	"Doheny"		8.1	2D Gaussian $\sigma = 2$	f_{Nyq}	$(0.3\ \text{mm})^3$

spectrum (NPS). The NPS was obtained by calculating the 2D NPS inside a $97x97$ ROI for 11 consecutive slices and averaging the results. Each slice was detrended to remove any remaining cupping artifact by fitting a plane. To obtain the 1D NPS from the 2D NPS, rotational averaging was performed. Sample distribution before signal averaging was propagated to estimate error bars.

The next two scans, RUMC P1 & P2, were acquired with a Koning 1st generation breast CT scanner. We suspect that they are affected by sudden motion of large amplitude during the scan. A rough estimate based on the double appearance of the skinline indicates motion of approximately 1.6 mm and 4 mm, respectively. The main differences relevant for our experiments between the Koning 1st generation system, used for these patient images, in comparison to the Koning 2nd generation system, used for the phantom images, are the larger pixel size (0.388 mm vs. 0.152 mm), smaller magnification factor (1.42 vs. 1.58), and higher number of projections (300 vs. 225).

The last two scans, UC Davis P1 & P2, were acquired at University of California, Davis, with the "Doheny" breast CT system [5, 32] using relevant acquisition parameteres as detailed in Table 5.1. Also here the motion is suspected but not confirmed. The main difference between the Doheny scans and the RUMC cases is the higher noise level in the Doheny cases resulting from inherently higher spatial resolution utilized during the acquisition.

Clinical results in this work were generated using the FDK algorithm, although our motion-artifact reduction method is compatible with any reconstruction algorithm as shown in the simulation study. This is because motion is only relevant to the geometric transformations performed by the projector and backprojector operations, and is therefore effectively independent from any other aspects in a reconstruction algorithm. As mentioned before, for b-spline fields estimation the FDK algorithm was applied with a ramp filter, additionally, except for the log-conversion, we did not apply any pre-corrections. We only applied a beam hardening correction to the projections for visualisation of final results according to [33].

5.4 RESULTS

5.4.1 SIMULATION STUDY

MOTION-ARTIFACT REDUCTION

Analysis of the average L1 loss over all 300 projections of our simulation experiment is shown in Fig. 5.3a. We observe fast decrease of the L1 loss value in the beginning before it slowly approaches the lowest value. After a certain number of iterations it starts to increase again but does not reach the original L1 loss value. This graph validates our stopping criterion defined by the first iteration when the L1 is not decreasing by more than 1%. In that case, it is the sixth iteration.

We investigated the intensity and volume preservation of the method in Fig. 5.3b. It shows that if the final result is obtained without additional transformation to one of the angles, intensity starts decreasing after a couple of iterations due to volume shrinking. However, when this additional transformation is applied, the overall intensity remains relatively stable for longer periods of time. Small increase of the overall intensity is caused by better delineation of the linear attenuation coefficient on the interfaces between skin,

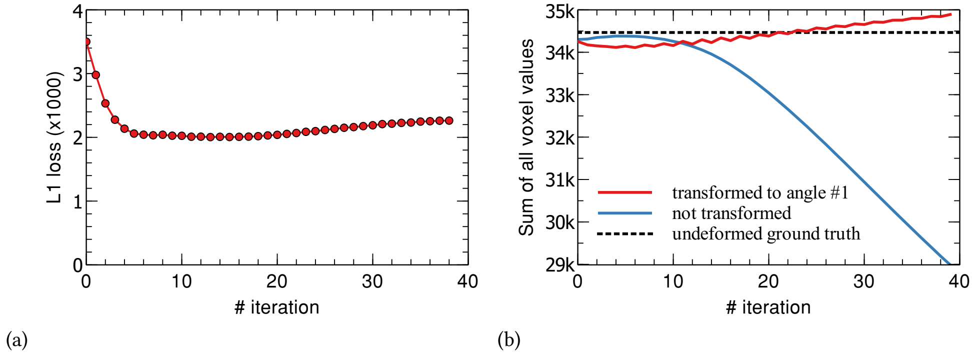

Figure 5.3: (a) L1 loss as function of iteration number in the projection domain. (b) Sum of all intensities inside the breast in the image domain as function of iteration number. The dotted line corresponds to the sum of all intensities in the motion-free reconstruction.

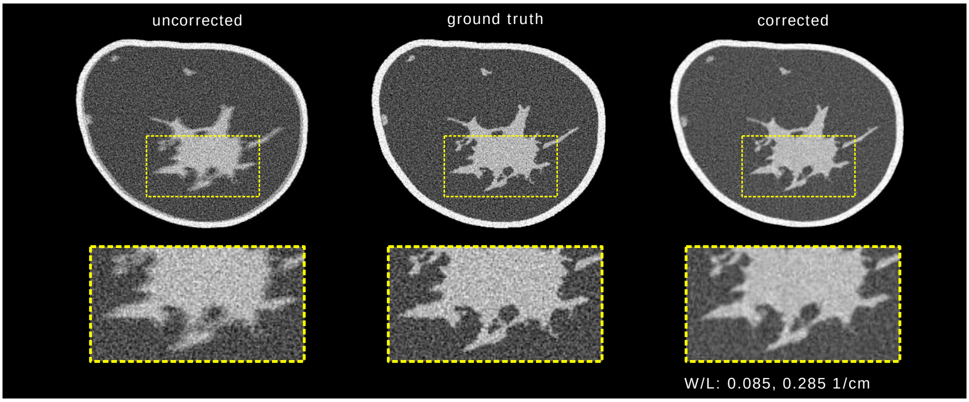

Figure 5.4: Application of proposed algorithm to remove motion artifacts due to non-rigid deformation in a digital breast phantom. Structures in the corrected image appear similar to the ground truth image.

adipose, and fibro-glandular tissues.

Fig. 5.4 presents the results of application of six iterations of the proposed method to remove the motion artifacts from the transformed digital phantom. The characteristic blur caused by motion is reduced, enhancing the delineation of the structures. Therefore, the fibro-glandular structures recover their shape and become visually very similar to the ground truth. However, the corrected image does not attain the shape and location of the ground truth image entirely as is shown in the left difference image in Fig. 5.5. Black contours on the left and white contours on the right indicate that the corrected image is laterally shifted even after motion-artifact reduction. We performed b-spline registration of the corrected and ground truth images using the BRAINS framework in Slicer 3D 5.6.2. [34, 35]. When the difference image is calculated using registered images, the inner structures appear to be almost perfectly matched.

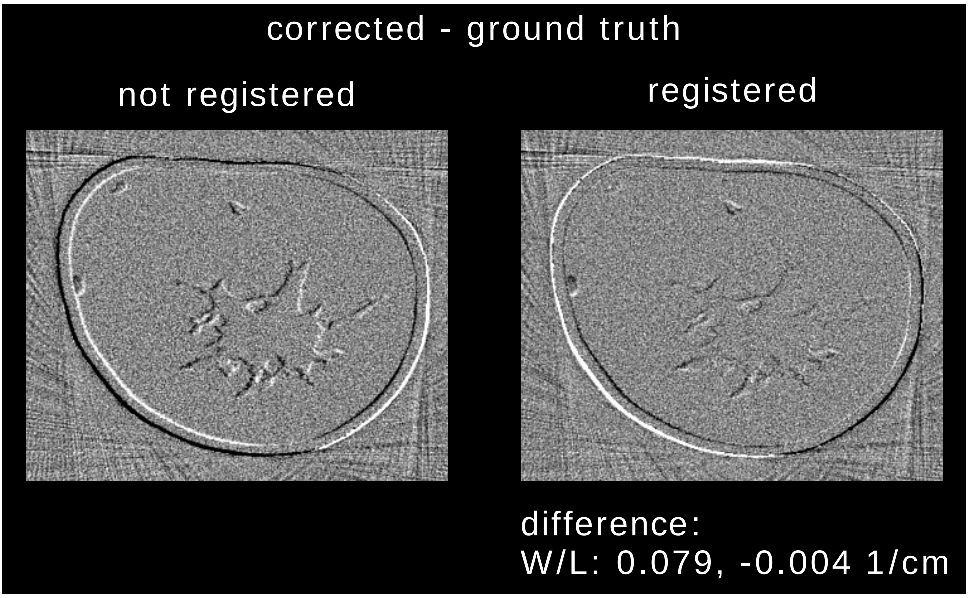

Figure 5.5: The difference between the corrected and ground truth images with and without non-rigid registration to ground truth.

DEFORMATION PARAMETERS

Three raster images in Fig. 5.6 display X, Y, and Z components of the resulting sparse b-spline field at projection angle 150. The bottom right panel shows the X, Y, and Z components of one specific sparse parameter as function of projection angle indicating that the found parameters do not match with the motion vector field applied to the breast phantom, thus leading to unmatched position in Fig. 5.4. The raster images support this claim since Z and Y components would need to be zero everywhere if the b-spline fields correlated to the deformation matrix in (5.4). Additionally, despite the mismatch there is no visual indication of this in the reconstructed volume.

RECONSTRUCTION USING ITERATIVE METHOD

Fig. 5.7 shows the results of iterative reconstruction using the estimated b-spline fields. The number of iterations was 50 in all three cases. Notably, the image blur due to motion is significantly reduced and the shape of the inner structure closely resembles the inner structure of the motion-free ground truth image.

5.4.2 PHYSICAL PHANTOMS

RIGID MOTION

Fig. 5.8 shows uncorrected, corrected, and static scans of the dried orange. All three images were reconstructed with the same FDK kernel according to table 5.2 and are displayed with the same window width/level (W/L). Overall, we observe sharpening of skin and fibers inside the orange. Ghosting artifacts, especially in the right part of the image, are significantly reduced and fibers do not appear as doubled lines anymore.

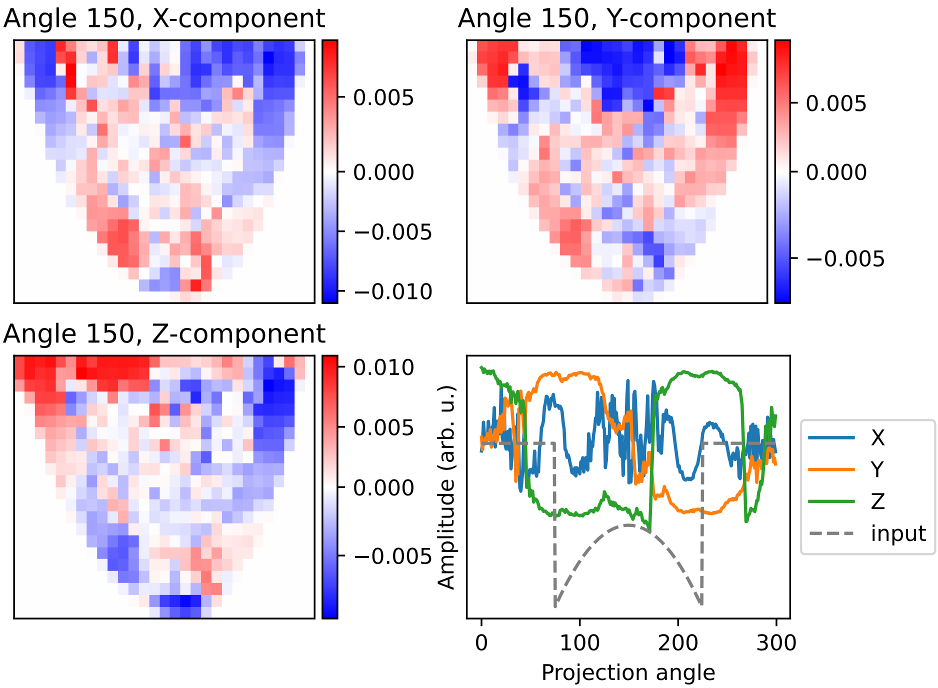

Figure 5.6: Raster images show the three estimated components in the center of the breast. The bottom right plot shows each component of a selected parameter as function of projection angle.

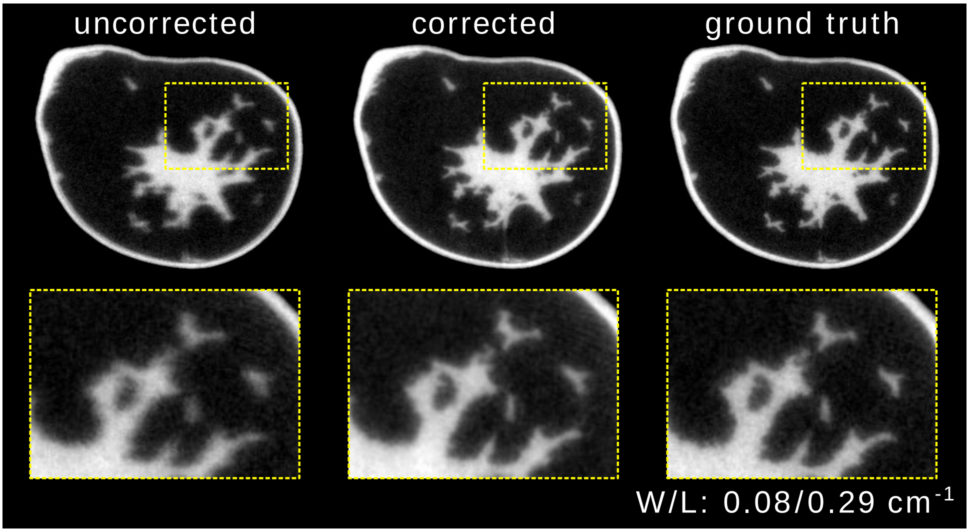

Figure 5.7: ML-TR reconstructions of the breast phantom without and with motion-artifact reduction. Ground truth image is shown for comparison. Blurring due to motion is reduced when motion-artifact reduction is applied.

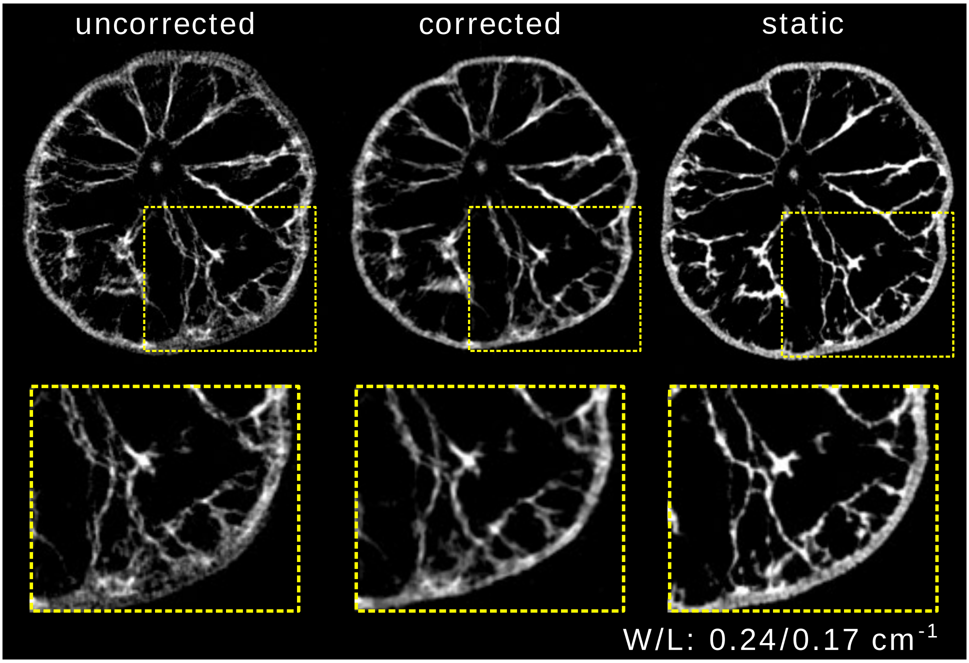

Figure 5.8: Application of empirical motion-artifact reduction to a scan of a rigidly moving dried orange. Comparable slice made with a static acquisition is depicted in the right column. The skin appears thinner and brighter on the corrected image. Additionally, ghosting artifacts of the inner structures are removed.

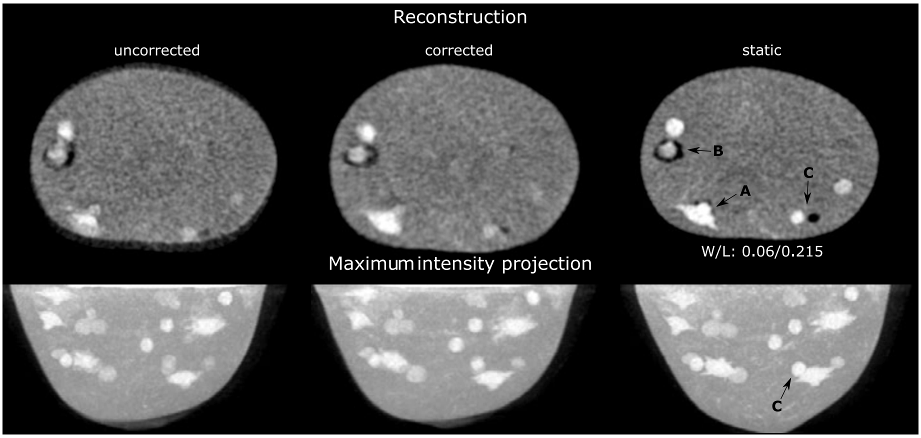

Figure 5.9: Application of motion-artifact reduction to a scan of a non-rigidly deformed ultrasound phantom. Comparable slice made with a static acquisition is depicted in the right column. The line profiles are shown in Fig. The bottom panel shows maximum intensity projection through each volume.

Non-rigid motion

Fig. 5.9 demonstrates that our method successfully reduces motion artifacts from a scan of an object that was non-rigidly deformed during the acquisition. We observe notable improvement in the lower part of the image, where a denser spiculated object **A** separates clearly from the skin and looks more similar to its depiction on the static phantom when motion-artifact reduction is used. Another improvement is visible in the left part of the image, where a nodule **B** takes the correct form and is connected to the rest only in two places.

Because our method reconstructs the image at some (weighted average) position that is different from the position of the object during the static scan, as shown in section 5.4.1, some anatomical structures in the corrected images are not aligned with the static scan but can be found in other slices in the image.

The quantitative analysis of noise texture using NPS confirms the visual impression that the homogeneous region in the corrected image appears to have the same structure when compared to the uncorrected image. Fig. 5.10 shows that, in fact, the frequency content remains the same when the method is applied.

However, not all structures could be restored and there still exists large differences in appearance between corrected and static images, e.g., a combination of bright and dark circles **C** in the lower right part of Fig. 5.9. Even a few slices above or below the same structure does not take the same shape as in the static scan, although it can be found. To verify that all structures that can be detected in the uncorrected scan can also be found in the corrected scan, we resliced the images and calculated the maximum intensity projection (MIP). The resulting MIP projection is shown in the bottom panel of Fig. 5.9. By comparing all three images it becomes apparent that no additional structures are added to the scan when the method is applied but also no structures are missing. The shapes of structures at the bottom and at the right approach their true shape. Nevertheless, the static scan still outperforms the corrected image in both sharpness and visibility of structures. The

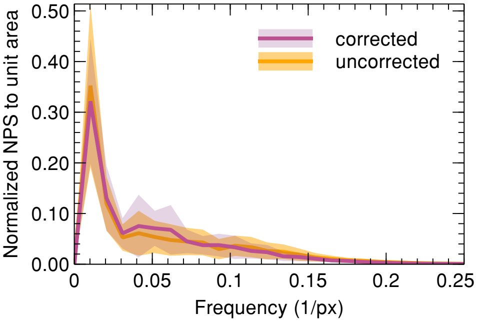

Figure 5.10: Noise power spectrum in the homogeneous region of the ultrasound phantom. The proposed empirical motion-artifact reduction does not change the noise structure.

depiction of the denser sphere in **C** improves only slightly, with a brighter core in the center and it remains prolonged even when the empirical motion-artifact reduction is applied.

5.4.3 CLINICAL CASES

MOTION-ARTIFACT REDUCTION

The result of the application of our method to two clinical cases acquired with the 1st gen. Koning breast CT is displayed in Fig 5.11. In both cases we observe repaired skin line depiction when reconstructed with motion-artifact reduction. Furthermore, the glandular structures become sharper and take on a well-defined shape. It is important to point out that high-contrast objects such as a calcification or blood vessels in the second patient are not affected by the method and remain unchanged both in intensity and location. We also want to point out that the black stripe of what seems adipose tissue along the skin is not an artifact due to our method but is caused by a non-linear detector response at the higher exposure levels. It is evident since this effect is also visible in uncorrected scans although not as clear as on corrected images due to motion blur.

Fig. 5.12a shows the uncorrected and corrected images of the UCD_P1 clinical case after six iterations. The skin line has uniform high contrast appearance when reconstructed with our method. Moreover, all inner structures can be found at the same locations but they appear sharper. The majority of interrupted lines of fibro-glandular structures in the uncorrected image are connected again in the corrected one.

The mean loss for this case has decreased from 0.4312 to 0.4107 as shown in top panel in Fig. 5.12b. The bottom panel of the same figure shows that the improvement was achieved for all projection angles. This figure tells us that the reconstruction was consistent with the first third of the projection angles as the loss at the first iteration has barely decreased in this region. However, after only three iterations the improvement is substantial for all angles. We can see almost no change between iterations five and six, indicating that the iterative process can be stopped as is does not lead to improvement.

Image quality improvement after application of the empirical motion-artifact reduction method can also be observed for the second clinical case from UC Davis in Fig. 5.13. Again, we see sharpening of the structures and better depiction of the skin line. From the qualitative analysis of the projections, we hypothesize that the patient most likely exhibited

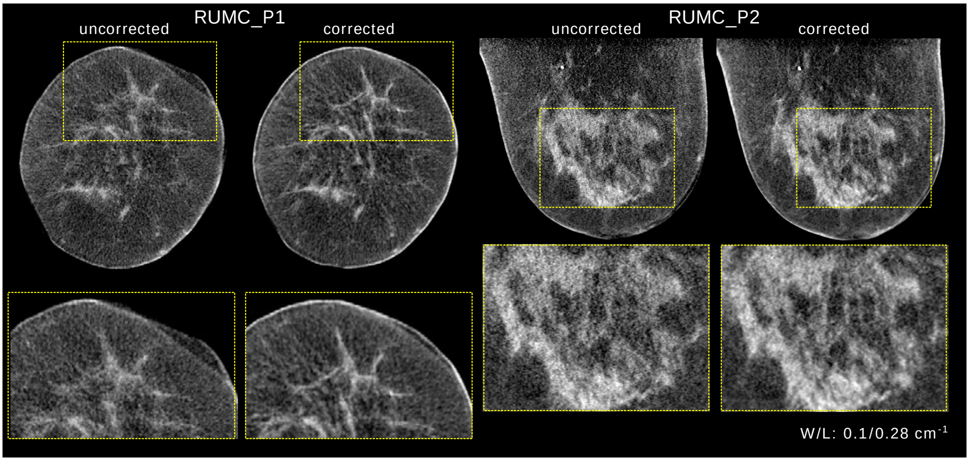

Figure 5.11: Application of empirical motion-artifact reduction to the two clinical scans acquired with a first generation Koning breast CT system. Restored skin line and increased sharpness of inner structures can be observed for both patients in different views.

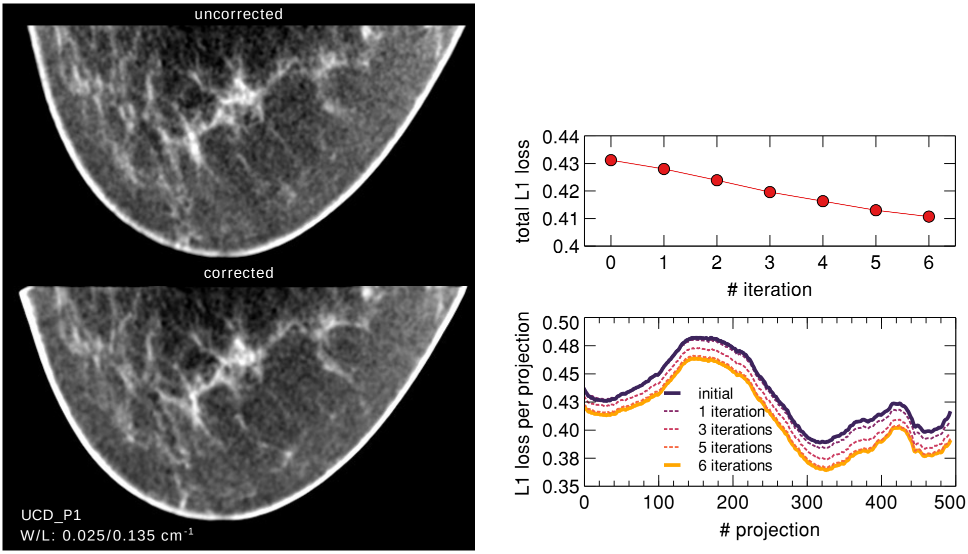

(a) (b)

Figure 5.12: (a) Application of motion-artifact reduction to clinical scan UCD_P1. (b) Recorded loss as function of iteration number and loss per projection when applying the method to clinical case UCD_P1.

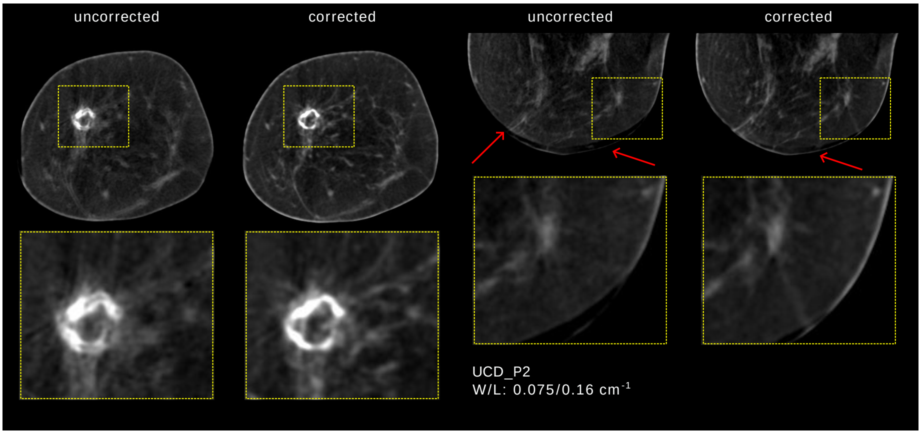

Figure 5.13: Application of empirical motion-artifact reduction to clinical scan UCD_P2. There is clear improvement in the side view where the locations of the blood vessels remain unchanged but the tissue is better attached to the skin towards the nipple as indicated by red arrows.

sudden vertical motion towards the end of the scan that resulted in "detached" skin near the nipple that can be seen in the axial view highlighted by red arrows. Our method identified the location of this artifact but could not fully resolve it in the anterior region of the breast near the nipple where this detachment is the largest. The sudden vertical motion is the likely cause of the double appearance of a calcification ring shown in the axial view. Our method restores the shape of the ring and surrounding fibro-glandular tissue.

RUNTIME ANALYSIS

We performed the detailed runtime analysis of patient UCD_P1. This acquisition consisted of 492 projections of 924×768 pixels, and the reconstructed volume size was $300 \times 500 \times 500$ voxels. Sparse b-spline field describing the image transformation at each projection angle had the following shape: $3 \times 27 \times 44 \times 44$ - the sparse coefficients were placed approximately every 11 voxels in each direction. The overview of the time required for each operations is presented in table 5.3. In each iteration, the estimation of b-spline field took around 400 seconds and the image update took 255 seconds with non-optimized research-level code: all 492 projections were processed sequentially on a Nvidia A6000 GPU without any operations performed on the CPU. When estimating the b-spline fields, the FP operation required in total less than 0.5 seconds, while combined time to densify the b-spline fields, apply scaling and squaring, and deform the image took 177 seconds. The error backpropagation took 220 seconds, which included the BP operation and error backpropagation through scaling and squaring algorithm and 3D convolution of sparse b-spline fields. The remaining time was spent mainly on algebraic operations and memory manipulations. In the image update phase, the BP operation took in total 5 seconds. The time needed for image transformation remained unchanged.

Table 5.3: Time required to perform a single iteration of the method on UCD_P1 case.

Component	Time (s)
FP (τ_{FP})	0.5
BP (τ_{BP})	5
Transformation (τ_{trafo})	177
Backpropagation transformation (τ_{trafo}^{backpr})	220
Total time: b-spline fields estimation	400
Total time: image update	225

5.5 Discussion

We developed a data-driven non-rigid empirical motion-artifact reduction algorithm for dedicated breast CT that can potentially also be applied to any type of cone-beam CT. Our method is model-free and, therefore, does not require knowing the type or amplitude of the motion in order to fit its parameters during motion-artifact reduction. Unlike deep learning-based motion compensation methods, we also do not require large datasets of cases with known motion for training. Therefore, the presented method can be applied to novel imaging modalities for which the size of existing datasets do not allow training of neural networks. We demonstrated the independence from both acquisition parameter and motion type by applying the developed method to three different implementations that use different x-ray spectra and dose levels as well as to a digital phantom [36]. In all cases, the method was able to correctly recognize and significantly reduce motion artifacts.

We verified the correctness of this method by performing a simulation study with known motion pattern. The comparison of the shape of fibro-glandular tissue in Fig. 5.4 suggests that the method is able to correctly identify artifacts due to non-rigid patient motion and remove them while analysis of the estimated b-spline fields in Fig. 5.6 and the difference image in Fig. 5.4 indicates that the developed method does not need to recover the ground truth motion in order to do so.

The evolution of the reconstructed volume seen in Fig. 5.3b shows that we do need to select a reference position instead of using the average position found by the algorithm. Practically we used the first position even though this does not correspond to the starting position during the scan. The reason to select this reference position is that the b-spline fields that we find are not guaranteed to preserve volume. The existence of the inverse only guarantees that the Jacobian of the local transformation at a specific point is non-zero while volume preservation would require the Jacobian to be equal to 1. At the chosen reference position there is the additional constraint that the projection of the volume needs to match the measured data, effectively adding a constraint to the volume that does not exist at the average position.

The versatility of our approach to variations in motion type are evident when examining scans of physical phantoms that exhibited motion, whether it was rigid or deformable in nature. Fig. 5.8 and 5.9 show the reduction of motion artifacts. The outer skin line and overall shape of the objects could always be successfully restored. Moreover, the ghosting of inner structures disappeared and a large part of the objects inside the phantoms took their original form. We achieved independence on motion type by using a very general combination of b-spline representation combined with diffeomorphic transformation. Although it puts some limitation on the type of motion for which artifacts can be reduced for, it most probably fully covers anatomically-feasible motion.

Fig. 5.9 and 5.13 suggest that although the method can reduce artifacts with great success, major motion artifacts cannot be completely eliminated. For example, in the UCD_P2 case a residual artifact is still present towards the nipple where the skin is still not attached to the rest of the breast tissue. The performance of our method on the CIRS 073 phantom is satisfactory at best. We explain it by the fact that the method is unlikely to arrive at the global minimum when performing the optimization. In other words, the method cannot recover the true object form and motion field that is applied to the scanned item in the object space. Consequently, not all parts of the reconstructed image become equally artifact-free as can be seen in Fig. 5.9, where the objects only in the left part of image took the same shape as in the corresponding static scan. The features on the right are the average between all acquired discrete realizations of the phantom.

Another limitation of our method that makes its introduction into clinical use difficult is the long computational time and high GPU memory usage. We believe that the runtime can be improved in two ways. First, due to the linear run time dependency on the number of angles, multiple angles could be processed in parallel in both phases. Depending on the available GPU memory, the runtime could be at least halved in this manner. Second, the runtime of the most time consuming operation – densification of sparse b-spline fields – can be decreased further by adjusting the spatial spacing between the elements so that 3D convolution is faster without negatively affecting the performance. Third, the scaling factor in the scaling and squaring method can be reduced on the fly if the resulting image transformation is expected to be diffeomorphic. Smaller scaling factor will reduce the time required to densify the parameters and backpropagate the error through this operation.

In our opinion, future research needs to address issues with convergence by, e.g., performing both phases of the method using cleverly selected subsets of projections as it might lead to better results when motion amplitude is large or when high-contrast objects are present in the tissue. Execution of the method with subset processing could also be combined with motion detection based either on external sensors or camera data or by checking that projections are consistent. When motion is detected using a dedicated method, subsets could be selected based on how well they cover the timepoint when motion has happened. Given the limited performance observed in regions with the largest motion amplitude, starting with a rigid pre-correction that can then be further refined by the presented method and would likely help both processing quality and speed. Another important aspect to be investigated is the selection of the stopping criterion. The L1 loss used in this work was chosen empirically, and although it leads to good results on our datasets, it is not guaranteed to be the best since we lack mathematical proof. Moreover, the selection of the cost function could depend on anatomy and, hence require a different stopping criterion. For example, in the presence of highly-attenuating objects such as bones, in our opinion an additional term in the cost function accounting for edges could be advantageous.

Finally, we only presented a technical and visual evaluation of our method and the question whether the improvements we have shown translate to a clinical benefit remains to be answered. For example, the doubled appearance of the calcification ring in Fig.

5.13 coupled with visible skin detachment towards the nipple would be a rather obvious indication of motion artifacts for experienced clinicians that is unlikely to result in misinterpretation of the image. The same holds true for the region near the nipple in the corrected image where the skin is still detached. However, despite the limitations, our method now allows us to perform motion-artifact reduction in clinical breast CT scans when needed and thus creating the data that will allow us to demonstrate potential clinical benefit and to train deep-learning based motion compensation methods if desired. Moreover, perfusion imaging that relies on artifact-free prior images is now more likely to succeed since the alignment of artifact-free images is in general an easier task.

5.6 CONCLUSION

We presented an empirical motion-artifact reduction algorithm that follows joint model-free motion estimation and image update scheme for cone beam breast CT. We demonstrated in simulations, phantom studies, and clinical cases that this method is able to reduce major motion artifacts and substantially improve image quality in all examined cases. However, there are still remaining residual artifacts due to severe motion present in some cases; therefore, improvement in the performance on such cases still needs to be addressed.

DECLARATION OF COMPETING INTEREST

Ioannis Sechopoulos has research agreements with Siemens Healthcare, Canon Medical, ScreenPoint Medical, Sectra Benelux, Volpara Healthcare, Lunit, iCAD, and E-COM, and speaker agreements with Siemens Healthcare and Canon Medical and is a Scientific Advisory Board member of Koning Corp.

Andrew M. Hernandez has research agreements with Canon Medical and is a consultant for Izotropic Corp.

John M. Boone has research agreements with Canon Medical and is a founder and shareholder for Izotropic Corp.

Sjoerd A.M. Tunissen is a shareholder for Izotropic Corp.

DATA AVAILABILITY

The authors do not have permission to share patient data. The phantom data can be shared upon reasonable request.

ACKNOWLEDGMENT

We thank Chris de Korte and Jan Menssen for helping us with the transitional stages and providing the CIRS 073 phantom.

REFERENCES

[1] M. Mikerov, K. Michielsen, N. Moriakov, J. J. Pautasso, S. A.M. Tunissen, A. M. Hernandez, J. M. Boone, and I. Sechopoulos. Empirical motion-artifact reduction for non-rigid motion in dedicated breast ct. *IEEE Transactions on Biomedical Engineering*, 72(10):3133–3145, 2025.

[2] J. Ferlay, I. Soerjomataram, R. Dikshit, S. Eser, C. Mathers, M. Rebelo, D. M. Parkin, D. Forman, and F. Bray. Cancer incidence and mortality worldwide: Sources, methods and major patterns in GLOBOCAN 2012: Globocan 2012. *International Journal of Cancer*, 136(5):E359–E386, March 2015.

[3] H. Sung, J. Ferlay, R. L. Siegel, M. Laversanne, I. Soerjomataram, A. Jemal, and F. Bray. Global Cancer Statistics 2020: GLOBOCAN Estimates of Incidence and Mortality Worldwide for 36 Cancers in 185 Countries. *CA: A Cancer Journal for Clinicians*, 71(3):209–249, May 2021.

[4] Y. Zhu, A. M. O'Connell, Y. Ma, A. Liu, H. Li, Y. Zhang, X. Zhang, and Z. Ye. Dedicated breast ct: state of the art—part i. historical evolution and technical aspects. *European Radiology*, pages 1–11, 2022.

[5] P. M. Gazi, K. Yang, G. W. Burkett Jr., S. Aminololama-Shakeri, J. Anthony Seibert, and J. M. Boone. Evolution of spatial resolution in breast ct at uc davis. *Medical Physics*, 42(4):1973–1981, 2015.

[6] A. M. Hernandez, A. E. Becker, S. H. Lyu, C. K. Abbey, and J. M. Boone. High-resolution μCT imaging for characterizing microcalcification detection performance in breast CT. *Journal of Medical Imaging*, 8(5):052107, 2021.

[7] M. Mikerov, K. Michielsen, N. Moriakov, and I. Sechopoulos. Adding patient motion from DCE-MRI to anthropomorphic phantoms for dedicated breast CT. In Ivana Išgum and Olivier Colliot, editors, *Medical Imaging 2022: Image Processing*, page 73, San Diego, United States, April 2022. SPIE.

[8] A. Melbourne, J. Hipwell, M. Modat, T. Mertzanidou, H. Huisman, S. Ourselin, and D. J. Hawkes. The effect of motion correction on pharmacokinetic parameter estimation in dynamic-contrast-enhanced MRI. *Physics in Medicine and Biology*, 56(24):7693–7708, December 2011.

[9] L. Brombal, L. M. Arana Peña, F. Arfelli, R. Longo, F. Brun, A. Contillo, F. Di Lillo, G. Tromba, V. Di Trapani, S. Donato, R. H. Menk, and L. Rigon. Motion artifacts assessment and correction using optical tracking in synchrotron radiation breast CT. *Medical Physics*, 48(9):5343–5355, September 2021.

[10] T. E. Marchant, G. J. Price, B. J. Matuszewski, and C. J. Moore. Reduction of motion artefacts in on-board cone beam ct by warping of projection images. *The British Journal of Radiology*, 84(999):251–264, 2011. PMID: 21081580.

[11] A. Z. Kyme and R. R. Fulton. Motion estimation and correction in spect, pet and ct. *Physics in Medicine & Biology*, 66(18):18TR02, sep 2021.

[12] S. Capostagno, A. Sisniega, J. W. Stayman, T. Ehtiati, C. R. Weiss, and J. H. Siewerdsen. Deformable motion compensation for interventional cone-beam ct. *Physics in Medicine & Biology*, 66(5):055010, feb 2021.

[13] H. Huang, J. H. Siewerdsen, W. Zbijewski, C. R. Weiss, M. Unberath, T. Ehtiati, and A. Sisniega. Reference-free learning-based similarity metric for motion compensation in cone-beam ct. *Physics in Medicine & Biology*, 67(12):125020, jun 2022.

[14] T. Sun, J. Kim, R. Fulton, and J. Nuyts. An iterative projection-based motion estimation and compensation scheme for head x-ray ct. *Medical Physics*, 43(10):5705–5716, 2016.

[15] B. Flach, M. Brehm, S. Sawall, and M. Kachelrieß. Deformable 3d–2d registration for ct and its application to low dose tomographic fluoroscopy. *Physics in Medicine & Biology*, 59(24):7865, nov 2014.

[16] H. Huang, J. Siewerdsen, A. Lu, Y. Hu, W. Zbijewski, M. Unberath, C. Weiss, and A. Sisniega. Multi-stage adaptive spline autofocus (masa) with a learned metric for deformable motion compensation in interventional cone-beam ct. volume 12463, page 42, 04 2023.

[17] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache. Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage*, 45(1, Supplement 1):S61–S72, 2009. Mathematics in Brain Imaging.

[18] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS 2017 Workshop on Autodiff*, 2017.

[19] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[20] J. Wang and X. Gu. Simultaneous motion estimation and image reconstruction (smeir) for 4d cone-beam ct. *Medical Physics*, 40(10):101912, 2013.

[21] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.

[22] N. Moriakov, J. Sonke, and J. Teuwen. End-to-end memory-efficient reconstruction for cone beam ct. *Medical Physics*, n/a(n/a).

[23] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu. Spatial transformer networks, 2016.

[24] N. J. Higham. The scaling and squaring method for the matrix exponential revisited. *SIAM Journal on Matrix Analysis and Applications*, 26(4):1179–1193, 2005.

[25] A. V. Dalca, G. Balakrishnan, J. Guttag, and M. R. Sabuncu. Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical Image Analysis*, 57:226–236, 2019.

[26] Hsieh H. and H. Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26(6):508–517, 1978.

[27] R. Sandkühler, C. Jud, S. Andermatt, and P. C. Cattin. Airlab: Autograd image registration laboratory, 2020.

[28] M. Caballo, R. Mann, and I. Sechopoulos. Patient-based 4d digital breast phantom for perfusion contrast-enhanced breast ct imaging. *Medical Physics*, 45(10):4448–4460, 2018.

[29] R. Beare, B. Lowekamp, and Z. Yaniv. Image segmentation, registration and characterization in r with simpleitk. *Journal of Statistical Software*, 86(8):1–35, 2018.

[30] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens. Reduction of metal streak artifacts in x-ray computed tomography using a transmission maximum a posteriori algorithm. *IEEE Transactions on Nuclear Science*, 47(3):977–981, 2000.

[31] Multi-modality breast biopsy and sonographic trainer. `https://www.sunnuclear.com/uploads/documents/datasheets/Diagnostic/073_datasheet_072622.pdf`. Accessed: 2023-01-12.

[32] A. M. Hernandez, C. K. Abbey, P. Ghazi, G. Burkett, and J. M. Boone. Effects of kv, filtration, dose, and object size on soft tissue and iodine contrast in dedicated breast ct. *Medical Physics*, 47(7):2869–2880, 2020.

[33] G. T. Herman. Correction for beam hardening in computed tomography. *Physics in Medicine & Biology*, 24(1):81, jan 1979.

[34] H. Johnson, G. Harris, and K. Williams. Brainsfit: Mutual information registrations of whole-brain 3d images, using the insight toolkit. 10 2007.

[35] A. Fedorov, R. Beichel, J. Kalpathy-Cramer, J. Finet, J. Fillion-Robin, S. Pujol, C. Bauer, D. Jennings, F. Fennessy, M. Sonka, J. Buatti, S. Aylward, J. V. Miller, S. Pieper, and R. Kikinis. 3d slicer as an image computing platform for the quantitative imaging network. *Magnetic Resonance Imaging*, 30(9):1323–1341, 2012. Quantitative Imaging in Cancer.

[36] M. Mikerov, K. Michielsen, N. Moriakov, and I. Sechopoulos. Non-rigid motion compensation for breast ct. In *2022 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–2, 2022.

6

APPROXIMATE VS. EXACT: THE (UN)EXPECTED ADVANTAGE OF APPROXIMATIONS

Functional 4D CT imaging including contrast injection can be beneficial when evaluating pathology. However, acquisition of long image sequences implies the need to use a low dose per phase to limit total radiation dose. This calls for development of advanced processing methods that can handle the resulting high image noise efficiently. Noise filters specific for dynamic imaging can use temporal information to achieve higher noise reduction but have a high computational cost. In this work, we accelerated an existing clinical 4D noise filter by making changes to its algorithm that enabled it to run approximately 13 times faster on GPU compared to a highly optimized CPU implementation without any obvious reduction in its effectiveness.

This chapter appeared as
M. Mikerov, S.A.M. Tunissen, K. Michielsen, E.J. Smit, and I. Sechopoulos.
Adaptation of a 4D noise filter for implementation on GPU
8th International Conference on Image Formation in X-Ray Computed Tomography [1].

6.1 INTRODUCTION

Despite significant progress in CT hardware and software that improved image quality and clinical decision making, standard CT provides only anatomical information of the patient condition. In many clinical situations it would be beneficial to also obtain functional information by studying the enhancement with iodinated contrast agents, especially when evaluating pathology for treatment planning or response.

Usually three- or four-phase functional imaging protocols are implemented in the clinic for patients with stroke, cardiac ischemia, or liver lesions. They can still be performed at relatively low doses but have the risk of mistimed acquisitions, resulting in missing important aspects of contrast dynamics. More densely sampled sequences would eliminate the risk of missing important points on perfusion curves.

Functional imaging with long sequences consisting of 10 to 30 phases depending on the protocol comes at the cost of an increased dose. To reduce this, each phase in the sequence needs to be acquired at considerably lower dose than a standard anatomical CT image. Noise reduction filters can then be applied as a postprocessing step to enable the extraction of useful information from otherwise too-noisy images.

Traditional denoising filters can be used in each individual phase, however they often use information from a small spatial neighborhood, leading to loss of resolution. To solve this, information from the other images in the sequence can also be used to decrease noise in the current phase. Therefore noise filters specific for dynamic imaging that leverage the temporal information such as the 4D similarity filter have been proposed [2].

In this work, we aim to improve runtime of the 4D similarity filter by introducing a small modification to the algorithm that enables its implementation on GPU. We performed the first comparison with the exact CPU version and did not notice a performance drop.

6.2 FILTER IMPLEMENTATION

6.2.1 4D SIMILARITY FILTER

Prokop and Smit introduced a 4D similarity filter to reduce noise in low dose 4D imaging [2]. The idea of this filter is to find, for each voxel v in the image at a given timepoint t, a set of voxels $\{V_c\}$ that have similar perfusion curves and use those with the least differences to calculate an average voxel value that can then be used to update the image.

The search for similar voxels is performed by comparing the perfusion curves of the voxel of interest v (the one being updated) with a large pool of candidate voxels $v_c \in \{V_c\}$ from the entire image according to how close the temporal mean of each voxel v_c is to that of v. If the perfusion curve of a candidate voxel v_c is similar according to some threshold TH, it is saved in an array A. When the array is full, it is sorted based on the difference of perfusion curves and only the best FS candidate voxels are used to calculate the average value for this voxel. This filter can be implemented both on CPU and GPU, but the latter has not been achieved before. There exist multi-threaded CPU realizations of this filter, which are already used in clinical practice [3–5].

One of the drawbacks of the 4D similarity filter is long computational times since every voxel in the image needs to be processed independently. In a typical 320x512x512 volume, with just 1/3 of all voxels being in the mask, the total number of voxels is around 27 million. On a 32-core CPU, each core would need to process more than 870 thousand voxels. The independence of each voxel invites use of the GPU for this task, since it has many more cores, which can process even more voxels in parallel. However, we encountered major problems when implementing the 4D similarity filter on the GPU.

6.2.2 Difficulties of implementation on GPU

In general, low computational times on the GPU are achieved when it is possible to execute one thread per output element that can run independently from other threads and consists mainly of arithmetic operations with very limited memory accesses. This is due to the fact that arithmetic operations need up to 20 times fewer cycles to complete than operations on device (global) memory. If a high number of memory accesses are required to compute a single output element, several techniques can be employed. First, one wants to ensure that memory reads are coalesced, i.e., during the calculation of neighboring output elements, memory reads from the threads in the same block should be performed in accordance to memory layout. Second, if multiple threads in the same block can reuse some of the information, shared memory can be used to specify a user-defined and -maintained cache. Third, if memory reads are random but there is a high chance of having close spatial proximity, textures can be utilized due to their large cache, which is checked first before reading the data from the (slow) device memory.

Unfortunately, none of these optimization techniques are applicable in our problem since each candidate voxel can be located in a completely different part of the image, therefore making coalesced memory reads impossible. Moreover, the probability of hitting the same element is very low, which removes the advantages of using the fast texture cache. Shared memory cannot be used either since it is not known at execution time whether a candidate voxel in one of the threads will be reused in another thread within the same block and, if so, in which one specifically. Thus, it is not possible to save this value in a structured cache for future use.

Another major limiting factor is the need to sort the selected similar voxels in array A according to the difference in perfusion curves. This is difficult due to two reasons. To begin with, keeping track of sometimes thousands of data pairs (difference and voxel value) requires sufficient memory that quickly exceeds the amount of fast on-chip shared memory. Moreover, since shared memory needs to be shared between all threads in a thread block, it greatly reduces the number of voxels that can be processed in parallel. Thus there is no other choice left other than accept writes into slow device memory. The fastest sorting algorithms require on average $\mathcal{O}(n \log n)$ comparisons, which can become a large number for typical settings of the filter. From the computational point of view this method becomes inefficient on the GPU, since memory reads and writes are done in device memory and if not all threads in a thread block (*warp*) have finished their sorting, the core cannot be released to process another block of threads.

A very straightforward approach to reduce the computational burden due to sorting at the end is to sort each element at insertion by using, e.g., min-heap structure of size FS. With this approach, one does still need to keep the differences and values of FS elements. However, FS is almost always smaller than the size of A and it can be possible to store all FS elements in fast memory. Theoretically, on average the insertion takes again $\mathcal{O}(n \log n)$ comparisons, but now the memory throughput is very high since we are working with fast shared memory only. In our case, this number will be even lower since the min-heap

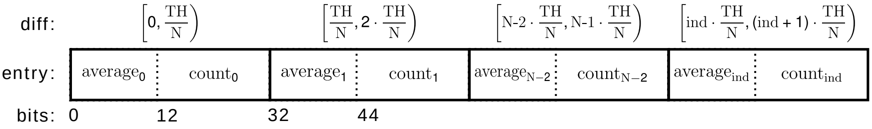

Figure 6.1: A sorted array structure is used to store the average value of voxels within a specific range of differences, along with the corresponding voxel count required for this average calculation. The insertion and extraction operations have a worst-case performance of $\mathcal{O}(1)$.

structure is sorted before each insertion. Nevertheless, the insertion must be performed as many times as the number of elements we wanted to have in our original array A. Furthermore, element insertion into a heap has multiple if-statements, which increases the likelihood of branching and therefore they could lead to the cores being idle more often.

6.2.3 OUR SOLUTION

Based on these algorithmic bottlenecks hindering an efficient implementation on GPU, we propose a simple change to the algorithm that enables us to work around these issues and attain fast computational times on the GPU at the cost of a small reduction in the numerical accuracy of the solution.

Figure 6.1 schematically represents our way of storing the same number of elements as in the original A with limited space requirements and a fixed number of operations ($\mathcal{O}(1)$) on insertion and reading in the worst-case scenario.

As in the original version, a candidate voxel v_c is considered for future processing only if the difference between its and voxel v's perfusion curves is below a given threshold TH. By carefully selecting TH, it is possible to save all good candidate values in a sorted manner with limited memory resources. Since the end-result of each kernel execution is an average of FS best candidate voxels v_c, it makes sense to store the average and corresponding count utilizing as little memory space as possible. Assuming that a typical CT image is in the range from -1024 to 3071 HU, storing each value using only positive numbers requires just 12 bits. In case of 32-bit integers it leaves us with 20 bits that can be used to store metadata corresponding to this number. For us, the only metadata we have for the average value is the number of voxels that were used to calculate it. Now we only need to store it in a sorted manner in a small data structure. We propose to create an array H of size N where each n-th bin stores the average value of the candidate voxels v_c with the differences between $n \cdot TH/N$ and $(n+1) \cdot TH/N$. For example, if $N = 100$, $TH = 2000$, the first bin is responsible for storing the average value of candidate voxels with the difference between 0 and 19. The size of the array N is a trade-off between the intended bin width for a certain TH and the number of threads that can run in parallel. For example, for the Nvidia A6000 GPU with 99 KB of shared memory, one can theoretically have 193 parallel threads with $N = 128$ or just 96 threads with $N = 256$ in each block.

6.2.4 LIMITATIONS

The changes introduced to the algorithm have a negative effect on the numerical accuracy of the 4D similarity filter, which is important to understand and to explore possible solutions. To start, some accuracy is lost when calculating the average due to rounding, since only integer values can be stored. Theoretically, one could use fixed point number arithmetic to mitigate this problem, if one can accept that fewer bits will be used to represent the integer part of the number or the maximum count will be smaller. Following that, the width of the bin could be too broad so that almost all of FS best voxels can be found in one of the first bins. If the number of candidate voxels in one of the bins near the start of the array H greatly exceeds FS at some locations in the image, the noise reduction will not be homogeneous across the whole image. When it is not possible to increase the number of bins anymore due to memory constraints, one could move the array H into device memory, at the cost of increasing memory latency. Alternative solution could be to use the array H with non-uniform bin width. This way, the first bins can be made smaller to enable more granular summation of FS.

6.3 MATERIALS AND METHODS

We compared the performance of our and commercial CPU implementations of the 4D similarity filter on simulated CT images of an anthropomorphic XCAT perfusion phantom. The acquisitions were simulated in low-dose mode ($\text{CTDI}_{\text{vol}} = 1.3$ mGy per phase) using a previously developed and validated CT simulator [6]. The images were reconstructed using an own implementation of the standard FDK algorithm. To avoid negative values and allow the custom array structure we applied the following linear scaling to arbitrary units: $\text{AU} = (\text{HU} + 3000)/6$. n this preliminary work, we performed very simple analysis that included quantification of noise reduction in terms of signal-to-noise ratio (SNR) in a homogeneous region of the liver and estimation of the bias.

The GPU implementation has fewer and slightly different parameters than the CPU version. Even though we adjusted the parameters according to our understanding of the filter in order to achieve comparable performance, we note that both filters are not expected to lead to identical results. The size of array H was 128 bins to have a multiple of 32 to avoid bank conflicts when writing and reading from H. Mean squared error (MSE) TH was 12800, leading to each bin being 100 MSE units wide. FS was set to 100. For each voxel in the mask, 40 000 candidate voxels were checked.

6.4 RESULTS

Figure 6.2 shows a slice of the original reconstructed phantom and two slices filtered using the CPU and GPU versions of the algorithm. Both results look very similar in terms of representing anatomical structures.

We calculated the SNR in the entire liver as we had perfect segmentation from the XCAT phantom. The GPU version has achieved slightly higher noise reduction than the GPU version, (SNR increased by a factor of 7.3 and 7.7, respectively). Since the noise is expected to decrease mainly as a function of the number of averaged voxels FS provided that they are sampled from the same distribution, and unlike in the CPU version, FS is not fixed, we calculated the actually used FS for each voxel in the output. Figure 6.3a shows the actually used FS to calculate the average for each voxel in the image. In the whole liver, we found FS to be 126 ± 19 which is higher than the intended $FS = 100$ used in the CPU version.

No visible bias is observed when comparing images filtered on CPU and GPU. To verify

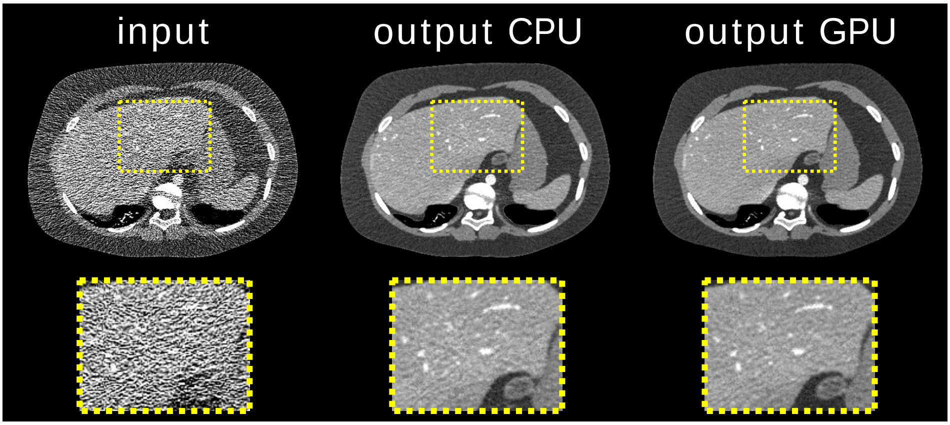

Figure 6.2: Comparison of exact CPU and approximate GPU 4D similarity filters. W/L: 200/50 HU.

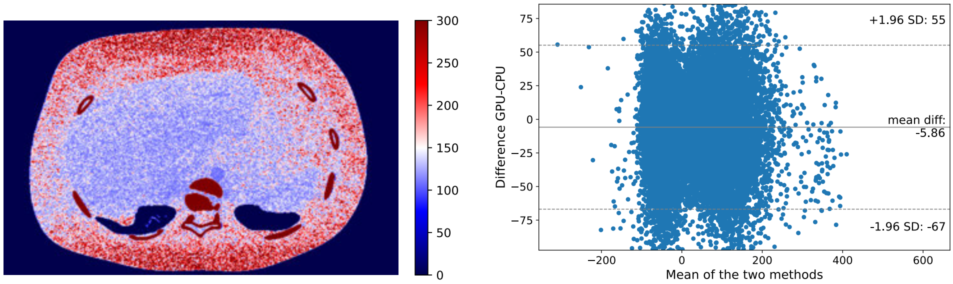

(a) (b)

Figure 6.3: (a) Actually used *FS* for each voxel measured as a number of voxels used to calculate the average. The values are clipped at 300. (b) Bland-Altman plot comparing GPU and CPU versions of the filter. The output of the GPU version is on average lower.

this we constructed the Bland-Altman plot shown in Figure 6.3b. There is a small negative bias of -6 HU. There appears to be no clear dependency of the difference between two versions of the filter and voxel grayscale value. Moreover, we observed that both the CPU and GPU versions introduce either positive or negative bias depending on the location in the liver or other organs. There is no clear dependency on anatomy, functional or image properties.

The GPU result can be obtained within 44 seconds on an Nvidia A6000 GPU, which is on average 13 times faster than the multi-threaded commercial CPU implementation that is executed on 32 cores, based on the mean of three runs. The Nvidia profiler has also shown that the L1 and L2 caches are not fully utilized. In order to improve reading from device memory, we implemented memory prefetching – eight candidate voxels are loaded from the device memory into registers and are then processed. This modification improved the utilization of the caches. Better memory utilization increased compute throughput by

385% and enabled, thus, the execution of the filter within 27 seconds, which is roughly 23 times faster than the CPU version.

6.5 Discussion

We modified the original algorithm of the 4D similarity filter in order to avoid GPU-specific computational constraints. With this we were able to run it on GPU several times faster than with highly-optimized multi-threaded code on the CPU. The main change was replacing the sorting operation by using an ordered hash-like structure that can reside in fast but limited on-chip shared memory and thereby enabling $\mathcal{O}(1)$ writing and reading. However, we needed to sacrifice accuracy to use this data structure. We performed first tests with both versions of the similarity filter to study the effect of this accuracy loss.

Overall, both versions produce very similar outputs with comparable set of parameters. The main difference is a higher level of denoising than anticipated, which is explained by the fact that FS is not a fixed parameter anymore. Therefore it leads to different level of noise reduction in various parts of the image, as demonstrated. To counteract this effect, the bin width in the hash H could be chosen smaller, therefore enabling more granular division of candidate voxels and thus increasing the chance of using the sought number of voxels FS for averaging.

An important advantage of our implementation is significantly reduced complexity of the code. This simplification will enable us to add modifications to the algorithm and explore additional features such as looking for neighbors in a local region or using different metrics to select the voxels. Furthermore, our implementation enables the filtering of single phase in the sequence. An interesting scenario to investigate would be the application of the filter to reduce the noise in every second or third image in the sequence for calculation of approximate perfusion maps in clinical situations in which the time of diagnosis is the main concern, while more images are being processed in the background to update the perfusion maps with more accurate data.

In this work we performed an initial evaluation of how the filter is affected by the algorithm changes in a digital phantom. At the moment it is not clear whether the introduction of slightly non-uniform noise reduction throughout the image and different anatomies we found will have negative effects on the performance of radiologists or further processing of the images, e.g., to calculate perfusion maps. Based on these initial results we plan to continue with further analysis of the performance by, e.g., quantifying bias of the perfusion curves in different anatomies. Such an analysis combined with hyperparameter tuning must be done in order to ensure correct and trustworthy performance.

6.6 Conclusion and outlook

We implemented a GPU version of the 4D similarity filter that is on average approximately 13 times faster than a highly optimized commercial multi-threaded CPU realization by introducing a change into the original algorithm to make it more GPU-friendly at the cost of lowering accuracy. Our modification to the algorithm removed a significant bottleneck for GPU implementation. Our further non algorithm-specific modification enabled further speedup by factor of 1.7 compared to the already modified algorithm.

REFERENCES

[1] M. Mikerov, S.A.M Tunissen, K. Michielsen, Ewoud J.S., and I. Sechopoulos. Adaptation of a 4d noise filter for implementation on gpu. In *Proceedings of the The 8th International Conference on Image Formation in X-Ray Computed Tomography*, 2024.

[2] E. J. Smit and W. M. Prokop. Noise reduction in image data, WO 2017/194787 A1, 2017.

[3] T. Kouchi, Y. Tanabe, E. J. Smit, T. Kido, A. Kurata, Y. Kouchi, H. Nishiyama, T. Uetani, S. Ikeda, O. Yamaguchi, M. Prokop, and T. Mochizuki. Clinical application of four-dimensional noise reduction filtering with a similarity algorithm in dynamic myocardial computed tomography perfusion imaging. *The International Journal of Cardiovascular Imaging*, 36:1781–1789, 2020.

[4] O. Sliwicka, Z. Swiderska-Chadaj, M. Snoeren, M. Brink, K. Salah, L. Peters-Bax, T. Stille, M. J. van Amerongen, I. Sechopoulos, and J. Habets. Multireader image quality evaluation of dynamic myocardial computed tomography perfusion imaging with a novel 4d noise reduction filter. *Acta Radiologica*, 64(3):999–1006, 2023.

[5] K. Murayama, E. J. Smit, M. Prokop, Y. Ikeda, K. Fukui, I. Nakahara, S. Hanamatsu, K. Katada, Y. Ohno, and H. Toyama. A bayesian estimation method for cerebral blood flow measurement by area-detector ct perfusion imaging. *Neuroradiology*, 65:65–75, 2022.

[6] S. A. M. Tunissen, L. J. Oostveen, N. Moriakov, J. Teuwen, K. Michielsen, E. J. Smit, and I. Sechopoulos. Development, validation, and simplification of a scanner-specific ct simulator. *Medical Physics*.

7

GENERAL DISCUSSION

The work I presented in this thesis lays the technical foundation from the algorithmic point of view for application of dynamic contrast-enhanced breast CT to personalize breast cancer treatment and monitor the treatment response to neoadjuvant therapy. These foundational developments were tested and validated using both computer simulations and imaging of physical phantoms, with the expectation that they will be tested in humans once patient testing of this new modality begins. Therefore, I can only speculate about the potential impact of the developed imaging modality on breast cancer treatment and will do so at the end of this chapter. There I will explain in more detail why the developed modality can be used for treatment planning and monitoring and what future work needs to be performed in order to make the research results in this thesis applicable to the clinic. In the meantime, however, I will recap the key findings and discuss the influence of the developed methods on low dose DCE CT imaging not necessarily limited to breast CT.

I will start by examining the applicability of Midgley's decomposition in iterative model-based image reconstruction. In **chapter 2**, we demonstrated in a theoretical study with restrictive assumptions that Midgley's parameterization into energy-dependent S-curves and material specific a_i-coefficients[1] is the most accurate decomposition of the linear attenuation coefficient in the low photon energy range for typical breast tissues. In accordance with the initial research by Midgley [1], I am confident that it will remain the most accurate decomposition that separates dependencies on energy and atomic numbers also for other soft tissues and bones, as long as they do not contain materials with a K-edge in the relevant energy range. I identify two directions on how this decomposition can be used in X-ray imaging.

On the one hand, this parameterization can be used in simulations of tissue variability across different regions of the organ. Since the linear attenuation coefficient in each voxel is given by five parameters, only five forward projections are required to calculate the signal on the detector, independent of the number of variations in tissue composition present in the organ of interest. This is more efficient than the classical approach of

---

[1] In chapter two, we combined electron density N_e and material specific R_i-coefficients into a_i-coefficients to reduce the number of unknowns from six to five.

representing certain tissues in each voxel of the image by discrete values, since a very large number of forward projections is required to calculate corresponding thickness maps of each material if multiple materials are present. Researchers in my group have already successfully applied Midgley's decomposition to simulate tissue variability for applications in digital breast tomosynthesis [2]. Even when tissue variability is not a requirement but multiple organs are used in simulations, e.g., in full-body CT, Midgley's parameterization will already provide a computational advantage. This parameterization is also superior compared to decomposition into the effects of photoelectric absorption and Compton scatter on attenuation of X rays, since there are constraints on the selection of its parameters. a-parameters in this parameterization are defined in a such way that $a_{k+1} \geq a_k$ and the maximum possible a-coefficient corresponds to the largest atomic number in the modeled tissue. Therefore, the anatomical feasibility can be easily ensured. Furthermore, the use of different parameterization methods in simulation and reconstruction is advantageous when developing new reconstruction algorithms to avoid the exact match between the models used to simulate projection data and reconstruct it.

On the other hand, the incorporation of this parameterization into iterative reconstruction of soft tissue without iodine in double-energy setups has shown promising initial results. This might sound counter-intuitive since the problem to be solved is underdetermined - five unknowns must be determined from two sets of measurement data. The explanation why this decomposition still works in reconstruction is the fact that there is more than a single unique way to achieve the same polychromatic linear attenuation coefficient, which in combination with the spectrum used, leads to the right value at the detector. Since the linear attenuation coefficient at some energy rather than the exact R-parameters is sought, this decomposition leads to the correct result. From the technical point of view, the easiest way to apply this parameterization is to use automatic differentiation and gradient descent optimization to minimize the data discrepancy between the forward projected and measured data. The use of automatic differentiation will also permit incorporation of advanced regularization priors, therefore leading to very accurate polychromatic reconstruction with desired visual properties in the image domain.

Although this model can be used to reconstruct images of organs not containing K-edge materials, it fails to do so when such materials are present. In this case, the fact that the problem is underdetermined makes accurate estimation of the concentration of the material with the K-edge not possible. Even if the concentration of such a material is incorrectly estimated, the remaining five parameters can still compensate for it leading to matching projected and measured data.

The inability to accurately estimate low iodine concentration in breast imaging using iterative model-based methods and the fact that it would have been possible only at very high patient dose resulted in us investigating other methods to reconstruct post-contrast images. In **chapter 3**, we optimized the image acquisition and reconstruction settings using the well-known PICCS method. Using PICCS we were able to significantly reduce the radiation dose required. When I look back at the original problem, I firmly believe that this was the right decision, not only due to reduced patient dose, but also due to reduced heating of the X-ray tube in the current realization of breast CT. When we had started working on this project, we underestimated the latter.

The main technical novelty of our PICCS implementation is the use of the IMPACT

- single-energy polychromatic reconstruction method. Although it is a computationally expensive method, its polychromatic nature enables numerically accurate reconstructions without beam hardening artifacts such as cupping. It opens the doors to intrapatient longitudinal comparisons even if the studied organ exhibits large variability in size and shape. It becomes especially important in breast imaging, since the breasts vary largely not only in size, but also in shape and inner structure. These effects cannot be accounted for efficiently with a simple single-material beam-hardening correction.

Another advantage of IMPACT as the backend method for PICCS is that it is based on MLTR, and therefore it models the detector signal as samples drawn from Poisson distribution. Therefore, it achieves lower noise levels in the image domain than other iterative methods without this assumption or working in log-domain. Despite that, I think that future research should aim towards incorporating noise-reducing priors while maintaining numerical accuracy. This becomes especially important since we have demonstrated that post-contrast images reconstructed with PICCS fully inherit the noise properties of the prior image. Moreover, we found that high spatial resolution in the prior image leads to more numerically accurate results in post-contrast PICCS images. We studied spatial resolution and noise level systematically as a function of the number of iterations and selected the optimal operating point based on these two characteristics. It is important to point out that the spatial resolution was optimized using the available data only. One of the issues with this dataset is that it contained only one object with a very large diameter. Therefore, we could perform MTF measurements in the image domain on the edges of this object. The major flaw of this method is the reporting of the final result, i.e., the maximum spatial resolution. The standard way to report the results of this measurement in literature is to use the frequency that corresponds to 10% MTF. Although this definition enables protocol optimization based on a quantitative metric, it remains an open question whether 10% MTF is sufficient to see small objects. In my opinion, in order to define the spatial resolution of this system, one needs to perform additional measurements with objects of different sizes filled with various concentrations of contrast agent and define the spatial resolution as the minimum size that can be detected for each concentration level. This type of study is known as contrast-detail study and, in general, requires a sophisticated phantom that, to the best of my knowledge, does not exist for breast CT, yet.

A rather unexpected result presented in **chapter 3** was the decrease of spatial resolution defined as frequency corresponding to 10% MTF with increasing iodine concentration in the post-contrast images reconstructed with PICCS. However, this has a simple explanation. The higher the concentration of iodine, the larger the difference between the linear attenuation coefficients in corresponding voxels between non-contrast and iodine-containing phases and the larger the values in the difference image. Since PICCS is trying to reduce the L1 metric on this difference image, it finds it more difficult to not oversmooth the peak. In my opinion, this presents the main limitation for the use of PICCS in DCE imaging, where high numerical accuracy at high spatial resolution is sought. I see two ways to improve the performance in this case. First, one should aim to achieve the highest possible spatial resolution in the prior image. Second, the use of a different sparsification transformation could lead to better performance. However, at this moment it is unclear which properties would be required for this different sparsification method.

My colleagues are currently working on verifying and testing PICCS reconstruction on phantoms with dynamically changing iodine concentration that is known at each time point. I contributed to the development of this kind of phantoms, although not as a part of the work presented in the previous chapters. Together with colleagues we developed a setup that can be used to estimate concentration of iodinated solutions using absorption of visible light. My contribution to this development was the help with the selection of optical components and writing software to acquire data from photodetectors. First results obtained with this new setup have already been published and presented at various conferences [3–5]. They are well aligned with the conclusions regarding strong points and drawbacks of the optimized protocol presented in **chapter 3.**

The main limitation of the study in **chapter 3** is that all experiments were performed on an immobile phantom. It is obvious that patient motion will drastically affect the performance of PICCS because the change in the linear attenuation coefficients will not be explained solely by differences in the iodine concentration. While working on this chapter, we developed a rather straightforward method to compensate for the intraframe motion by rigid registration of the post-contrast projections to the projections of the prior image. That this method achieves reasonable efficiency can only be verified once this algorithm is applied to clinical cases. At the present moment, we have limited knowledge of the amount of motion and how it would affect the accuracy of the PICCS reconstructions. Brombal *et al.* concluded in their study, in which they tracked motion in volunteers in synchrotron breast CT using optical markers, that the typical range of motion in the vertical direction is between 0.4 mm and 1 mm over a time span of 40 seconds [6]. Moreover, they found that the motion patterns are periodic and are caused by breathing. This result is rather unexpected as it is believed that prone positioning of the patient prevents breathing motion. This assumption is usually valid only for MRI where breathing is unlikely to affect low-resolution scans. Despite our own data being limited only to single-phase scans of 10 seconds, we could not observe periodic patterns due to breathing.

After reading the entire thesis, the reader might argue that the results presented in **chapter 4** can be used to perform the first estimation of intraframe motion, but I have to disagree. The important difference between the motion extracted from DCE-MRI in **chapter 4** and potential patient motion in breast CT is the positioning of the patient. In breast MRI, both breasts are pending through openings while both shoulders lie on a special support, offering stability. This stability is missing in breast CT, where only one shoulder is resting on the table, while, in general, the second shoulder is not touching the surface of the table. This is not an issue for most patients when a single breast CT scan is performed. But it might become an issue when the woman has to remain in the same position for more than five minutes. The next factor that makes the extracted motion not a very good predictor of motion in breast CT is the presence of motion compensation in DCE-MRI. We do not have any information about the type of motion compensation methods applied to the data we used to extract the motion field. Hence, it is quite possible that the data used in **chapter 4** as representative of patient motion is actually the residual motion left after the DCE-MRI system corrected for the actual patient motion.

Nevertheless, **chapter 4** provided a good starting point to add motion to existing digital phantoms in order to obtain ground truth data for testing motion compensation methods. However, I do not see further application for this method at this time. As already mentioned above, the motion vector fields that are extracted from DCE-MRI sequences are usually processed with motion compensation methods and do not represent the entire range of motion. Moreover, the intraframe motion needs to be parameterized to be used as interscan motion. It is impossible to simulate all possible cases since motion of various magnitude can happen at any time during the scan. I believe that the methods that are are based on finite element simulations are superior for adding motion to digital phantoms.

In **chapter 5**, we introduced a non-rigid motion-artifact reduction method for single-phase breast CT. This method reduces the motion artifacts by non-rigorously modifying the volume in order to reduce the L1-metric between the forward projection and measured data. Although the concept of minimizing data discrepancy is certainly not new, the use of automatic differentiation to estimate the partial derivative for the update step of each motion parameter independently is novel. The use of automatic differentiation enabled the optimization of the motion parameters in a reasonable time. We demonstrated that this method is successful at mitigating the motion artifacts that can be encountered in breast CT. I expect that outside of breast imaging this motion-artifact reduction method can also be successfully applied when scanning anatomies with high contrast using cone-beam CT. The presence of high-contrast objects simplifies the use of the L1 loss in the projection domain and gradient descent since there is an easily-findable single update direction that needs to be followed to reduce the motion artifacts. The application of the developed algorithm in lung or trauma imaging is an interesting research direction worth pursuing.

Severe motion still presents a great challenge for this method and additional strategies need to be developed and tested to improve the method. I believe that the most promising solution to this problem would be rigid precorrection. It should eliminate the largest motion components while still maintaining the shape of the scanned object. Along this direction, future research should be focused on speeding up the algorithm and on incorporating the motion parameters into iterative model-based reconstruction methods, such as MLTR or IMPACT.

As it often happens in research, the best results are achieved when they are least expected. Work in **chapter 6** certainly falls into this category. I started working on the acceleration of the 4D similarity filter to learn a new skill and help a colleague who was working on another project. In the end, we were able to significantly improve the runtime of the noise filter while still maintaining its denoising performance. The GPU version of the 4D similarity filter uses a hash-like data structure to store candidate values for calculation of the average voxel value. This hash-like data structure can be implemented efficiently on modern GPUs. However, due to its approximate nature it is not possible to match the outputs of GPU and CPU versions of the filter. In **chapter 6**, we have shown that both implementations lead to similar results and I believe that this difference will not have negative impact. In fact, it is possible that the resulting dynamic data-driven effective filter strength resulting in the GPU implementation results in better outcomes, but this needs to be investigated further. Another advantage of the new GPU-based version of the 4D similarity filter is simplified code, since there are no low-level hardware specific instructions that make code modification a non-trivial task. Incorporation of neighborhood search or use of different similarity metrics is rather straightforward. Currently, the modified version of the filter is already used in research and shows very promising results in stroke imaging in brain CT scans [7]. The GPU version of the filter presented in this chapter enables accelerated research on 4D filtering and can find its application in the clinic, when time needed to make the diagnosis is crucial. Unfortunately, due to the high degree of time-domain correlation in PICCS images, we cannot use this filter to reduce image noise in the developed DCE breast CT. However, if we will find later that the PICCS approach results in sub-par performance on clinical cases, the use of the 4D similarity filter may become a viable alternative for low-dose DCE breast CT imaging.

When I started working on this thesis a bit more than five years ago, our goal was to develop a breast imaging modality to enable numerically accurate dynamic contrast-enhanced breast imaging at high spatial and temporal resolutions. Now, as I approach the writing of the end of this chapter, I can say that this goal was achieved. We have a solid technical foundation to start investigating whether DCE breast CT provides additional clinical value to other imaging modalities or can even replace them when it comes to personalizing treatment and monitoring treatment response. I will quickly repeat the main characteristics of the developed modality and speculate why it can succeed at this task.

Spatial resolution depends on the iodine concentration in the post-contrast images. The lowest in-plane spatial resolution will be around 0.9 mm^{-1} at 10 mg I/mL. However, this resolution is still almost twice than what can be achieved in MRI. At lower iodine concentrations, the spatial resolution will be higher, making the imaging of large cancers at the subtumor level possible.

Temporal resolution is limited mainly by the rotational speed of the gantry. In the current implementation of the breast CT, the gantry rotates at 0.1 Hz, leading to 10 seconds gap between post-contrast images when in 360° reconstruction mode and around 6 seconds when in 180° + fanbeam reconstruction mode. Although this temporal resolution is not significantly better than what is possible in DCE-MRI, it is only limited mechanically and can be improved by focusing development towards installation of a more stable and faster mechanical setup without affecting spatial resolution.

Radiation dose is lowered by using sparse sampling and PICCS reconstruction. We could reduce both the number of projections and the fluence per projection.

Accurate numerical output is provided by the IMPACT reconstruction method that models the polychromatic linear attenuation coefficient and a scatter correction algorithm developed by colleagues in the lab.

Low dose[2] and high spatial and temporal resolutions are essential components of DCE breast CT that set it apart from other X-ray based breast imaging modalities or DCE-MRI. An additional advantage is the accurate numerical output that enables longitudinal studies. However, all these advantages will not result in DCE breast CT being used for the envisioned clinical applications if intraframe patient motion plays an important role and cannot be mitigated.

---

[2] Although, many patients will receive radiotherapy, not all women will undergo it. Therefore, we should follow the ALARA principle. ALARA stands for As Low As Reasonably Achievable and dictates to use the lowest possible radiation dose that still leads to accurate diagnosis.

Future work must first and foremost focus on performing a large patient study. The study must answer the question of whether patient motion has a substantial effect on numerical accuracy and how well the same image quality is maintained across images of the same and different patients. Furthermore, relevant perfusion maps need to be determined. This process will likely involve development and further improvement of blood perfusion models in the breast. Currently, the enormous amount of output data cannot be easily and efficiently interpreted and used to draw clinical conclusions and, therefore, must be simplified.

To wrap up this work, I want to emphasize that researchers and breast CT manufacturers need to work more closely to introduce DCE breast CT in the clinic. However, even when this imaging modality reaches its final stage and is proven to produce consistently accurate numerical results, radiologists will need to be convinced to give it a chance. At this moment there is a general reluctance towards CT perfusion when imaging other organs largely caused by concerns about radiation dose, and it will not be different in breast imaging. For DCE breast CT to be successful it either needs to excel in treatment monitoring, which can only be determined by multiple clinical studies or be the most cost-effective solution at comparable level of performance to other modalities. In other words, it must significantly improve the clinical outcome for women with breast cancer. Technical superiority alone will not be enough, but it is absolutely necessary for this. I firmly believe that the results in this work can be used to move it several steps closer to achieving this goal.

References

[1] S. M. Midgley. A parameterization scheme for the x-ray linear attenuation coefficient and energy absorption coefficient. *Physics in Medicine and Biology*, 49(2):307–325, January 2004.

[2] M. Nassi, M. Mikerov, A. Tomal, K. Michielsen, and I. Sechopoulos. Efficient simulation of biological tissue variability using midgley's decomposition. In *Proceedings of the The 8th International Conference on Image Formation in X-Ray Computed Tomography*, 2024.

[3] L. Goris, M. Mikerov, S. Manohar, and I. Sechopoulos. In-line spectroscopy for the quantification of iodine in a breast perfusion phantom for dynamic contrast-enhanced dedicated breast CT imaging validation. In Rebecca Fahrig, John M. Sabol, and Ke Li, editors, *Medical Imaging 2024: Physics of Medical Imaging*, volume 12925, page 1292547. International Society for Optics and Photonics, SPIE, 2024.

[4] L. Goris, J. J. Pautasso, M. Mikerov, K. Michielsen, and I. Sechopoulos. Perfusion phantom for the optimization of dynamic contrast-enhanced dedicated breast CT: iodine contrast curves in a simplified breast phantom. In Maryellen L. Giger, Heather M. Whitney, Karen Drukker, and Hui Li, editors, *17th International Workshop on Breast Imaging (IWBI 2024)*, volume 13174, page 131740U. International Society for Optics and Photonics, SPIE, 2024.

[5] L. C. Goris, M. Mikerov, J. J. Pautasso, and I. Sechopoulos. In-line spectroscopy for iodine quantification in dynamic contrast-enhanced dedicated breast ct. *Medical Physics*, 52(2):1037–1044, 2025.

[6] L. Brombal, L. M. Arana Peña, F. Arfelli, R. Longo, F. Brun, A. Contillo, F. Di Lillo, G. Tromba, V. Di Trapani, S. Donato, R. H. Menk, and L. Rigon. Motion artifacts assessment and correction using optical tracking in synchrotron radiation breast ct. *Medical Physics*, 48(9):5343–5355, 2021.

[7] S. A. M. Tunissen, E. J. Smit, M. Mikerov, and I. Sechopoulos. Four-dimensional similarity filter for cerebral ct perfusion maps in acute stroke. *Radiology*, 315(2):e241704, 2025. PMID: 40358453.

SUMMARY

Although breast CT is the only imaging modality that can deliver fully 3D images at high spatial resolution and without causing patient discomfort, it is not suited for screening due to its higher radiation dose compared to mammography and DBT. However, because unlike in MRI the spatial resolution is decoupled from the temporal resolution, the addition of dynamic perfusion imaging to breast CT opens up new possibilities of using breast CT to monitor neoadjuvant treatment of breast cancer. This thesis discusses the technical advancements necessary to accomplish highly accurate dynamic perfusion imaging using breast CT.

Chapter one provides a General Introduction to the monitoring of breast cancer treatment, dynamic CT imaging, dedicated breast CT, image reconstruction in 3D and 4D, as well as image noise reduction techniques.

Chapter two explores theoretically whether a more accurate model of linear attenuation coefficients using Midgley's decomposition into five coefficients is more useful than classical decompositions into just two components. It comes to the conclusion that Midgley's decomposition is more accurate for the low-energy range of breast CT, but not practical for application in reconstruction.

Chapter three presents an optimized protocol for dynamic perfusion imaging in breast CT to achieve high numerical accuracy in the estimation of iodine content. This is made possible while using polychromatic reconstruction from sparse data involving just 40 projections. The chapter also discusses limitations causing lower spatial resolution than in comparable static scans.

Chapter four describes an approach of combining patient motion data that can be obtained from DCE-MRI sequences to create digital breast CT phantoms with motion. The generated phantoms are then used to create projection data for better understanding of motion artifacts in breast CT. Similar artifacts to the ones observed in real clinical data can be achieved with these simulations.

Chapter five introduces a non-rigid motion compensation algorithm for single-phase breast CT scans based on transforming the volume using automatic differentiation to match the reprojected volume and measured projections. The correctness of the method was verified in a simulation study and its performance was evaluated on real clinical cases from two institutions. Overall, the algorithm reduces the appearance of motion artifacts. However, artifacts caused by motion with large amplitude cannot always be fully removed.

Chapter six explores the possibility of using approximations to speed up a 4D noise filter by making it possible to execute it on a GPU. Old and new versions of the filter are tested on the same simulated data. The GPU version performs up to 23 times faster.

The thesis ends with a General Discussion, containing the author's opinion about the methods presented in each chapter and their impact on dynamic perfusion breast CT imaging. Moreover, this thesis has a dedicated data management plan, this summary, author's CV, and his scientific contributions during his PhD trajectory.

SAMENVATTING

Hoewel borst-CT de enige beeldvormingsmodaliteit is die echte 3D-beelden kan leveren met hoge spatiale resolutie en zonder ongemak voor de patiënt, is het niet geschikt voor bevolkingsonderzoek vanwege de hogere stralingsdosis in vergelijking met mammografie- en DBT-onderzoeken. De toevoeging van dynamische perfusiebeeldvorming aan deze modaliteit biedt echter nieuwe mogelijkheden voor het gebruik van borst-CT om de neoadjuvante behandeling van borstkanker te controleren, omdat, in de tegenstelling tot MRI, spatiele resolutie onafhankelijk is van de temporele resolutie. In dit proefschrift worden de technische ontwikkelingen besproken die nodig zijn om nauwkeurige dynamische perfusiebeeldvorming met borst-CT te realiseren.

Hoofdstuk één biedt een algemene inleiding over het opvolgen van de behandeling van borstkanker, dynamische CT-beeldvorming, de borst-CT, beeldreconstructie in 3D en 4D en technieken voor beeldruisonderdrukking.

Hoofdstuk twee onderzoekt theoretisch of een nauwkeuriger model van lineaire verzwakkingscoëfficiënt met behulp van Midgley's decompositie in vijf coëfficiënten voordeliger is dan klassieke decomposities in slechts twee componenten. De conclusie van dit hoofdstuk is dat Midgley's decompositie nauwkeuriger is bij de lage fotonenenergie van borst-CT, maar niet praktisch toepasbaar in beeldreconstructie.

Hoofdstuk drie presenteert een geoptimaliseerd protocol voor dynamische perfusiebeeldvorming in borst-CT. Het onderzoek is gericht op het bereiken van een hoge numerieke nauwkeurigheid van de aanwezige jodiumsconcentratie en een lage stralingsdosis met behulp van polychromatische reconstructie van slechts 40 projecties. Het hoofdstuk bespreekt ook beperkingen die een lagere spatiale resolutie veroorzaken dan in vergelijkbare statische scans.

Hoofdstuk vier beschrijft hoe beweging van patiënten uit DCE-MRI-sequenties kan worden geïntroduceerd in digitale CT-fantomen. De gegenereerde fantomen worden vervolgens gebruikt om projectiedata te creëren voor een beter begrip van bewegingsartefacten in speciale borst-CT. De gecreëerde bewegingsartefacten zijn vergelijkbaar met de artefacten in klinische beelden.

Hoofdstuk vijf introduceert een algoritme voor niet-rigide bewegingscompensatie voor borst-CT, gebaseerd op het transformeren van het volume met behulp van automatische differentiatie om het gereprojecteerde volume en de gemeten data op elkaar af te stemmen. De nauwkeurigheid van de methode werd geverifieerd in een simulatiestudie en verder geëvalueerd op echte klinische scans uit twee instellingen. In het algemeen vermindert het algoritme de bewegingsartefacten, maar wanneer ze het gevolg zijn van beweging met grote amplitude, kunnen de artefacten niet altijd volledig worden verwijderd.

Hoofdstuk zes onderzoekt de mogelijkheid om specifieke benaderingen te gebruiken om een 4D-beeldruisfilter te versnellen door het uitvoeren op een GPU. Oude en nieuwe versies van de filter worden getest op dezelfde gesimuleerde gegevens. De GPU versie is tot 23 keer sneller.

Het proefschrift eindigt met de Algemene Discussie. Deze bevat voornamelijk de mening van de auteur over de methoden die in elk hoofdstuk worden gepresenteerd en hun invloed op dynamische CT-beeldvorming van borstperfusie. Daarnaast bevat dit proefschrift een datamanagementplan, deze samenvatting, het CV van de auteur en zijn wetenschappelijke bijdragen tijdens zijn promotietraject.

ACKNOWLEDGMENTS

More than five years have passed since I joined AXTI. During this time I met and worked with a lot of amazing people who assisted me in this journey.

I will, of course, begin with my supervision team. I would like to thank *Ioannis* for helpful advice on various aspects of medical physics, breast imaging, and, especially, presenting and writing scientific texts, including the usage of articles in English, the choice of active/passive tense, Bland-Altman plots, adding zoom-in inserts in figures, and

- bullet points.

I enjoyed our Christmas dinners, AXTI weekends, and cocktail evenings at conferences, even though I wasn't always there. I valued your availability and willingness to help no matter where you were at any given moment. I also hope that you learned a thing or two about LaTeX and are not so much against it anymore. Thank you, *Ritse*, for valuable discussions about clinical aspects of this work. Finally, I would like to thank *Koen* for being my weekly supervisor, which meant dealing with me every single week, going through first drafts of everything I wrote, and accepting that we have different views on many things. Because of you I know so much about CT reconstruction, polychromatic models, scatter correction, (distance-driven) projectors, and CUDA programming. I always appreciated that I could come to you with any silly question and you had time and patience to go through them with me. I am still considering putting a rubber duck named "Koen" on my desk for Quack debugging. Despite not being an official member of my supervision team, I am glad that *Marjolein* was always there to assist me with solving any bureaucratic problem.

Next, I would like to thank my fellow colleagues with whom I closely worked on the breast 4D project. Although we had completely different approach to the project and disagreed on many things, I valued your structured and methodical attitude to research, *Juan*. It was reassuring to know that I could always rely on you to get data that I deleted previously. Working with *Liselot* meant having lots of fun with transfusion of olive oil, doing at least something with hands, and even building optical setups and programming hardware. It was an entertaining distraction from sitting in front of a PC whole day. I always looked forward to the days when you came to the lab.

I'm not allowed to forget my co-authors. Thank You, *Andrew* and *John*, for helping me with the motion compensation (sorry, motion-artifact reduction) project. Even though it did not lead to anything, Thank You, *Jim*, for some days that you dedicated to helping me with Midgley's decomposition. Finally, I would like to thank *Nikita* for sharing his awesome projector library with me, interesting discussions about autodifferentiation, grid sampling, Hermitian adjoints, and other mathematical properties.

Other PhD candidates, with whom I spend the majority of time in three different offices, including dark nine months long exile on sunless and lifeless -1, are next! Thank You, *Wendelien* and *Suzanne*, for helping me with studies involving MRI. I am very happy that I could continue the work of *Marco* on breast 4D. Also Thank You for storing my stuff for a few days when I was moving to Nijmegen and picking up the keys for me. I also enjoyed having *Luuk* at my side, the real MVP of our department, when I needed help with medical physics or taking scans.

I am grateful to *Marta* and *Miguel* for introducing me to tennis - the game that helped me to stay distracted from everyday difficulties during my PhD, visiting tennis tournaments and Christmas markets, going to salsa parties, organizing boardgames evenings and being, in general, very often to spending time together. I also liked our Corona-walks, which were often joined by *Joana* and *Olga*. I learned a thing or two about Portugal and the importance of drinking without slurping. However, I am glad it was not me who caused this discussion.

Sjoerd, helaas is het ons ondanks onze dromen en hard werken niet gelukt om medische beeldvorming te veranderen. Ik geloof ook niet dat wij ooit echt dichtbij zijn geweest. Maar ik denk nu terugkijkend naar onze PhD, dat dat ook niet het belangrijkste was. Ik ben blij dat we door de tijd dat we samen op het kantoor en buiten Radboud hebben doogebracht vrienden zijn geworden. Bedankt dat je me hebt laten kennismaken met wielrennen. Bedankt ook voor niet switchen naar het Engels vanaf het begin. Dankzij jou en *Sophie* kan ik Nederlands praten. Ik waardeer het dat ik je bruiloft heb mogen bijwonen.

Maranda, ik neem aan dat je blij bent dat je naar vlakke Utrecht bent verhuisd. Maar gezien je vooruitgang op Strava, ben ik ervan overtuigd, dat je geen huivels hoeft te vrezen. Ik moet toegeven dat ik vanwege jou nu ook klimmen en afdalen (of afdalen en klimmen) op dezelfde weg vermijd.

Vielen Dank, *Franziska*, für die Möglichkeit, mit dir Deutsch zu sprechen, nette Fahradausflüge rund um Nijmegen, und gemeinsames Besuchen von verschiedenen Museen in den Niederlanden.

Jessie and *Sarah*, two people who liked to stay late in the office, which often lead to interesting conversation not only about science cannot be forgotten, too. Thank You, Jessie, for familiarizing me with must-know Dutch cultural heritage of the 00' and 10' through Dumpert.

Thank You, *Marialena*, for helping me with the clinical part of this work. It was always fun to start MRI vs. breast CT discussions with you because of how passionate you were about the entire subject. I also hope that you are more used to Dutch stairs now.

Thank You, *Daan* and *Leo*, for being participants and subjects of very interesting discussions and events. It was always a joy to listen to Daan's crazy stories, maybe a bit less, but still very entertaining to his complaints about pretty much everything when he visited our office. Although tennis wasn't really your favorite sport, I enjoyed the few games we had, Leo, even though I was without any doubts intentionally hit by balls. But I accepted that it is just a part of your highly-competitive nature.

Gustavo, *Martina*, *Raneim*, and *Hanne* - you joined us later but somehow managed to be part of most unforgettable memories of my PhD, be it the Venezuelan bank, complete disrespect for shark teeth or not locking the bike when going to a supermarket and many more. I am glad that I noticed your drawing talent, Gustavo. Also I think that light blue background in PowerPoint is absolutely brilliant! I always enjoyed talking to you, Martina. I valued you honest yet fun way of saying things. Thank You for teaching me how to walk slowly, although, I must admit, I am still very bad it. You, on the other hand, were a much more capable student, given how fast you managed to take over my workstation admin

job. And I am happy that you could be on that flight despite all the obstacles from DB!

Big Thank You to my two friends from Marburg - *Jochen* and *Alex* - for wonderful and quite eventful trips to Norway and Corsica, for binge-watching Clarkson's farm, and always being there to discuss all aspects of the PhD live.

I also appreciate the support of my colleagues at SeeThrough in the last part of this journey.

Finally, I would like to thank my grandmother, who sadly passed away during this time, and my mother for believing in me and providing support.

RESEARCH DATA MANAGEMENT

In this thesis different types of data were used, consisting of phantom images, simulated data, and retrospective patient images. Each chapter is discussed below.

Chapter two uses simulation data only. Simulation code is published in a Data Sharing Collection (DSC) in the Radboud Data Repository (RDR) with open access (`https://doi.org/10.34973/27m3-c202`, CCO). Chapter two is published with open access.

Chapter three uses only measured data of a physical phantom. Raw projection data and geometry specifications of the breast CT scanner are archived in their original form in a Data Acquisition Collection (DAC) in the RDR under closed access (`https://doi.org/10.34973/4ptk-ez10`).

Chapter four uses retrospective pseudo-anonymized patient data performed under CMO 2016-2774, Project 20068. Two extracted usable motion vector fields in this chapter are archived in a DSC in the RDR (`ru.rumc.mmrict_t0000541a_dsc_950`, CCO). Chapter four is published without open access.

Chapter five uses simulation data, measured data of physical phantoms, and retrospective anonymized patient data from RadboudUMC and UC Davis. The data from RadboudUMC used in this research was obtained during the study "Dedicated Breast CT for Diagnosis of Breast Cancer" (NL55378.091.15). The participants gave consent for the re-use of their data for other research purposes. The data from UC Davis was approved by the Institutional Review Board (#214750) at UC Davis. Raw projection data of phantom acquisitions and simulations, as well as geometry specifications of the breast CT scanner are archived in in their original form in a DAC in the RDR under closed access (`https://doi.org/10.34973/4g4r-j351`)

Chapter six uses simulation data only using XCAT phantoms. The data underlying chapter six are not suitable for reuse because of vendor's intellectual property. The simulation tool used to create this data is archived in the RDR: `https://doi.org/10.34973/4zgg-5j18`

LIST OF PUBLICATIONS

Publications as first author:

1. 🗎 *Mikhail Mikerov*, Juan J. Pautasso[3], Liselot Goris, Koen Michielsen, and Ioannis Sechopoulos: 4D Dynamic Contrast-Enhanced Breast CT: Phantom-based Reconstruction Parameter Optimization for Iodine Quantification

   Med Phys. 2025; 52: 2212-2223

2. 🗎 *Mikhail Mikerov*, Koen Michielsen, Nikita Moriakov, Juan J. Pautasso, Sjoerd A.M. Tunissen, Andrew M. Hernandez, John M. Boone, and Ioannis Sechopoulos: Non-rigid motion compensation in dedicated breast CT

   Transactions on Biomedical Engineering 2025; 72(10): 3133-3145

Publications as contributing author:

1. 🗎 Juan J. Pautasso, Marco Caballo, *Mikhail Mikerov*, John M. Boone, Koen Michielsen, and Ioannis Sechopoulos: Deep learning for x-ray scatter correction in dedicated breast CT

   Med Phys. 2023; 50: 2022–2036

2. 🗎 Sjoerd A.M. Tunissen, Nikita Moriakov, *Mikhail Mikerov*, Ewoud J. Smit, Ioannis Sechopoulos, and Jonas Teuwen: Deep learning-based low-dose CT simulator for non-linear reconstruction methods

   Med Phys. 2024; 51: 6046–6060

3. 🗎 Sjoerd A.M. Tunissen, Ewoud J. Smit, *Mikhail Mikerov*, Mathias Prokop, and Ioannis Sechopoulos: Performance evaluation of a 4D similarity filter for dynamic CT angiography imaging of the liver

   Med Phys. 2024; 51: 8814–8827

4. 🗎 Liselot C. Goris, *Mikhail Mikerov*, Juan J. Pautasso, and Ioannis Sechopoulos: In-line spectroscopy for iodine quqntification in dynamic contrast-enhanced dedicated breast CT

   Med Phys. 2024; 52: 1037-1044

---

[3] joint first author

5. Sjoerd A. M. Tunissen, Ewoud J. Smit, *Mikhail Mikerov*, and Ioannis Sechopoulos: Four-dimensional similarity filter for cerebral CT perfusion maps in acute stroke

   Radiology 2025; 315(2): e241704

Accepted conference abstracts as first author:

1. *Mikhail Mikerov*, Koen Michielsen, Nikita Moriakov, and Ioannis Sechopoulos: Adding patient motion from DCE-MRI to anthropomorphic phantoms for dedicated breast CT

   SPIE Medical Imaging 2022

2. *Mikhail Mikerov*, Koen Michielsen, Nikita Moriakov, and Ioannis Sechopoulos: Rigid motion compensation for breast CT

   16th International Workshop on Breast Imaging (IWBI 2022)

3. *Mikhail Mikerov*, Koen Michielsen, James J. Nagy, and Ioannis Sechopoulos: Effect of attenuation model on iodine quantification in contrast-enhanced breast CT

   7th International Conference on Image Formation in X-Ray Computed Tomography (CT Meeting 2022)

4. *Mikhail Mikerov*, Koen Michielsen, Nikita Moriakov, and Ioannis Sechopoulos: Non-Rigid Motion Compensation for Breast CT

   2022 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC 2022)

5. *Mikhail Mikerov*, Sjoerd A.M. Tunissen, Koen Michielsen, Ewoud J. Smit, and Ioannis Sechohpoulos: Adaptation of a 4D noise filter for implementation on GPU

   8th International Conference on Image Formation in X-Ray Computed Tomography (CT Meeting 2024)

Accepted conference abstracts as contributing author:

1. Liselot C. Goris, *Mikhail Mikerov*, Srirang Manohar, and Ioannis Sechopoulos: In-line spectroscopy for the quantification of iodine in breast perfusion phantom for dynamic contrast-enhanced dedicated breast CT imaging validation

   SPIE Medical Imaging 2024

2. Liselot C. Goris, Juan J. Pautasso, *Mikhail Mikerov*, Koen Michielsen, and Ioannis Sechopoulos: Development of a breast perfusion phantom including in-line spectroscopy measurements for the validation and testing of dynamic contrast-enhanced dedicated breast CT

   European Congress of Radiology (ECR 2024)

🗎 3. Liselot C. Goris, Juan J. Pautasso, *Mikhail Mikerov*, Koen Michielsen, and Ioannis Sechopoulos: Perfusion phantom for the optimization of dynamic contrast-enhanced dedicated breast CT: iodine contrast curves in a simplified breast phantom

17th International Workshop on Breast Imaging (IWBI 2024)

🗎 4. Martina Nassi, *Mikhail Mikerov*, Alessandra Tomal, Koen Michielsen, and Ioannis Sechopoulos: Efficient simulation of biological tissue variablity using Midgley's decomposition

8th International Conference on Image Formation in X-Ray Computed Tomography (CT Meeting 2024)

🗎 5. Liselot C. Goris, Juan J. Pautasso, *Mikhail Mikerov*, Koen Michielsen, and Ioannis Sechopoulos: In the spotlight: 4D dynamic contrast-enhanced dedicated breast CT, a phantom study for the validation of a novel imaging technique

10th Dutch Bio-Medical Engineering Conference (BME 2025)

PhD portfolio

Department: Medical Imaging
PhD period: 01/06/2020 – 31/05/2024
PhD Supervisor(s): Prof. Dr. Ioannis Sechopoulos
PhD Co-supervisor(s): Dr. Ritse Mann, Dr. Koen Michielsen

Training activities	**Hours**
Courses	
– RU - Presentation Skills (2020)	42.00
– RIHS - Introduction course for PhD candidates (2020)	15.00
– Dagcursus Nederlands voor beginners II (A1-A2) (2021)	57.00
– Fundamentals of Accelerated Computing with CUDA Python (2021)	8.00
– RU - Academic English Conversation and Pronunciation (2021)	43.00
– Radboudumc - Scientific integrity (2021)	20.00
– RU - Writing Scientific Articles (2022)	96.00
Conferences	
– 6th International Conference on Image Formation in X-Ray CT Reconstruction (2020) (no contribution, online)	25.00
– 16th International Meeting on Fully Three-Dimensional Image Reconstruction (2021) (no contribution, online)	25.00
– 16th International Workshop on Breast Imaging (2022) (poster)	35.00
– IEEE NSS/MIC/RTSD Conference (2022) (poster)	35.00
– 8th International Conference on Image Formation in X-Ray CT (2024) (poster)	35.00
Teaching activities	
Supervision of internships / other	
– Supervision of the bachelor student (2024)	50.00
Total	**486.00**

CURRICULUM VITÆ

Mikhail Mikerov was born in Perm, Russia, on April 26, 1993.

In 2013, he started his B. Sc. program in physics at Philipps University Marburg in Germany from which he graduated in 2016. During his bachelor's degree he developed strong interest in optics and semiconductor physics.

In 2016, he continued his education in physics at Philipps University Marburg. In 2019, he obtained his M. Sc. with a thesis on algorithmic removal of unwanted signal components from terahertz time-domain spectroscopy data.

In the same year, he started working as a researcher in the same group where he completed his master's thesis.

In June 2020, he joined AXTI (Advanced X-Ray Tomographic Imaging) lab at Radboud University Medical Center in Nijmegen, The Netherlands.

During his PhD, he focused on various CT reconstruction problems. His main projects were quantification of iodinated contrast agents in breast perfusion scans and motion compensation in single-phase breast CT.

Moreover, he contributed to other projects predominately revolving around understanding and optimizing 4D similarity filter as well as developing in-line optical spectroscopy for perfusion phantoms.

After finishing work on his thesis, he stayed committed to research and development in image reconstruction in CT by joining dental CT manufacturer W&H (See Through s.r.l.) as a scientist.

9 789465 151731 >

Radboudumc
university medical center

Radboud University