

Master Thesis


Feasibility and Quality Assessment of Model-Based Respiratory Motion Compensation in Positron Emission Tomography

Manon Verra

Student number: 3938344

Study: Experimental Physics

Utrecht University



Supervisor:

Dr Andre Mischke

Utrecht University

Institute of Subatomic Physics

Daily supervisor:

Remco Bastiaannet, MSc

UMC Utrecht

Department Radiology

Second examiner:

Dr Hugo de Jong

UMC Utrecht

Department Radiology

Faculty of Science

August 2017



Utrecht University



UMC Utrecht

Front page image retrieved from [1].

Abstract

Quantitative information resulting from PET scans, such as tumor volume and radionuclide uptake, is crucial to radiologists for the diagnosis and staging of cancer, as well as for treatment planning. In recent years, many advances have been made in the development of quantitative reconstruction models for Positron Emission Tomography (PET). These reconstruction models correct for quantitative effects such as absorption, scatter and random coincidences.

However, the possibilities to include respiratory motion correction in these quantitative reconstructions models are still very limited and rely on the assumption that patient breathing is stationary. It is however known from radiotherapy that patient breathing is highly variable. This poses a major problem in quantitative PET, as respiratory motion causes image artifacts, such as blurring of the tumor, which results in overestimation of the tumor volume and underestimation of radionuclide uptake.

Therefore, the aim of this study was to develop a novel motion correction method for PET acquisitions that can deal with breathing variability. This was achieved by incorporating a model for the internal breathing motion in the iterative image reconstruction. The internal breathing motion was assumed to correlate with certain indicators of the respiratory phase of the patient (e.g. end-expiration), which could for instance originate from a respiratory belt. This correlation is described by a motion model which, given such an indicator of the respirator phase, predicts the internal motion. In this study, a linear motion model was used. The motion model parameters were assumed to be constant during an acquisition. Given the perfect motion model for a specific acquisition, the internal motion at each time is exactly known and can perfectly be corrected for. However, in practice these motion models are not a priori known so the motion model parameters are optimized during the iterative image reconstruction.

Internal motion can consist of motions such as translations, rotations and deformations. This study focused on the correction of translational motion. As the motion of the liver mainly consists of translations in the head-foot direction, this new model-based motion correction method (MBC) was applied to the liver. MBC was tested for both variable and non-variable breathing dynamics through simulations and experiments. Simulations and experiments were also performed to test the behavior of the model error if the motion model parameters are not constant during the acquisition.

All results were compared to state-of-the-art motion correction methods. In respiratory gating, only the data that was acquired in a limited window of respiratory phases (e.g. around end-inspiration) is used for image reconstruction. In another motion correction method, all of the PET data is divided over a number (often five or eight) of such gating windows that are reconstructed independently. These individual images can then be registered to form one final motion corrected image.

The results showed that in the case of variable breathing, MBC performed similar to the

other methods in terms of lesion volume and mean radionuclide uptake. The contrast-to-noise ratio (CNR) of MBC is always higher than or similar to those of the other methods. In the case of non-variable breathing dynamics, all methods underestimated the CNR and radionuclide uptake and overestimated tumor volume with respect to a completely motion-free reconstruction. This was unexpected for MBC as it was hypothesized that, given a perfect estimate of the motion model parameters, MBC would perform equally well for variable and non-variable breathing dynamics. The problem is most likely caused by problems in the optimization of the motion model parameters. If the parameter values resulting from the optimization highly deviate from the correct parameter values, the internal motion cannot be fully corrected for, causing residual motion to be present in the final reconstructed image.

The results show that the model error as a function of different model parameters is smooth and convex and has a clear optimum near the perfect model parameter values, making it very suitable for optimization. Contrary to expectation, the model error is not increased for acquisitions with non-constant motion model parameters. This indicates that the current model error cannot be used as an absolute measure to compare the amount of residual motion in different acquisitions.

In conclusion, the results showed that it is feasible to optimize a motion model using PET data. An abstract reporting this result that I contributed to was accepted for a poster presentation at the 2017 IEEE Nuclear Science Symposium and Medical Imaging Conference. A better quantitative analyses should be performed to draw a definite conclusion about the performance of model-based motion correction and its benefit in the case of variable breathing dynamics. The chosen model error proved to be suitable for optimization of the motion model, but cannot currently be used to compare the goodness of motion correction of different PET acquisitions.

Contents

1	Introduction	1
2	Background	3
2.1	Radioembolization	3
2.2	Positron Emission Tomography (PET)	7
2.3	Image reconstruction	13
3	Respiratory motion compensation	26
3.1	Breath-hold	26
3.2	Respiratory gating	26
3.3	Registration of gated images	29
3.4	Respiratory motion models	30
3.5	Research question	33
4	Method	35
4.1	Motion correction	35
4.2	Motion estimation	38
4.3	Simulations	45
4.4	Experiments	48
4.5	Quantitative analysis	52
5	Results	54
5.1	Simulations	54
5.2	Experiments	65
6	Discussion	81
7	Outlook	87
8	Conclusions	90
	Acknowledgements	91
A	Unsuccessful techniques	92
	References	96

Chapter 1

Introduction

Positron Emission Tomography (PET) serves its main applications in oncology [2]. Image quality and in particular quantitative information about tumors, such as tumor volume and radionuclide uptake, are crucial to the radiologist. These parameters can be used for the diagnosis and staging of cancer, as well as treatment planning.

In radiotherapy for instance, quantitative PET is used to determine the required tumor dose and target volume [3]. Accurate delineation of the tumor volume is required to avoid irradiation of healthy tissue, as well as to improve tumor coverage [4]. Other applications include radionuclide therapies, such as radioembolization. Radioembolization is a palliative liver cancer treatment that uses micron-sized radioactive spheres that selectively irradiate tumors, thus minimizing damage to healthy liver tissue [5]. In radioembolization, quantitative PET can be used to determine the patient-specific microsphere distribution and tumor dose [6]. These quantities can be used to optimize treatment planning, such as the amount of administered activity [7, 8].

Quantitative models for PET image reconstruction are essential to determine these quantitative parameters, such as radionuclide uptake and tumor volume, needed for the aforementioned applications. These models correct for quantitative effects such as absorption, scatter and random coincidences [9]. However, the possibilities to include respiratory motion correction in these quantitative reconstructions models are still very limited. This poses a major problem in quantitative PET, as respiratory motion causes image artifacts, such as blurring of the tumor. Tumor blurring results in overestimation of the tumor volume and underestimation of radionuclide uptake [10].

Several solutions to the problem of respiratory motion in quantitative PET have been proposed. In respiratory gating the acquired data is retrospectively binned in several respiratory phases or gates. Only one of these gates is reconstructed. This reduces motion artifacts, but increases image noise. Moreover, due to the width of the bin motion is not fully removed from the image. Another possibility is to reconstruct all gates independently and combine the result-

ing images into one motion-corrected image using registration. This method has the advantage of using all available statistics.

Both respiratory gating and registration of gated images rely however on the assumption that patient breathing is stationary, but it is known from radiotherapy that breathing is patient specific and highly variable, containing for instance amplitude and baseline drifts [11, 12]. Currently a common solution to overcome the problem of breathing variability is to coach patients to breathe more regularly. However, for some patients coaching is not suitable due to their medical condition or because they have difficulty following the coaching instructions [13, 14].

Therefore, the aim of this study was to develop a novel motion correction method for PET that can deal with breathing variability. This will be achieved by explicitly modeling the internal motion dynamics during the image reconstruction. The goal of this study is to test the performance of such a model-based respiratory motion correction method in PET and compare this with respiratory gating and registration of gated images through simulations and experiments.

The remainder of this thesis is structured as follows. In **Chapter 2**, the principles of radioembolization, Positron Emission Tomography (PET) and image reconstruction that are relevant for this study are discussed. In **Chapter 3**, a brief overview is given of the literature concerning respiratory motion compensation techniques and is concluded by motivating the aim for this study. In **Chapter 4**, the respiratory motion compensation technique that is proposed in this study is described, as well as how its performance was tested under variable breathing dynamics using simulations and experiments. In **Chapter 5** the results of the studies that were performed are presented. In **Chapter 6** the results of these studies are discussed, as well as some limitations of the studies that were performed. In **Chapter 7**, some possible directions for future research are discussed. In **Chapter 8**, this thesis is ended by presenting the conclusions.

Chapter 2

Background

This chapter discusses the required background for this study, starting with radioembolization as this is an important possible application. This is followed by discussions of the working principles of Positron Emission Tomography (PET) and image reconstruction techniques.

2.1 Radioembolization

Radioembolization is a liver cancer treatment that is mainly used in a palliative setting with patients that suffer from unresectable liver cancer [15]. Radioembolization uses micron-sized resin or glass spheres that contain the radioisotope yttrium-90 [16, 17]. These spheres are administered in the hepatic artery via a catheter. The microspheres end up primarily in tumor tissue, thus selectively irradiating tumor tissue over non-tumor tissue. This effect is caused by the liver's unique vascular system.

2.1.1 Vascular system of the liver

The liver's vascular system consists of two blood supplying vessels: the portal vein and the hepatic artery. The *portal vein* carries blood from the intestines and the stomach, to the liver. This blood is rich in nutrients and is processed by the liver that filters it from toxins. The *hepatic artery* is part of the normal blood circulation and supplies the liver with oxygenated blood. All blood leaves the liver via the *hepatic vein*. Figure 2.1 shows the liver with the portal vein, hepatic artery and hepatic vein.

2.1.2 Tumor to non-tumor (T/N) uptake ratio

Radioembolization is based on the fact that tumor tissue receives its blood primarily from the hepatic artery, whereas normal liver tissue receives its blood primarily from the portal vein. Therefore, administering the microspheres in the hepatic artery should lead to a higher accumulation of microspheres in tumor tissue than in normal liver tissue [18]. The ratio between

radionuclide uptake in tumor and non-tumor tissue (T/N ratio) will then be larger than 1. Preferably the T/N ratio should be as large as possible, thus maximizing the damaging effect on the tumor, while minimizing the radiation damage on healthy liver tissue.

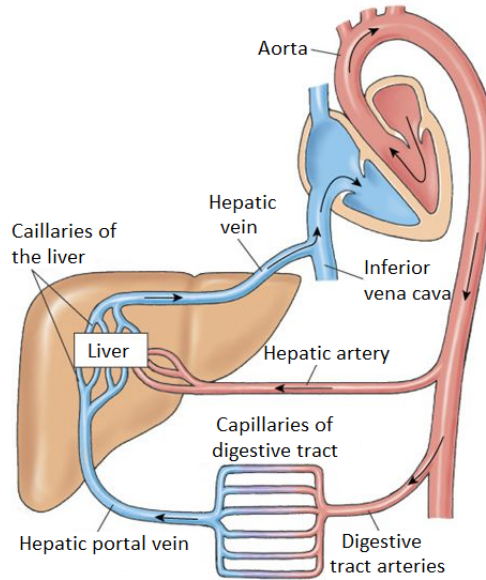


Figure 2.1: *The vascular system of the liver, showing the portal vein that carries blood from the digestive tract to the liver and the hepatic artery that carries oxygenated blood from the heart to the liver. The hepatic vein carries blood away from the liver, back to the heart [19].*

2.1.3 Beta radiation

Tumor destruction is caused by radiation originating from the yttrium-90 decay in the microspheres. Yttrium-90 has a half-life of 64.1 hours and its primary decay mode is the following β^- -decay to the ground state of Zirconium-90,



Figure 2.2 shows the energy spectrum of the electrons resulting from the yttrium-90 decay. These electrons have a mean energy of 930 keV with a maximum of 2270 keV. The electrons interact with the tumor tissue, depositing their energy and thereby inducing radiation damage. The mean range of the electrons in tissue is 2.5 mm with a maximum of 11 mm. It is important that this range is not too high, as it limits the accuracy of dose deposition.




Figure 2.2: *Energy spectrum of the electrons resulting from the yttrium-90 decay [20]. The electrons have a mean energy of 930 keV and a maximum energy of 2270 keV.*

2.1.4 Patient eligibility

As mentioned before, not all microspheres end up in tumor tissue. Part of the microspheres end up in healthy liver tissue, causing unwanted radiation damage. The amount of radiation damage done to healthy liver tissue depends on the patient's vascular system. Before the treatment, the patient's vascular system is therefore examined using pre-treatment imaging to assess its fitness for the treatment.

Only a small number of patients is found to be unsuited for radioembolization, usually due to vascular abnormalities or extensive lung shunting [5]. In the case of lung shunting, patients have an abnormal, direct connection between blood vessels of the liver and the heart, causing part of the microspheres to reach the lungs [15]. Patients with insufficient liver reserve are also not suited for radioembolization, as they risk radioembolization induced liver failure and treatment related morbidity due to the inevitable non-tumor dose [5]. The next section will discuss how pre-treatment imaging is used to assess the fitness of a patient's vascular system for a radioembolization treatment.

2.1.5 Pre-treatment imaging

The screening procedure for radioembolization starts with angiography. Angiography is a technique that is used to visualize the blood vessels. This is achieved by making two scans using X-ray based imaging techniques. The first scan is regular, but prior to the second scan a radio-opaque contrast agent is injected into the blood vessels. By digitally subtracting the two images, only the blood vessels remain. Angiography is used in radioembolization to determine the optimal position for the catheter and find any vascular abnormalities that could compromise the treatment.

After the angiography, a safety procedure is performed using Tc-99m MAA particles to represent the microspheres. These particles contain the radionuclide technetium-99m, which is a commonly used radionuclide in planar imaging. Tc-99m is a gamma emitter, where the

‘m’ indicates that the radionuclide is metastable. The structure of the particles is provided by macro aggregated albumin (MAA), which is a human albumin aggregate particle that is used to resemble the resin or glass microspheres used for radioembolization. The final distribution of the Tc-99m MAA particles is visualized using planar imaging. This way, the vascular system of the liver can be tested and the amount of unwanted dose to the non-tumor tissue, as well as the amount of lung shunting can be determined. In this process, the assumption is made that the Tc-99m MAA distribution correlates well with the final Y-90 distribution after treatment. Recent research has however shown that this assumption might be incorrect [21].

In some hospitals, a SPECT/CT scan is also made of the Tc-99m MAA distribution. SPECT/CT has the advantage that it is much better at localizing extrahepatic depositions [22, 23]. If no contraindications are found in the pre-treatment imaging, the patient is found to be suitable for injection with yttrium-90 microspheres. The catheter position determined during the angiography will be used to administer the spheres.

2.1.6 Post-treatment imaging

After the radioembolization treatment, some hospitals evaluate the treatment efficacy using quantitative PET imaging. Treatment evaluation can for instance be used to investigate the minimal effective tumor dose [24, 25]. The PET imaging is based on the existence of a very small branching ratio of β^+ -decay to the zirconium-90 ground state. In these decays, yttrium-90 decays to an excited state of zirconium-90 while emitting an electron, as is shown in Figure 2.3. Internal pair production causes a positron and an electron to be emitted while the zirconium-90 decays to its ground state. This transition is however very rare and occurs in only 32 of every 1 million decays of yttrium-90 [26]. Using this transition a post-treatment yttrium-90 PET scan can be made. Treatment efficacy is then assessed by determining quantities like the microsphere distribution, T/N ratio and tumor dose [6].




Figure 2.3: The decay scheme of yttrium-90, showing the large β^- -branch that is used for treatment and the much smaller branch of the transition containing internal pair production. The latter gives rise to the positrons employed in post-radioembolization PET imaging [26].

2.1.7 PET dosimetry based treatment planning

Currently, empiric methods are used to calculate the amount of activity that needs to be administered in a radioembolization treatment. These methods take into account limited patient-specific factors and have resulted both in ineffective patient doses, as well as 20 procedure related deaths in one institution [27]. A few cases have been presented in which quantitative PET imaging is used for activity calculation, which has the benefit of taking into account patient specific physiology such as T/N ratio [7, 8]. In these cases, activity calculation is based on a quantitative PET scan that was either acquired after a previous treatment, or after part of the activity of the current treatment was already administered for the purpose of activity calculation. Activity calculation based on quantitative PET can then be used for the planning of future treatments or for planning the second part of the current treatment, using patient specific physiology.

2.2 Positron Emission Tomography (PET)

Radiologists specialize in using medical imaging techniques for diagnosis, treatment planning and image-guided interventions. Many different imaging modalities exist, such as CT, MRI and PET. All of these modalities rely on different physical principles for image formation, such as nuclear decay or magnetic fields. They also serve different applications, such as anatomical or functional imaging. *Anatomical* imaging is used to visualize the internal structure of a patient, such as the locations of organs. *Functional* imaging is used to visualize functional information about organs or tissue, such as blood flow or glucose metabolism.

Positron Emission Tomography (PET) uses *radiotracers*, which consist of biologically active molecules that are labelled with a positron emitting radionuclide. The radiotracer is administered to the patient via injection, inhalation or ingestion. The radiotracer will then spread through the body and naturally accumulate in locations that take up a lot of the biologically active molecule. A commonly used radiotracer for diagnostic PET is *fluorine-18-fludeoxyglucose* (F18-FDG), which consists of the radionuclide fluorine-18 and the glucose analog fludeoxyglucose. This radiotracer will accumulate in locations with a high glucose metabolism, such as the brain or a tumor. The final radiotracer distribution directly reflects biological activity, which in the case of F18-FDG is the glucose metabolism. This distribution can be visualized using tomographic image reconstruction.

PET serves its main applications in oncology, such as the diagnosis and staging of cancer [2]. Other applications include the diagnosis of neurodegenerative diseases, dementia, epilepsy, neurodevelopmental disorders and psychiatric disorders [28].

2.2.1 Positron annihilation

Imaging of the radiotracer distribution is based on the annihilation of positrons that were emitted in β^+ -decay. The positrons will interact with the surrounding tissue and lose their kinetic energy through the ionization and excitation of atoms. After losing sufficient kinetic energy, the positron can *annihilate* upon encountering an electron, releasing two gamma photons. Annihilation is schematically depicted as a Feynman diagram in Figure 2.4.

In the annihilation process, all positron and electron rest mass is converted into energy (the gamma photons) according to the relativistic mass-energy relation $E^2 = m^2c^4 + p^2c^2$, with m the mass of the particle, c the speed of light in vacuum and p the momentum of the particle. Due to momentum and energy conservation, the two gamma photons will be emitted under a 180° angle and will both have an energy equal to the electron rest mass ($511 \text{ keV}/c^2$).




Figure 2.4: A positron annihilates with an electron, resulting in two gamma photons. The horizontal axis represents time.

2.2.2 Photon interactions

Scatter, absorption and attenuation

The two gamma photons that were released in the annihilation process interact with various types of tissue before they reach the detector. When interacting with tissue, or matter in general, a photon can either be *absorbed*, *transmitted* or *scattered*. When a photon is scattered, its direction of motion changes. These three effects can originate from various physical principles, which will be discussed in a moment.

Absorption and scatter both cause photons to be *attenuated*, i.e. be removed from a beam of photons. For a narrow mono-energetic beam of photons the attenuation is known to behave exponentially as $I(x) = I(0)e^{-\mu x}$, with I the beam intensity, μ the linear attenuation coefficient and x the distance in the medium.

These effects of scatter and absorption are caused by the photons interacting with tissue.

Photons can interact with tissue through three types of interactions: the photo-electric effect, scatter and pair production. As pair production is not possible at the photon energies that are relevant for medical imaging, only the photo-electric effect and scatter will be discussed.

Photo-electric effect

In the *photo-electric effect*, the incoming photon interacts with an orbital electron. This electron fully absorbs the photon energy, causing it to be emitted and leaving a vacancy. For this process the photon energy must at least be equal to the binding energy of the electron. All excess energy is converted into electron kinetic energy. The repopulation of the vacancy that was created causes a characteristic X-ray or Auger electron to be emitted.

Compton scatter

Different types of scatter exist, of which Compton scatter is the most relevant for medical imaging and will thus be discussed here. In *Compton scattering* the incoming photon interacts with a valence electron. The direction of the photon is diverted and part of its energy is absorbed by the electron. Figure 2.5 shows the Feynman diagram that describes the process of Compton scatter.

In 1923, an experiment conducted by A.H. Compton showed that the energy of a photon that is scattered from a particle at rest is shifted. A photon with initial energy E_0 leaves this process with scattered energy

$$E_{SC} = \frac{E_0}{1 + \frac{E_0}{m_e} (1 - \cos(\theta))}, \quad (2.2)$$

with m_e the electron mass and θ the scatter angle [29]. For the interested reader, this formula can also be derived by treating light as a particle of zero rest mass with energy given by Planck's equation and applying the laws of conservation of (relativistic) energy and momentum [30].




Figure 2.5: A photon can interact in tissue with orbital electrons through Compton scatter, resulting in energy loss and a change of direction for the photon.

Photo-electric effect vs. Compton scatter

Figure 2.6 shows the likelihoods of the photo-electric effect and Compton scatter as a function of photon energy for absorbers of different atomic numbers. For the relevant range of atomic numbers ($Z \ll 80$) the Compton effect clearly dominates for photons with an energy of 511 keV.




Figure 2.6: *Photons can interact with matter via the photoelectric effect, Compton scattering and pair production. For 511 keV photons the Compton effect dominates [31].*

2.2.3 Detector system

After the two gamma photons traversed through tissue and possibly interacted through Compton scatter and the photo-electric effect, as described in the previous section, they reach the PET scanner. A PET scanner consists of several rings of detectors surrounding the patient. Figure 2.7 shows an example of a PET scanner, showing its cylindrical geometry. Each detector in a ring has a *scintillation crystal* that is read out by a *photomultiplier tube*.




Figure 2.7: *The Siemens Biograph TruePoint PET/CT scanner [32]. Patients lie on the bed inside the cylindrical scanner.*

Scintillation crystal

A gamma photon that reaches a scintillation crystal will be absorbed by the photoelectric effect and releases a photo-electron. This high-energy electron ionizes and excites thousands of other electrons in the crystal, which lose their energy upon recombination or de-excitation. Most of the energy is dissipated as lattice vibrations (i.e thermal energy). A part of the energy is however released in the form of visible photons. The amount of light that is produced in a scintillation crystal is proportional to the energy that was deposited by the gamma photon. These visible photons are then transmitted to the photomultiplier tube.

Commonly used materials for scintillation crystals of PET scanners are bismuth germanium oxide (BGO) and lutetium oxyorthosilicate (LSO). BGO has a high detection efficiency at a photon energy of 511 keV. LSO is slightly less efficient than BGO at 511 keV, but may sometimes be preferred over BGO if the counting rate is high, fast timing information is needed or small scintillator elements are required. The LSO crystal is however expensive to grow [28].

Photomultiplier tube (PMT)

A photomultiplier tube (PMT) consists of a photocathode, several dynodes and an anode. First, the visible photon reaches the photocathode. Due to the photoelectric effect, the photon is absorbed and an electron is emitted from the cathode. Due to a positive potential, this electron is accelerated and focused onto the first dynode, where it causes two to five low energy electrons to be released. These electrons are accelerated to the second dynode, which is held at a more positive potential than the first, and multiplied as well. This process keeps repeating itself for up to 14 dynodes in a PMT. Overall, amplification factors of up to 10^9 can be achieved, causing a current pulse in the anode that is easily detectable [33]. Figure 2.8 schematically depicts the working principles of a PMT.




Figure 2.8: A photomultiplier tube (PMT) converts a visible photon coming from a scintillation crystal into an electron. The electron is then multiplied up to a factor of 10^9 by the dynodes, producing a measurable current pulse [34].

Annihilation coincidence detection

If analysed correctly, these measurements of individual photons in a detector can be used to reconstruct the annihilation site. The first step is to determine which photons might originate from the same annihilation event. Since annihilation events produce two gamma photons under an angle of 180° , the simultaneous measurement of two photons in two opposing detectors indicates a positron annihilation event.

Each measured photon therefore gets a time stamp, typically with a resolution of 1 to 2 nanoseconds. Photons that were detected within a certain *timing window*, typically 6 to 12 nanoseconds, are considered to originate from the same annihilation event. This timing window is necessary to allow for different transmit times in the electronics, different travel distances (and thus different arrival times) for the two photons and the finite timing resolution of the detectors [28].

Line of response (LOR)

If two photons are considered to be coincident, the annihilation event must have taken place somewhere along the path between the two detectors that measured these photons. This path is called the *line of response* (LOR) and is schematically depicted in Figure 2.9. The PET scanner does not just consist of two opposing detectors, but of a whole ring of detectors such that LORs in all directions can be measured simultaneously. LORs with different directions will intersect at the annihilation site and can thus be used for image reconstruction.

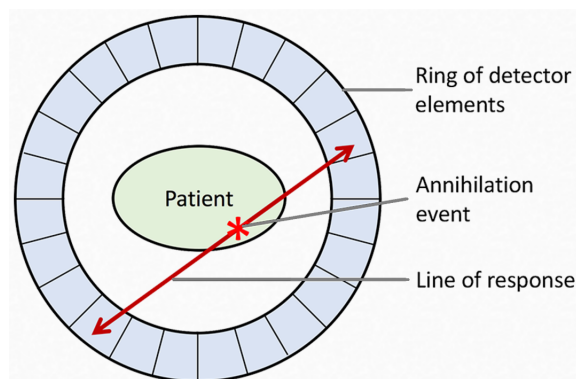


Figure 2.9: *The coincident measurements of two annihilation photons in opposing detectors allows the reconstruction of a line of response, along which the annihilation event must have taken place [35].*

2.2.4 Hybrid PET/CT scanners

PET scanners can be combined with anatomical modalities, such as MRI or CT, to profit both from functional and anatomical imaging. In *PET/CT*, a PET and a CT scan are acquired inde-

pendently but directly after each other without moving the patient. By overlaying information from both scans, zones of high biological activity indicated by the PET scan can be located in the body using the anatomical detail of the CT scan. Figure 2.10 shows an example of separate PET and CT scans, as well as the result of superimposing these scans. High biological activity does not necessarily indicate a functional anomaly. In this figure for instance, the lower black spot is the bladder which lights up as all of the radiotracer that was not taken up elsewhere accumulates in the bladder. Besides localization, the CT scan also contributes in correcting for the aforementioned photon interactions that cause the annihilation photons to be attenuated.




Figure 2.10: *Left: a CT scan showing anatomical details. Middle: a PET scan showing in black zones of high biological activity. Right: the PET scan superimposed on the CT scan, showing the locations of these active zones [36].*

2.3 Image reconstruction

The LORs that were found using annihilation coincidence detection can be regarded as projections. By combining projections under several angles a 3D object can be reconstructed. As a PET scanner consists of a ring of detectors and measures coincidences under all angles, this information can be used to reconstruct a 3D image.

Mathematically, the projection data is described as a Radon transform. By applying the inverse Radon transform to the projection data, an image of the original object can be reconstructed. The Radon transformation is exact, so theoretically the reconstructed image exactly resembles the underlying object.

However, in practice physical effects in the patient and the PET scanner such as noise, scatter and absorption (see Section 2.2.2) influence the projection data. Due to these effects, the projection data is not a perfect Radon transform of the patient that was imaged and the

inverse Radon transform will thus not yield a perfect reconstruction of this patient. Therefore, in practice images are not reconstructed using the inverse Radon transform.

In fact, these physical effects are too complicated to incorporate analytically into image reconstructions. Therefore, images are instead reconstructed using iterative reconstruction algorithms. All of the aforementioned physical effects are taken into account by explicitly modeling them in the iterative reconstruction.

This section first discusses projection (i.e. a PET measurement) and also its inverse, back-projection, in more detail. These concepts will then be of use for the discussion of iterative reconstruction algorithms later in this section.




Figure 2.11: A projection $p(s)$ is defined as the line integral of the image density over a path perpendicular to the detector at detector position s (dashed lines). In this image, only the light-gray disk had a non-zero (homogeneous) image density [37].

2.3.1 Projection

In projection, a three-dimensional volume is reduced to a two-dimensional volume by taking line integrals over one of the directions. Structures that lie in the direction over which the projection is integrated will thus be superimposed. An example of a projection is an X-ray image.

Mathematical description

Mathematically, a projection can be described as a *line integral* over the image density, also called a *Radon transform*. In the case of a two-dimensional image, integrating over one direction will result in a one-dimensional image and the detector is thus one-dimensional (i.e. a line). For a detector that is placed under an angle θ , for each position s a line integral $p(s, \theta)$ of the image density $f(s, \theta)$ is computed over the line perpendicular to the detector at that position.

Mathematically, this can be expressed as

$$p(s, \theta) = \int_{-\infty}^{\infty} f_{\theta}(s, t) dt, \quad (2.3)$$

with f_{θ} the density distribution rotated clockwise by an angle θ (instead of rotating the detector counter-clockwise) and t a parametrization of the line that is integrated over [37]. Figure 2.11 illustrates this principle using two two-dimensional disks.




Figure 2.12: Projections $p(s, \theta)$ (left figure) for all detector positions s and all detector angles θ can be combined in a sinogram (right figure). A point source traces out a sinus in a sinogram [37].




Figure 2.13: The PET data that is measured along LORs with the same angle θ is all stored in a column in the sinogram.

2.3.2 Sinogram

By acquiring projections at all angles θ , a full set of projection data $p(s, \theta)$ is obtained. One commonly used way to display such a full set of projection data is in the form of a matrix in the (s, θ) -coordinate system called a *sinogram*. In such a sinogram, each column represents a projection $p(s, \theta)$ under a specific angle θ . Each element in such a column corresponds to a specific detector position s and is given by the value $p(s, \theta)$. Projections under different angles are stacked together to form the matrix. Figure 2.13 illustrates where the PET data is stored in the sinogram, depending on the angle of the LOR along which the data was measured.

Each pixel in the sinogram is usually represented on a gray scale, with black representing the minimum voxel value and white the maximum value. Figure 2.12 shows an example of a sinogram, in this case for a point source. The sinogram was named after the sinusoidal pattern that a point source traces out in projection space. Figure 2.14 shows an example of a sinogram of non-point source.




Figure 2.14: *Left: a Shepp-Logan phantom. Right: a (noise-free) sinogram of the Shepp-Logan phantom. Pixels are displayed on a gray-scale, with black the minimum pixel value and white the maximum value.*

Projections and sinograms in PET

We already stated that a measurement of a PET scanner is a projection. A detector pair measures photons from annihilation events anywhere on the LOR connecting them without localizing the exact event location. The measured coincidences are thus superimposed for each LOR. All parallel LORs have the same ‘detector’ angle (in PET the detector is not actually rotated as it is capable of measuring all angles at once). Each parallel LOR has however its

own ‘detector’ position, as they are spatially separated. Each LOR therefore corresponds to one sinogram element $p(s, \theta)$.

2.3.3 Backprojection




Figure 2.15: In a backprojection, the number of coincidences measured in each pair of detectors is set back along the line of response. Backprojections (indicated in red here) will intersect in annihilation sites.

The inverse process of projection is backprojection. Backprojections can thus transform a projection back to image space and hence be used for image reconstruction. Since projections are line integrals, the individual contributions from each pixel on that line are summed up. The backprojection disentangles these contributions by assigning to each voxel in an LOR the total number of photons that was measured along that LOR. Figure 2.15 schematically depicts a typical backprojection.




Figure 2.16: Left: without TOF information the projection is backprojected along the entire LOR. Right: a backprojection with TOF information can be constrained to only part of the LOR. [38].

The mathematical description of a backprojection is the *inverse Radon transformation*, given by

$$b(x, y) = \int_0^\pi p(s, \theta)|_{s=\vec{x}\cdot\vec{\theta}} d\theta, \quad (2.4)$$

Note that the Radon and inverse Radon transformations are exact. Theoretically, using infinite integrals and infinitely small step sizes, no information is lost by applying both transformations sequentially.

Time-of-Flight

The conventional backprojection can be improved by incorporating the arrival times of the two photons. This is called *time-of-flight* and it is used to estimate the emission site of the two photons along the LOR. Given the arrival times t_1 and t_2 of the two photons, the distance of the emission site relative to the midpoint between the two detectors is given by

$$\Delta d = \frac{\Delta t \cdot c}{2}, \quad (2.5)$$

with $\Delta t = t_2 - t_1$ and c the speed of light.

In practice, the arrival times are only known up to a few hundred picoseconds, which allows localization to only within a few centimeters [28]. TOF therefore can be used to improve the accuracy of reconstruction algorithms by constraining the backprojection to only a few centimeters. Figure 2.16 illustrates the constrained backprojection.

2.3.4 Iterative reconstruction

Projections and backprojections are part of the process of image reconstruction. As mentioned in the introduction of this section, it is analytically not possible to take into account physical effects, such as absorption, scatter and noise in a reconstructions using backprojection or the Radon transform (see Section 2.2.2). Therefore, images are instead reconstructed using iterative reconstruction algorithms that can explicitly model these physical effects in the reconstruction.

The basic idea of iterative reconstructions algorithms is to iteratively find the image that best compares with the measured PET data. The algorithm works as follows: choose an initial guess, project this guess to a sinogram, compare this sinogram with the measured sinogram and, based on their comparison, update the image guess. The comparison of the sinogram of the guess and the measured sinogram is quantified in an *update error*. Iterative reconstruction can then be mathematically described as

$$\mathbf{x}^{next} = \mathbf{x}^{current} \cdot \frac{E}{N}, \quad (2.6)$$

with $\mathbf{x}^{current}$ the current image guess, E the *update error* resulting from the comparison, N a normalization factor and \mathbf{x}^{next} the resulting new image. The form of the update error E and normalization factor N depends on the specific iterative reconstruction algorithm that is used.

By iterating through this process the image will compare better to the measurements with every iteration. Figure 2.17 schematically depicts the principles of an iterative reconstruction algorithm. The first guess does not need to resemble the imaged object, but can for instance be

homogeneous. All of the voxel values of the first image guess can then for instance be chosen to be 1.

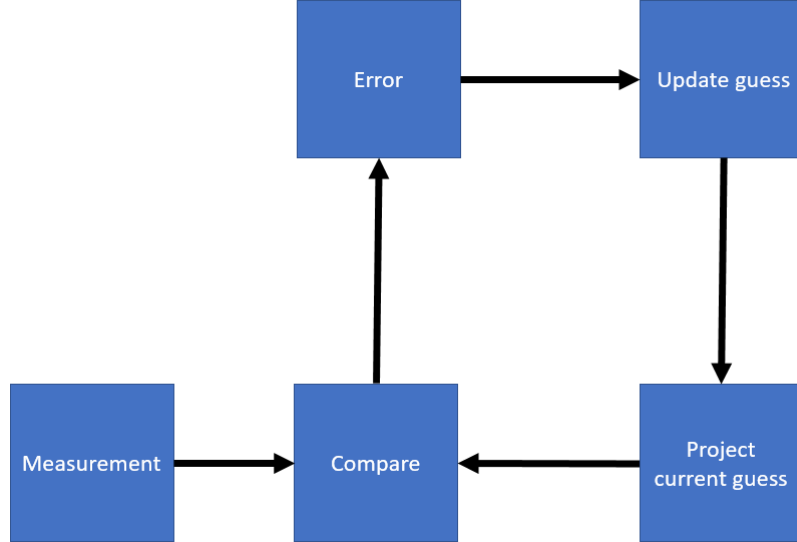


Figure 2.17: *Schematic depiction of an iterative reconstruction algorithm. All relevant physical effects, such as absorption, scatter and noise, are modeled in the projection of the guess.*

Physical effects such as scatter, absorption and noise are modeled in the projector that converts the guess to a sinogram, which is then compared to the measured sinogram. This way, if modeled correctly, both the sinogram of the guess and the measured sinogram contain the same physical effects.

A reliable and robust iterative reconstruction algorithm is *Maximum Likelihood Expectation Maximization (MLEM)*. This algorithm will not be derived here, but for the interested reader different derivations of the algorithm exist, such as the original statistical derivation of Shepp and Vardi [39] and the optimization transfer approach by De Pierro [40]. In the case of MLEM, the update error E and normalization N are given by

$$E = p^{-1} \frac{M}{p(\mathbf{x}^{current})} \quad (2.7)$$

$$N = p^{-1}(\mathbf{1}), \quad (2.8)$$

with p the projection operator, p^{-1} its inverse and $\mathbf{1}$ a vector with elements of 1's. The projection operator p contains modeling of all relevant physical effects, such as scatter, absorption and noise, in the patient and the PET scanner.

2.3.5 PET modes

A PET scanner consists of several rings of detectors, that are all centered around the patient. As the coincident gamma photons “fly off” in all directions, coincidences cannot only be measured

by pairs of detectors in the same ring, but also by pairs of detectors that are in different rings. Taking into account different parts of all of these possible coincidences gives rise to the different *modes* of PET: 2D and 3D.

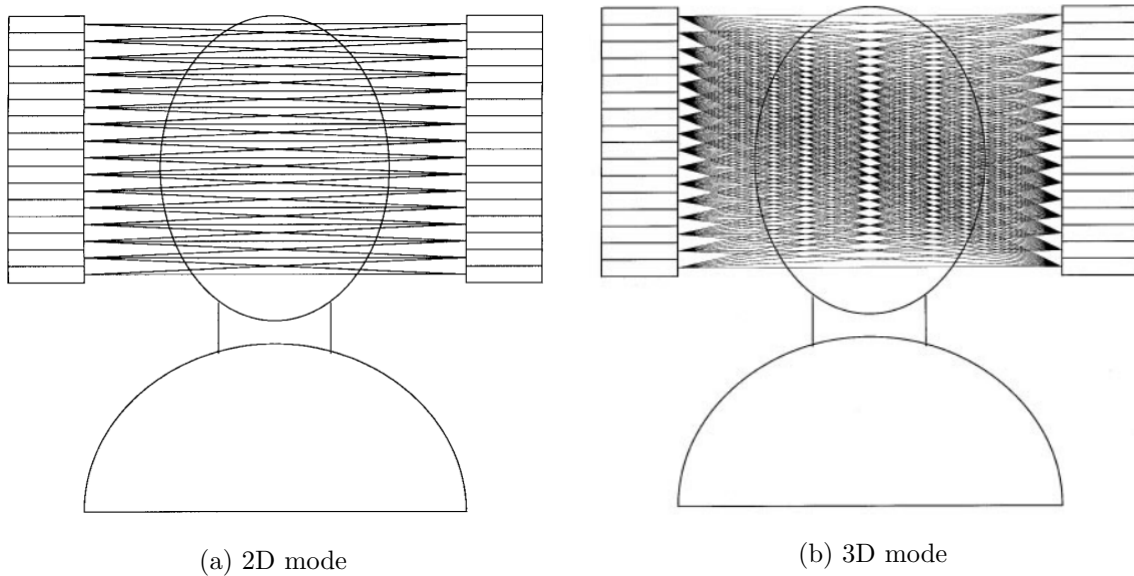


Figure 2.18: *PET modes.* a) In 2D PET, all of the detector rings are employed, but only coincidences within one ring, or close neighbouring rings, are used for reconstruction. b) In 3D PET, all detector rings and all possible coincidences between these rings are used for reconstruction. [41]

2D

In 2D PET, all of the detector rings are used, but only coincidences that were measured within one ring (or close neighbouring rings) are considered for reconstruction. All of the LORs now lie in one of several axial planes. By reconstructing all of these slices independently and stacking them behind each other a 3D volume can be acquired. Since for only a small part of the coincidence pairs both photons will be detected in the same ring (or a neighbouring ring), most coincidence pairs will not be used for the reconstruction in 2D mode.

3D

In 3D PET, coincidence pairs for which the two photons were detected in different rings that are not close neighbours are also employed. Compared to 2D PET, this has both advantages and disadvantages. 3D PET enhances the sensitivity, but a 3D reconstruction is also more complex and time-consuming [41]. Due to the larger number of photons that is detected the number of random coincidences in 3D PET is also larger. Random coincidences are discussed in more

detail in Section 2.3.6. Figure 2.18 schematically depicts the working principles of both scan modes.

2.3.6 Image degrading effects

Images consist of *voxels* (3D pixels) which, in the case of PET, have a size of typically a few mm. Each voxel has a certain intensity, reflecting the amount of relative or absolute radiotracer uptake in that volume. This radiotracer uptake distribution was reconstructed as described in the previous section.

So far it was assumed that all measured coincidences originate from an annihilation event somewhere on the reconstructed LOR. However, this is only the case for *true* coincidences. Several effects will be discussed here that cause the mislocalization of LORs and thus limit spatial resolution and compromise image quality.

Scattered and random coincidences

In a *scattered* coincidence one of the photons was Compton scattered and the assumed LOR will be misplaced. In the case of a *random* coincidence, two photons originating from different annihilation events were measured within the timing window and recorded as a coincident event. This could happen if one photon of each annihilation event was absorbed. Figure 2.19 schematically depicts scatter and random coincidences.




Figure 2.19: *Only in the case of a true coincidence the annihilation event is located on the assumed LOR (left panel). In the case of a scatter or random coincidence the LOR does not intersect the annihilation site (middle and right panel) [28].*

Positron range, non-collinearity and DOI effect

LORs can also be misplaced due to the positron's non-zero *range*, meaning that the positron travels a few mm before it annihilates with an electron. The annihilation location is thus not exactly equal to the radiotracer location.

In addition, the positron can have residual kinetic energy upon annihilation. In that case, due to momentum conservation, the angle between the annihilation photons will not be exactly 180° . This effect, called *non-collinearity*, causes the LOR to be slightly misplaced.

Another effect is the *depth of interaction (DOI)* effect, which is caused by the fact that it is not known at what depth in the detector the photons were measured. Due to this effect, the angle of the assumed LOR will deviate from the angle of the true LOR. Figure 2.20 schematically depicts the effects of the positron range, non-collinearity and the DOI effect on the LOR.

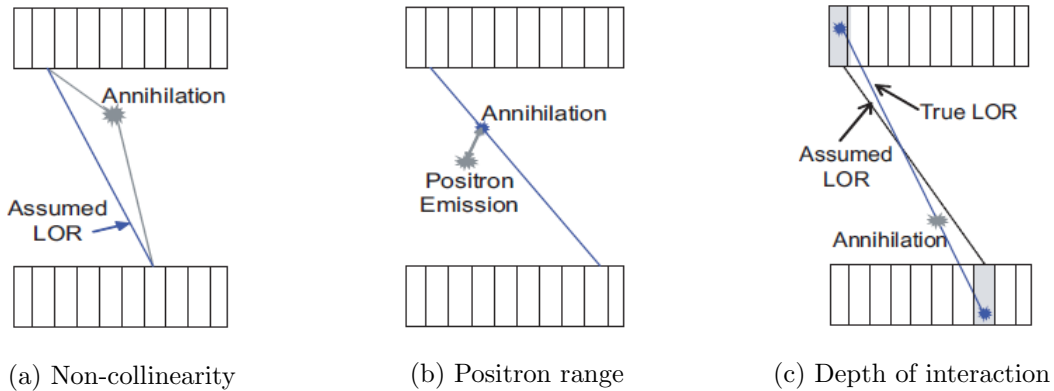


Figure 2.20: *The spatial resolution in PET imaging is decreased by non-collinearity, the non-zero positron range and the depth of interaction effect [42]. Figure 2.19 shows how the spatial resolution can also be decreased by scatter and random coincidences.*

Detector size, noise and respiratory motion

Not only the mislocalization of LORs, but also the finite *detector size* limits the spatial resolution, as structures smaller than the detector cannot be resolved. Moreover, *Poisson noise* due to the statistical nature of radioactive decay causes images to look mottled. Next to these physics and detector system related effects, the effective spatial resolution is also significantly compromised by respiratory motion. PET systems can nowadays reach a spatial resolution of 2-5 mm, but the presence of respiratory motion can drastically decrease the effective spatial resolution, especially in high resolution PET systems. For instance, for a PET scanner with a system resolution of 5 mm, the effective resolution due to a motion amplitude of 10 mm will be 11.25, which is an increase of 125% [43].

2.3.7 Quantitative PET

Without quantification the voxels in a reconstructed PET image only represent relative values. From these qualitative images it can be deduced that some regions in the patient body take up more radiotracer than others, but no quantitative parameters can be extracted from the

image. Quantification of PET images allows the radiologist to use quantitative parameters such as tumor volume and radiotracer uptake for various applications.

However, accurate quantification requires corrections for the relevant (physical) quantitative effects that were mentioned in the previous sections, such as scatter, randoms and crystal efficiency. This section will describe how these effects can be incorporated into image reconstruction, various applications of quantitative PET and the problems posed by the absence of respiratory motion correction in image reconstruction algorithms.

Quantitative reconstruction models

Corrections for quantitative effects, such as scatter, are added to the reconstruction to obtain a quantitative reconstruction model. This section will explain how these corrections can be incorporated. Scattered coincidences are difficult to correct in PET due to its low energy resolution. Scattered coincidences are therefore modeled using single scatter simulation and then subtracted from the measured data.

Randoms correction can be achieved by delaying the arrival times of the photons that were measured by one of the detectors. The delay should be significantly larger than the timing window. Due to the delay, the two photons of a true coincidence will not be measured within the timing window and will not be recorded as a coincidence. But one of the photons can accidentally be recorded as coincident with a photon of another annihilation event within the timing window. As all true coincidences were removed, the number of coincidences determined with the delayed window technique gives an estimate of the number of random coincidences. These can then be subtracted from the total number of measured coincidences .

Attenuation can be corrected for by acquiring a CT scan prior to the PET scan, as a CT scan quantifies the amount of attenuation for each LOR. Each element of the sinogram data corresponds to an LOR and can thus be corrected for its attenuation using a multiplicative correction factor $e^{-\sum_i \mu_i L_i}$, with μ_i the attenuation factor for tissue type i and L_i the distance that was traversed in that tissue.

Detector system related effects can be determined by measuring a point source in the PET scanner. The measured number of counts for each LOR will deviate from the expected number counts. This deviation is caused by detector related effects such as crystal efficiency. Each LOR can then be individually corrected for these effects in sinogram space by using the measured deviation for that LOR.

Applications

Quantitative parameters such as tumor volume and radionuclide uptake are of relevance for various clinical applications. Tumor volume is for instance of high importance for radiotherapy, where accurate delineation of the tumor is required to avoid irradiation of healthy tissue, as

well as to improve tumor coverage [3, 4].

Other applications include radionuclide therapies, such as radioembolization. Current empiric activity calculation methods for radioembolization take into account limited patient-specific factors and have resulted both in ineffective patient doses as well as 20 procedure related deaths in a single institution [27]. If quantitative PET would be used in the activity calculation, the patient-specific microsphere distribution could be determined as well as the resulting tumor dose [6].

As mentioned before, cases have been presented in which PET dosimetry based on either a previous treatment or the first part of the current treatment was used to adapt future treatment planning [7, 8]. This way, patient-specific information can be taken into account by using quantitative PET.

The problem of respiratory motion

As quantitative PET serves many applications and so many corrections have already been incorporated in quantitative reconstruction models for PET, this sounds as a very promising field. There is however one problem with these quantitative reconstruction models: patient breathing is not taken into account. Although all the physics and detector related effects are modeled, this is not yet the case for respiratory motion.

The diaphragm can move as much as 1 to 6 cm during an acquisition [10]. As a result, organs in the thorax and abdomen will also undergo motion, which complicates imaging of the thorax and abdomen. Organs can undergo translations, rotations and deformations due to respiratory motion. Especially for the lungs, the amount of deformation cannot be neglected. The motion of the liver however, which will be the focus of this study, mainly consists of translations in the head-foot direction. Studies have reported a mean liver tumor motion amplitude of 10-25 mm in the head-foot direction during shallow breathing. During deep breathing amplitudes of up to 80 mm can be reached [44].

This poses a major problem, as respiratory motion causes image artifacts that can result in reduced resolution, reduced tumor detectability, misalignment between the PET and CT and attenuation correction artifacts [45]. An important artifact is blurring of the tumor, which results in overestimation of the tumor volume and underestimation of radionuclide uptake. Liu et al. [10] found in their study a mean tumor volume overestimation of 130% and a mean tumor uptake underestimation of 28% for lung and liver tumors due to respiratory motion. It should be noted, however, that the amount of over- or underestimation depends on many parameters, such as tumor size and motion amplitude, and will thus differ per patient and tumor.

Several methods have been proposed for respiratory motion correction in PET, other imaging modalities, and applications outside image acquisition such as radiotherapy. These will be discussed in the next chapter.

Chapter 3

Respiratory motion compensation

Respiratory motion does not just pose a problem in PET imaging, but also in other imaging modalities and in treatments such as radiotherapy and liver ablation [12, 46]. Various solutions to the problem of respiratory motion have been proposed. This chapter will give a rudimentary overview of motion correction techniques that have been presented in the literature.

3.1 Breath-hold

Breath-hold is a technique that solves the problem of respiratory motion by requiring the patient not to breath during the scan. Breathhold is often used for CT scans that have an acquisition time of up to 30 seconds, such as the CT that is acquired independently in PET/CT.

A drawback of breathhold is that the patient will not hold his or her breath in exactly the same way in different acquisitions. This complicates the comparison of different scans, for instance for the purpose of treatment evaluation. Moreover, for some patients with lung problems it is not possible to hold their breath for such a long time. As a whole-body PET scan takes 15-30 minutes, breathhold is not a suitable respiratory motion handling technique for PET.

3.2 Respiratory gating

3.2.1 Description

In *respiratory gating*, the PET data is manipulated before the reconstruction to contain less motion. After the acquisition, the PET data is binned into several *gates* based on the respiratory phase of the patient. As the internal motion and thus the respiratory phase of the patient cannot be measured directly, a surrogate is used. This surrogate signal can for instance originate from a pressure sensitive elastic belt that is worn by the patient during the PET acquisition. This surrogate signal is then assumed to correlate with the internal motion. Only the subset of the

PET data that corresponds to one of the gates is used for the image reconstruction. This subset of the data contains only a fraction of the total motion that was present in the entire data set.

Figure 3.1 shows an example of a simplified, sinusoidal surrogate signal that was gated retrospectively using five gates. In respiratory gating, only the data from one of the gates will be used for reconstruction. For instance only the data from the fifth gate is used, which was acquired during end-expiration. It is also possible to use all of the gates for the reconstructed image, as will be explained in Section 3.3




Figure 3.1: *Using retrospective gating a simplified, sinusoidal respiratory signal was divided into five separate gates. Only the PET data that falls into one of the gates will be reconstructed. For instance only the PET data corresponding to the first gate, which was acquired during end-inspiration.*

3.2.2 Advantages

Respiratory gating reduces the amount of motion in the reconstructed image and hence reduces motion artifacts. This is noticeable in the tumor volume and in the maximum Standardized Uptake Value (SUV_{max}), which is a measure that is proportional to the maximum radiotracer uptake per voxel. Studies have reported an increase in SUV_{max} due to respiratory gating of 30% to 83% and a decrease in tumor volume of 45% [47–50].

3.2.3 Disadvantages

Respiratory gating may reduce motion artifacts, but at the cost of lower count statistics, thus increasing image noise. Due to this trade-off, the increase in image noise might unintentionally decrease the image quality. For realistic patient surrogate signals, the data of each gate will contain a different amount of noise and a different amount of residual motion. The trade-off between image noise and residual motion determines which gate is best suited for image reconstruction.

Moreover, the data that is used for reconstruction is not in the exact same respiratory

phase, but rather an average over a small variation of respiratory phases due to the width of the gate. This causes motion to not fully be removed from the image. Besides, the patient has to wear a respiratory belt or some other device that indicates the respiratory phase, which can be uncomfortable for the patient. For the clinical staff the addition of such devices in the clinical protocol costs a lot of extra time. These devices are often difficult to use and can give problems if the breathing dynamics of a patient deviate too much from the standard, for instance if the breathing amplitude is much larger.

In addition, respiratory gating relies on the assumption that patient breathing is stationary, but it is known from radiotherapy that breathing is patient specific and highly variable. A patient's breathing pattern can change both from treatment to treatment as well as within one treatment. Changes include amplitude, frequency and baseline shifts [11] and can be caused for example by the patient switching from thoracic to abdominal breathing [12]. Especially in the case of non-stationary motion, such as baseline and amplitude drifts, the difference between the minimum and maximum of the surrogate signal will be large. This causes the gates to be stretched. Due to their increased width the gates will contain more residual motion, which will reduce the image quality. This effect is illustrated in Figure 3.2, where each of the gate contains almost half of a respiratory cycle. In the case of constant breathing dynamics however, each gate contained only one fifth of a respiratory cycle (see Figure 3.1).

Another problem with breathing variability is that data of completely different respiratory phases can end up in the same gate. Amplitude drifts can for instance cause mid-inspiration data from one cycle and end-inspiration data from another cycle to have the same surrogate signal value. This data will thus end up in the same gate. This data is then assumed to originate from approximately the same respiratory phase, but in this case that assumption is incorrect. This effect is illustrated in Figure 3.2, where the maximum of the first respiratory cycle, mid-inspiration data of the second cycle and the minimum of the third cycle all fall into the third gate. This gate thus contains both end-inspiration, mid-cycle and end-expiration data that is all averaged.

Currently a common solution to deal with breathing variability is to control the variation. Using audio or visual feedback the patient can be coached to breathe more regularly. However, for some patients coaching is not suitable due to their medical condition or because they have difficulty following the coaching instructions [13, 14].




Figure 3.2: An example of retrospective gating of a surrogate signal in the case of variable breathing dynamics. Each gate now contains almost half of a respiratory cycle, rather than one fifth of a respiratory cycle.

3.3 Registration of gated images

3.3.1 Description

Another respiratory motion correction technique is *registration of gated images*, in which all of the data is divided into a number of gates that are reconstructed independently. These reconstructions will not be in the same average respiratory phase, but registration can transform the images to a common respiratory phase. This transformation can include for instance translations, rotations and deformations. The resulting images can then be combined into one single image.

3.3.2 Advantages

An advantage of registration of gated images is that more statistics are used for the final image, hence increasing the ratio between the mean radiotracer uptake in the tumor and the noise variation in the background (SNR). Registration of gated images also decreases the smearing of the tumor. Thorndyke et al. [51] found in a similar approach an increase in SNR of 15% and a decrease in tumor volume of 33% with respect to gated images.

3.3.3 Disadvantages

Although registration of gated images improves the SNR and reduces tumor smearing, the final images still suffer from some of the same drawbacks as respiratory gating. In the case of breathing non-stationarity, the gates will be stretched and the residual motion in each gated image increases. The quality of the gated images in the case of non-stationary breathing is thus suboptimal. Registration only guarantees that these suboptimal images are then registered in the most optimal way. But if the quality of the gated images themselves is already suboptimal,

this will affect the quality of the final image, which inevitably will also contain residual motion to some extent.

3.4 Respiratory motion models

To overcome the drawbacks of respiratory gating and registration of gated images, a new class of techniques has been developed in the past years that uses so called motion models. These motion models are used to *estimate* the motion, but do not describe how the motion is corrected. This Section will start with a description of motion models. This will be followed by a rudimentary literature overview and a discussion of advantages and disadvantages of motion models.

3.4.1 Description

A motion model takes the surrogate signal of a patient as input and then produces an estimate of the internal motion in the patient as output. This process is schematically depicted in Figure 3.3.




Figure 3.3: A motion model takes the surrogate signal s_t as input and then produces an estimate of the internal motion r_t as output

In general the motion model can depend on some parameters that we call \vec{A} here. These parameters can differ for different acquisitions, so they have to be optimized for each acquisition. In the form of an equation the motion model can then be described as

$$\vec{r}_t = \phi(s_t, \vec{A}), \quad (3.1)$$

with ϕ the motion model, s_t the surrogate signal at time t , \vec{A} the motion model parameters and \vec{r}_t the estimate of the internal motion at time t .

Example: linear motion model

This Section will present an example of a motion model. We let the patient wear a pressure-sensitive respiratory belt during the PET acquisition. The pressure that is measured by this

belt can then be used as the surrogate signal s_t . An example of such a belt is shown in Figure 3.4.




Figure 3.4: A pressure-sensitive respiratory belt that can be worn by the patient during the PET acquisition. The resulting pressure can be used as a surrogate for the internal motion.

Figure 3.5 shows in blue the surrogate signal s_t that was measured using this respiratory belt. In orange it shows the internal motion r_t that was present in the patient. There is clearly a correlation between the surrogate signal and the internal motion. In this case, the correlation can be described using a linear motion model with only one parameter A . This linear motion model is given by $r_t = A \cdot s_t$, with r_t the internal motion at time t , A the motion model parameter and s_t the surrogate signal. The model parameter is $A = 1.4$ a.u. in this case.

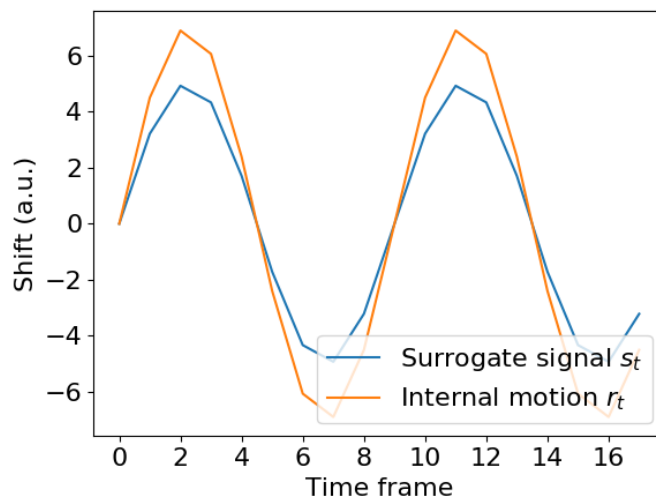


Figure 3.5: Blue: the surrogate signal s_t that was measured using the pressure-sensitive respiratory belt. Orange: the internal motion r_t in the patient. There clearly exists a correlation between the two signals.

Using the model parameter A and the surrogate signal s_t the internal motion r_t can be calculated at each time t . The estimate of the internal motion can then be used for motion correction. The latter can be done by binning the data into small time frames of, for instance, 0.5 seconds. The data of each time frame can then be corrected for the corresponding internal motion r_t . How this correction is performed exactly, can differ per imaging modality.

3.4.2 Literature overview

A limited overview of respiratory motion models ϕ , surrogate signals s_t and optimization methods for the model parameters \vec{A} that exist in the literature will be presented here. For the interested reader, McClelland et al. [12] gives an extensive overview. Motion models have been developed for different applications, which can be divided in image-guided interventions and image acquisition. The image-guided interventions include radiotherapy, cardiac interventions, liver ablation and bronchoscopy [52–55]. In the case of image acquisition, motion models have mainly focused on MRI, PET and CT [53, 56, 56, 57]. Motion models have been developed for different organs, but mostly for the lungs, heart and liver. Some models have also been presented that can be applied to the entire thorax or abdomen [56].

Different types of surrogate data have been used, which can in general be multi-dimensional and even gradient and phase information of the surrogate signal can be employed. Choices for surrogate data include the pressure on an elastic respiratory belt, chest or abdomen displacement measured with an optical tracking system and air flow measured with a spirometer [58–60]. Using some high-resolution modalities, like CT and MRI, it is also possible to use the motion of a clearly visible structure like the diaphragm as a surrogate for organ motion [56, 57]. The choice of surrogate signal affects the accuracy of the motion model, as some types of surrogate data correlate better with internal motion than others. It has for instance been shown that spirometry correlates better with internal motion than abdomen displacement [61].

Motion models that are often used in the literature are linear, polynomial and b-spline models [62–64]. A b-spline is a piecewise defined polynomial function. Examples of more complicated models include Fourier series, neural networks, fuzzy logic and support vector regression [65–67].

A common optimization method for the parameters \vec{A} of a motion model is linear least squares, where the optimal value of the model is a linear function of the parameters [62–64]. Other fitting methods include variants of linear least squares, principal component regression and support vector regression [52, 59, 68].

It should be noted that the choices for the surrogate signal, motion model and optimization method are not independent. The type of surrogate signal can for instance limit the amount of

variation that can be modeled, whereas the type of motion model influences the optimization method that could be used.

Although many motion models have been proposed, only a few have been proposed for PET. These are often specifically proposed for PET/MRI, as the motion model is estimated on the MRI data [56, 69]. In many clinics PET/MRI scanners are however not available. In PET/CT the motion model could be estimated from the CT data by acquiring CT scans at different respiratory phases [70]. This however increases patient dose and this solution is therefore considered to be suboptimal.

3.4.3 Advantages

An advantage of a motion model is that all breathing variability is contained in the surrogate signal. As long as the motion model is valid, it can be used to estimate the internal motion, which will contain the same variability. Moreover, the model can be fitted on the entire data of the PET scan, thus benefiting from all the information about motion dynamics that is available. In this process, it is not necessary to average parts of the data, so the time resolution is higher with respect to gating and registration of gated images. Furthermore, data from completely different respiratory cycles is not averaged, which enables the possibility to model inter-cycle variation, such as breathing non-stationarity. This makes a motion model possibly more accurate than respiratory gating or registration of gated images.

3.4.4 Disadvantages

A disadvantage of motion compensation using motion models is that the assumption is made that there is an underlying model that describes the internal motion dynamics. This model is assumed to hold for the entire duration of the scan. Although studies have been performed to test the validity of motion models, there is no guarantee that a certain model will hold for every patient nor for the duration of a clinical scan [44, 65].

3.5 Research question

The aim of this project is to develop a novel motion correction method for PET that is based on the PET data and can deal with inter-cycle breathing variability. In this method motion will be estimated by explicitly modeling the internal motion dynamics given a surrogate signal. Internal organ motion can consist of translations, rotations and deformations. This study will focus on a motion model for translational motion. This will already enable us to, to a great extent, describe the motion of the liver which mainly consists of translations in the head-foot direction. Such a model might already be accurate enough to serve certain applications that are specific to the liver, such as radioembolization. A quality measure will be assigned to the

motion model, indicating its goodness of fit. For the reliability of this method it is important that problems, such as an inaccurate motion model, will show up in this quality measure.

The goal of this study is to test the feasibility and performance of such a model-based respiratory motion correction method in PET acquisitions of the liver. The results will be compared with respiratory gating and registration of gated images through simulations and experiments. In this study we assume that there is an underlying motion model relating the internal motion to the surrogate signal. We will investigate the behavior of the quality measure in the case that this assumption is incorrect and there is no constant underlying motion model. In this case the quality measure is expected to increase, as the motion model that will be found cannot be suited for the data and motion correction will thus be suboptimal.

Chapter 4

Method

The goal of this study was to develop a novel motion correction method for PET acquisitions of the liver that can deal with breathing variability. This is achieved by modeling the entire motion dynamics as a function of a surrogate signal. Given the perfect motion model parameters for the motion model between the surrogate signal and internal motion, the motion at each time will be exactly known. It would then be possible to perfectly correct for this motion during the reconstruction. This chapter will start by explaining how motion correction can be incorporated in the reconstruction, given a perfect motion model.

Unfortunately, in practice, the motion model is not known. We therefore propose to not only optimize the unknown image, but also optimize the unknown motion model parameters in the iterative reconstruction algorithm. How this is achieved will be explained in the second section of this chapter.

The chapter will finish with a discussion of the simulations and experiments that were performed. For the interested reader, some unsuccessful alternative techniques that were tried before the final approach was established are discussed in Section A.

4.1 Motion correction

As described in the introduction, given a perfect motion model, the image can be perfectly motion corrected during the reconstruction. Motion correction was based on the adaptation of the update error of the regular MLEM algorithm as described in Equation 2.7. This update error is calculated using both the sinogram of the current guess and the measured sinogram. As the measured sinogram contains a different amount of motion for each time frame, the update error was calculated for each time frame independently.

To incorporate motion correction in this update error, the image guess was transformed to contain the estimated amount of motion before it was projected to a sinogram. The measured sinogram was then compared to the sinogram of the *transformed* guess rather than the sinogram

of the *untransformed* guess. This way, given the perfect motion model parameters, both sinograms contain exactly the same amount of motion. The comparison of these sinograms would then no longer be influenced by the presence of motion, but it is only influenced by the quality of the image guess. The update error was then backprojected just as in regular MLEM.

However, as the amount of motion that is now in both the measured and transformed guessed sinogram differed per time frame, the backprojected update errors for the different time frames were also calculated for different respiratory states. Therefore, the backprojected update errors for each time frame had to be transformed back to a common respiratory state, before they could be averaged. The resulting total backprojected error was used to update the image guess. Figure 4.1 schematically depicts how the iterative reconstruction algorithm was adapted to include motion correction.

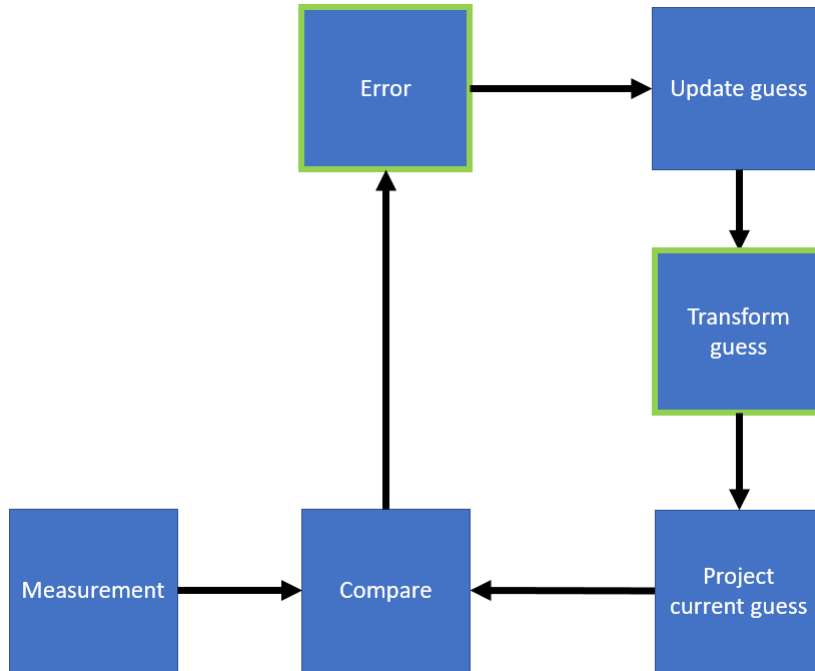


Figure 4.1: *In order to incorporate motion correction in the iterative reconstruction, the image guess was transformed using the motion model before it was projected. Using this transformed sinogram the update error was computed. The adaptations are indicated using green contours.*

This algorithm was assumed to be perfectly correct for respiratory motion, given a perfect motion model, and is summarized in the following expression for the update error

$$E_{\text{update}} = \frac{1}{N} \sum_t T(-\vec{r}_t) p^{-1} \left(\frac{M_t}{\tilde{G}_t} \right), \quad (4.1)$$

with N the number of time frames, p^{-1} the backprojection operator that converts a sinogram to an image, $T(\vec{x})$ a transformation operator that shifts the image by an amount \vec{x} , \vec{r}_t as determined by the motion model in equation 3.1, M_t the measured sinogram in time frame t

and \tilde{G}_t the sinogram of the *transformed* guess, which can be obtained as

$$\tilde{G}_t = pT(\vec{r}_t)G_t, \quad (4.2)$$


with G_t the *untransformed* image guess and p a projection operator that converts the image to a sinogram.

Example: update error calculation


This example will illustrate how the update error is calculated. We will use the same linear motion model as in the example in Section 3.4 and the liver phantom of Figure 4.10. We will assume we know the perfect model parameter to be $A = 1.5$. We can then use this parameter value to adapt the update error E_{update} for a specific iteration in the aforementioned way. We will assume that the sinogram of the measured PET data M was binned into three time frames of 1 second. This gives us three measured sinograms M_t , one for each time frame t . Now for each time frame we will calculate an error and in the end these errors will be averaged to give the total update error. This example will show how the error is calculated for the first time frame.

Figure 4.2a shows the measured sinogram for time frame 1 (M_1). For the first time frame the internal motion is calculated (using the motion model) to be $r_t = 10$ pixels, in the y -direction. The guessed image is then shifted this amount r_t and projected, resulting in the sinogram of the transformed guess for time frame 1 ($\tilde{G}_1 = pT(r_1)G_1$), which is shown in Figure 4.2b. The ratio of M_1 and \tilde{G}_1 is then calculated, as is shown in Figure 4.2c. This ratio is then backprojected to image space and shifted with an amount of motion $-r_t$, the result of which is shown in Figure 4.2d. Now this transformed and backprojected ratio was calculated for the first time frame.


By calculating this quantity for every time frame and averaging them, the total update error can be obtained. This update error has the same shape as the image guess, such that each voxel in the image guess will be multiplied with its corresponding motion corrected error in Equation 4.1.



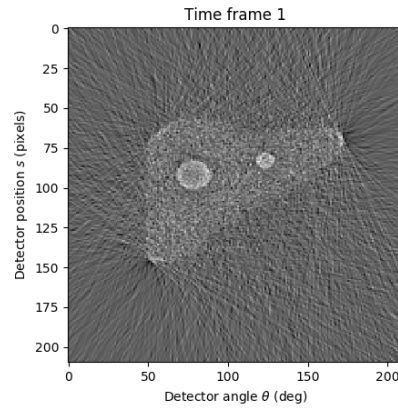
(a) The measured sinogram for time frame 1 (i.e. M_1).



(b) The sinogram of the transformed image guess for time frame 1 (i.e. \tilde{G}_1).



(c) The ratio M_1/\tilde{G}_1 .



(d) The backprojected ratio M_1/\tilde{G}_1 .

4.2 Motion estimation

Motion was estimated by modeling the internal motion dynamics using a motion model. This motion model correlates the internal motion with the surrogate signal. Given the perfect motion model, the image guess can be perfectly motion corrected during the reconstruction using the method as described in the previous section.

However, the motion model is not a priori known and therefore we propose to optimize the motion model parameters as part of the iterative reconstruction algorithm. A better motion

estimation facilitates a better image guess, whereas a better image guess increases the accuracy of motion estimation. Therefore, both the image guess and motion model are optimized in each iteration of the reconstruction algorithm. This section focusses on describing the choices that were made for the surrogate signal, motion model and optimization method.

4.2.1 Surrogate signal

The internal motion of the patient cannot be measured directly, so instead the measurement of a surrogate signal is used. This surrogate signal is assumed to correlate with the internal motion of the patient, and thus indicate the respiratory state of the patient. The surrogate signal can be measured by hardware, such as a pressure sensitive respiratory belt that is put around the patient's abdomen. In the case of PET, the surrogate signal can also be calculated directly from the raw data without loss of accuracy [71]. A benefit of a surrogate signal calculated from the raw data is that it is fully automatic and no additional hardware is required, so the regular scanning protocol does not need to be changed and the clinical staff does not have to perform any extra tasks. It also makes the technique more accessible for hospitals that do not have such hardware. Moreover, using a data-driven surrogate signal it is not necessary to synchronize the surrogate signal with the PET data. For these reasons, the surrogate signal in this study was also calculated from the raw data.

The surrogate signal was calculated using the center-of-mass method of Büther et al. [72], who showed that the movement of the center-of-mass in the direction into the scanner bore correlates with the signal of a respiratory belt and can thus also be used as a surrogate signal, without loss of accuracy.

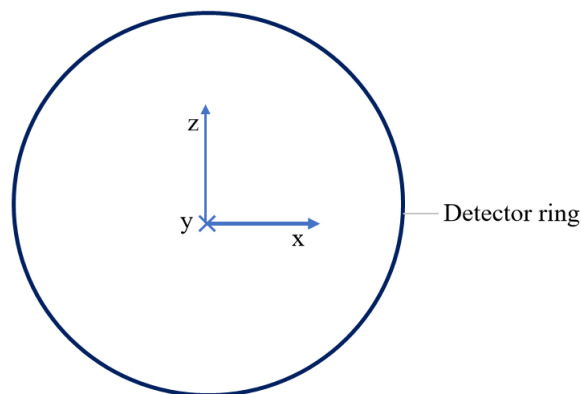


Figure 4.3: *Definition of the coordinate system of the cylindrical PET scanner. The y-direction is defined as the direction into the scanner bore.*

In order to explain the principle behind this method, we will first define a coordinate system for the cylindrical PET scanner as shown in Figure 4.3. To calculate the center-of-mass in the y -direction (into the scanner bore), the total number of counts in the sinogram of each axial

slice is determined. With $N_y(i, t)$ the total number of counts in the sinogram of axial slice i at time t , the center of mass in the y -direction can be calculated as

$$CM_y(t) = \frac{\sum_i i \cdot N_y(i, t)}{\sum_i i}. \quad (4.3)$$

4.2.2 Motion model

A linear motion model was used as many studies have shown its use, such as for tumor tracking in radiotherapy treatments [62, 63]. The linear motion model can be written as

$$\vec{r}_t = \vec{A} \cdot s_t, \quad (4.4)$$

with $\vec{r}_t = (x_t, y_t, z_t)$ the internal motion at time t , s_t the scalar surrogate signal at time t and $\vec{A} = (A_x, A_y, A_z)$ the model coefficients.

In general, a linear motion model could also contain extra coefficients \vec{B} for each direction, the motion model would then become $\vec{r}_t = \vec{A} \cdot s_t + \vec{B}$. However, since this extra term does not depend on time, it does not add to modeling the internal dynamics and it will only shift the entire image by a given amount. As this study focuses on just the liver, the absolute position of this one organ is not relevant and these coefficients were excluded from the study.

4.2.3 Optimization method

For each patient and each different PET acquisition, the optimal model parameters \vec{A} have to be determined. Doing this for every acquisition and every patient allows taking into account breathing variability between different image acquisitions. The optimization of the model coefficients was added at the start of each iterative reconstruction step, as indicated in Figure 4.4.

Each iteration, both the guessed image and the guessed motion model parameters will be updated, resulting in a better image guess and a better guess for the motion model parameters. A good image guess is needed to accurately estimate the motion model parameters, but likewise accurate motion model parameters are needed to estimate a good, motion-free image. Therefore, each iteration both the image and the motion model parameters are improved. Each iteration the current estimate for the model parameters is used for motion correction.

The optimization step in the reconstruction algorithm consists of the minimization of an error quantity that depends on the model coefficients, and is chosen to be minimal for the perfect model coefficients. The next sections will describe the model error and minimization algorithm that were used.

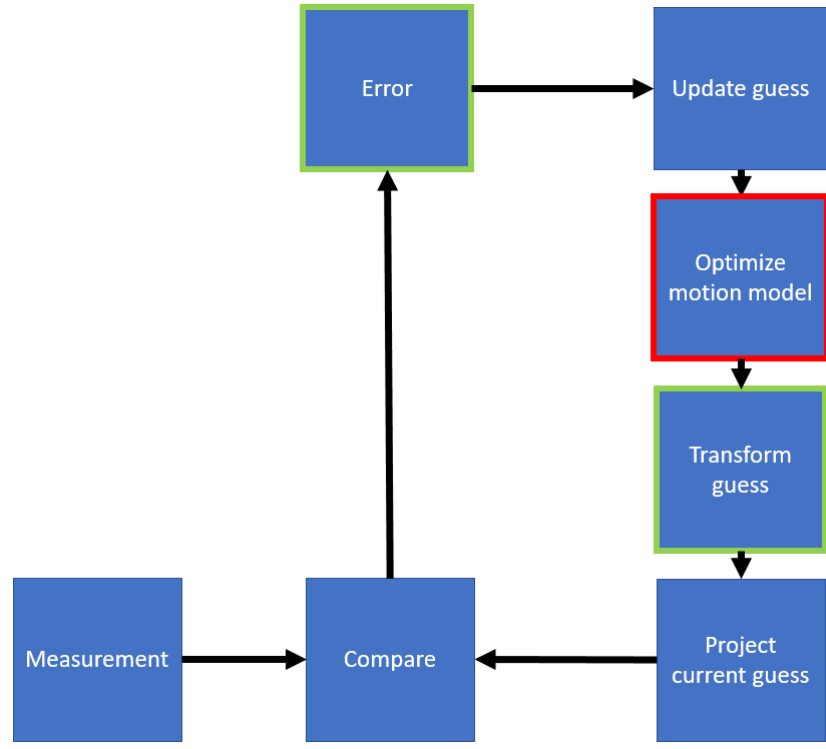


Figure 4.4: In order to optimize the model coefficients \vec{A} , a model optimization step (red contour) was added at the start of each iteration of the reconstruction algorithm.

Model error

A *model error* was constructed that could indicate the goodness of fit of the motion model parameters. This model error was used to optimize the unknown motion model parameters, and is different from the *update error* in Equation 4.1, that was used to optimize the unknown image.

The model error was chosen to be the absolute difference between the measured sinogram and the sinogram of the current *transformed* image guess. The measured data was divided into short time frames, yielding a measured sinogram M_t for each time t . The time-resolved sinograms can then be used to calculate the model error, which can mathematically be expressed as

$$E_{model} = \frac{1}{2N} \sum_{t,i} \text{abs} \left(\frac{M_{t,i}}{\sum_j M_{t,j}} - \frac{\tilde{G}_{t,i}}{\sum_j \tilde{G}_{t,j}} \right), \quad (4.5)$$

with N the number of time frames, $M_{t,i}$ the i 'th element of the measured sinogram in time frame t and $\tilde{G}_{t,i}$ the i 'th element of the current *transformed* guess (see Equation 4.2). The factor $\frac{1}{2N}$ ensures that $0 < E_{model} < 1$.

It is important to be able to quantitatively compare different PET acquisitions, for instance for treatment evaluation. Therefore, both $M_{t,i}$ and $\tilde{G}_{t,i}$ were normalized, such that the model error is influenced by the shape of the sinograms, but not the total number of counts in a sino-

gram. Otherwise, it would be difficult to compare scans of different duration, as the sinograms then contain a different number of counts.

Due to the size of the sinograms of typically tens of thousands of entries computing the model error is currently a very computationally intensive step.

Example: model error calculation

This example will illustrate how the model error is calculated. We will use the same linear motion model as in the example in Section 3.4 and the liver phantom of Figure 4.10. We will first calculate the model error for $A = 1$. Let's say we have a measured sinogram M . We bin the data in three time frames of 1 second. This gives us three measured sinograms M_t , one for each time frame t . Each of these three sinograms is then normalized, by dividing the sinogram M_t by the total number of counts in that sinogram ($\sum_j M_j$). Figure 4.5 shows these normalized sinograms.




Figure 4.5: *The normalized measured sinograms $M_t / \sum_j M_{t,j}$ for all time frames.*

Let's say that the image guess of the current iteration in the reconstruction is given by the image in Figure 4.6. For each time frame, this image guess needs to be transformed by the amount of estimated motion r_t , which is the same amount of motion that the measured sinogram in that time frame is estimated to contain. Given the surrogate signal s_t and the motion model ϕ , the motion model parameters imply an estimate for the internal motion r_t for each time frame (see Equation 3.1). For each time frame, the image guess is now shifted by an amount r_t , resulting in three shifted guesses $T(r_t)G$. These shifted guesses are then projected to a sinogram $\tilde{G}_t = pT(r_t)G$. These sinograms of the transformed guesses \tilde{G}_t are then normalized to one and are shown in Figure 4.7.




Figure 4.6: *The current image guess, that is used to calculate the model error.*




Figure 4.7: *The normalized sinograms of the transformed guess $(\tilde{G}_t / \sum_j \tilde{G}_{t,j})$ for all time frames.*

If the motion model parameter was guessed correctly, the internal motion r_t was calculated correctly for each time frame and the normalized sinogram of the transformed guess will contain the same amount of motion as the measured sinogram. We therefore calculate the absolute difference of these two sinograms for each time frame. These absolute difference sinograms are shown in Figure 4.8. For each time frame, the *total* absolute difference can be determined by summing over all the elements in the absolute difference sinogram. For the three absolute difference sinograms that are shown here, this yields 0.34, 0.35 and 0.35 respectively. These three values are then averaged and multiplied by 0.5. This yields for the model error the value 0.17.

This process can be repeated for different values of the model parameter, which will thus




Figure 4.8: *The absolute difference of the normalized measured sinogram $M_t/\sum_j M_{t,j}$ and the normalized sinogram of the transformed guess $\tilde{G}_t/\sum_j \tilde{G}_{t,j}$, for each time frame.*

imply different values of the estimate of the internal motion r_t for each time frame. The most optimal value of A will have the lowest model error, so by minimization of the model error the most optimal model parameter can be found.

If we choose a different value for A , for instance $A = 2$, the same process can be applied for $A = 1$. For $A = 2$ the model error was found to be 0.21, whereas for $A = 1$ it was 0.17. $A = 1$ is thus a more optimal parameter value than $A = 2$ in this case.

Minimization algorithm

The model error was minimized using Powell's method, which is a minimization algorithm for multi-variable functions [73]. Powell's method performs a sequence of line minimizations that each minimize only one variable, or model parameters in this case. For this, any line minimization method could be used. Powell's method then chooses the successive directions in the parameter space spanned by the model parameters to perform these line minimizations on. Many iterations through all directions are avoided by choosing a few good directions in the parameter space.

Powell's method starts by initializing the set of N directions u_i to the unit vectors,

$$\vec{u}_i = \vec{e}_i \quad i = 1, \dots, N. \quad (4.6)$$

After this, the following steps should be repeated until the function that is to be minimized stops decreasing [73].

1. Save the starting point as \vec{P}_0 .

2. Minimize the function along all directions \vec{u}_i successively.
3. Determine the direction \vec{u}_d that contributed most to the decrease of the function.
4. Discard \vec{u}_d from the set of directions.
5. Add the displacement vector $\vec{P}_N - \vec{P}_0$ to the end of the list of directions.
6. Move \vec{P}_N to the minimum along direction \vec{u}_N and call this point \vec{P}_0 .

Powell's method computes the model error typically a few tens of times for every iteration of the reconstruction, making the minimization step very computationally intensive. Without the optimization step, a reconstruction with 10 iterations takes only a few minutes. With the optimization step included the entire reconstruction takes 2-3 days.

4.3 Simulations

4.3.1 Overview

The goal of this study was to develop a respiratory motion correction method for PET that can deal with breathing variability. Simulations were therefore performed for both variable and non-variable breathing dynamics to test whether the method performed equally well in both cases. Figure 4.9 shows the surrogate signals that were used for stationary and non-stationary breathing dynamics. Simulations were also performed to test the behavior of the model error. In these simulations, the value of one of the model coefficients was changed halfway through the simulation. The internal motion could then no longer be accurately described with the same motion model for both halves of the PET data. As the optimization method assumes that the motion model is constant, motion can then not be perfectly corrected and the hypothesis was that the residual motion could be detected due to an increase in the model error. A simulation of a static phantom was also included, as this can serve as a reference for perfect motion correction. Table 4.1 summarizes these three simulations.

Table 4.1: *The simulations that were performed to test the performance of the model-based motion compensation method.*

Simulation	Motion	Breathing dynamics	Model parameters
1	None	N/A	N/A
2	2D-translational	Stationary	Constant
3	2D-translational	Non-stationary	Constant
4	2D-translational	Stationary	A_y changed at half the scan time




Figure 4.9: *The stationary and non-stationary surrogate signals that were used for the simulations.*

For comparison, all simulations were reconstructed four times, with different motion correction techniques:

1. No motion correction;
2. Gating;
3. Registration of gated images;
4. Model-based correction.

Model-based correction using motion models, as proposed in this study, was compared with the other correction techniques. The uncorrected reconstruction was used to test whether correction improved the image quality at all. Gating and registration of gated images are included for comparison as these are state-of-the-art motion correction techniques. All reconstructions were compared quantitatively, using the analyses described in Section 4.5.

4.3.2 Phantom

Simulations were performed with a 2D liver-like phantom, which is shown in Figure 4.10. The phantom was created such that it resembles the shape of a liver. A smaller and larger sphere were added to represent tumors of sizes 2 and 4 cm. The ratio between the voxel intensity in the tumors and in the background was 8.5:1.



Figure 4.10: *The liver-like phantom that was used for the simulations.*

4.3.3 Motion

Motion dynamics were simulated by generating a series of 18 phantoms that were shifted with respect to each other with a one- or two-dimensional sinusoidal motion. The 18 time frames covered two periods of the sine function. The motion amplitude was 5 voxels, or 1 cm in terms of physical dimensions. As discussed in Section 2.3.7, a motion of 10-25 mm is realistic for liver tumors. Non-stationary breathing patterns were simulated by shifting the second period of the sine function upwards by two amplitudes.

In the simulations, the surrogate signal was not computed by the center-of-mass method as this method cannot be applied with just one axial slice. The surrogate signal was assumed to be known instead.

4.3.4 Projection and backprojection

For projection and backprojection the Radon and inverse Radon transforms were used, respectively. These are the mathematical transforms that describe a projection and backprojection, as described in Sections 2.3.1 and 2.3.3. Since these are mathematical transformations, they do not include a detector response model. Two different approaches that did include modeling the detector were found unsuitable for this study and are described in more detail in Appendix A.

4.3.5 Noise

Poisson noise was added to the simulated measurements. In order to make sure that the amount of noise that was added was realistic, the sinograms were scaled such that they contain a realistic number of counts. This is an important step as the standard deviation of the Poisson distribution is \sqrt{N} with N the average number of counts. The sinograms of the experiments contain a realistic number of counts, including effects such as attenuation and crystal efficiency, and these were thus used to scale the simulated measurements.

4.3.6 Attenuation

Attenuation can easily be modeled by multiplying each element in the measured sinogram with an attenuation factor. As mentioned before in Section 2.3.2, each element in the measured

sinogram corresponds to a certain LOR. Therefore, the attenuation of an element of a sinogram corresponds to the amount of attenuation that was encountered along that LOR. The attenuation factor is thus $e^{\sum_i \mu_i L_i}$, with i the different tissue types that were encountered along the corresponding LOR, μ_i the attenuation coefficient of that tissue and L_i the distance that was traversed in that tissue (see Section 2.2.2).

4.3.7 Iterative reconstruction

The reconstruction consisted of ten iterations in total. The first three iterations were regular MLEM iterations, without motion estimation and correction. This was done such that the first guess that is used for motion model optimization already starts to resemble the (motion-blurred) image object.

4.4 Experiments

4.4.1 Overview

Six different setups were used to test the performance of the proposed model-based motion correction method. For the same reasons as mentioned in Section 4.3.1, the PET data of each of these setups was reconstructed four times, each with one of the following motion correction techniques:

1. No motion correction;
2. Gating;
3. Registration of gated images;
4. Model-based correction.

All reconstructions were compared quantitatively, using the analyses described in Section 4.5.

The first setup was a scan of a static phantom. As the phantom did not move during the scan, this scan could serve as a reference for a reconstruction with perfect motion correction. All motion corrected reconstructions of other experiments could then be compared to this scan to determine the quality of motion correction. Unfortunately, due to practical problems with the PET scanner the data of this scan was lost.

The second setup was used to test the hypothesis that it is feasible to optimize motion model parameters in a PET reconstruction. In this setup, the phantom moved sinusoidally in the y – *direction* (into the scanner bore, see Figure 4.3). The motion was stationary.

The third setup is similar, except that the phantom now moved in both the x - and y -direction. Using this setup it could be tested whether it is feasible to optimize the motion model parameters for motion in two directions. The optimization method used is suitable for

both one- and multi-dimensional problems, so different results from setups 2 and 3 were not expected.

The fourth setup was used to test the hypothesis that model-based motion correction performs better than gating and registration of gated images in the case of variable breathing dynamics. This setup also used two-dimensional motion, but the difference with the third setup is that the motion was now *non*-stationary to simulate variable breathing dynamics. Just as for the simulations, for non-stationary breathing dynamics the amplitude of the sine function was shifted by two amplitudes after the first respiratory cycle.

The fifth setup served to test the behavior of the model error if the assumption that the model parameters are constant throughout an acquisition is no longer valid. This setup is similar to the third setup, except that the values of the motion model parameters were changed after the first respiratory cycle. This was achieved by changing the direction of motion after the first respiratory cycle. As the surrogate signal was computed as the center of mass in the y -direction, the motion in the y -direction still correlated equally with the surrogate signal after the change in direction. The motion model thus retained the same model parameter value in the y -direction after the change in direction. However, the motion in the x -direction correlated with the surrogate signal through a different parameter value. The aforementioned assumption of constant motion model parameters was thus no longer valid. Therefore, assuming a motion model with constant parameters in the motion estimation, motion correction could not be accurate. It was thus expected that there would be residual motion in the reconstruction and the model error would therefore increase.

The sixth setup was used to test how well the current motion model, that only explicitly models translational motion, could correct for rotational motion. This setup used periodic, stationary, rotational motion. Table 4.2 summarizes these six setups.

The PET scan of each setup had a duration of 10 seconds and was performed on a Siemens mCT PET/CT scanner. In order to be able to model the motion dynamics, the raw data from the PET scanner was binned into sinograms with a time width of 0.5 seconds ¹. For each reconstruction, ten iterations were performed. Motion was not estimated nor corrected in the first three iterations, for reasons mentioned in Section 4.3.7.

¹This was achieved using an in-house developed software platform.

Table 4.2: *The setups that were used to test the performance of the proposed model-based motion compensation method.*

Setup	Motion	Breathing dynamics	Model coefficients
1	None	N/A	N/A
2	1D-translational	Stationary	Constant
3	2D-translational	Stationary	Constant
4	2D-translational	Non-stationary	Constant
5	2D-translational	Stationary	Changed at half the scan time
6	1D-rotational	Stationary	Constant

4.4.2 Phantom

A 1.4 L liver-like phantom was used. The compartment of the phantom was filled with water mixed with fluorine-18 FDG, a radiotracer frequently used in diagnostic PET imaging. To represent tumors in the liver, an 8 and a 20 mL sphere were also added to the compartment. The larger of the spheres had a cavity inside, such that it represents a necrotic tumor. In a necrotic tumor, part of the tumor tissue is dead, for instance in the center of the tumor. This can be caused by for instance fast growth of the tumor. The volume of 20 mL of the larger sphere is without the cavity. The spheres were also filled with water mixed with fluorine-18 FDG. The larger sphere contained 8.4 MBq of F18-FDG, the smaller sphere contained 4.7 MBq. The background was filled with 69.8 MBq, although an unknown amount of activity was leaked in the process. The small sphere was filled with a 11.8:1 activity concentration ratio with respect to the background, for the large sphere the ratio was 8.4:1 (not taking into account the leaked activity).

4.4.3 Surrogate signal

The surrogate signal was computed from the data using the center-of-mass method as described in Section 4.2.1.

4.4.4 Projection and backprojection

For the reconstruction of the experiments it is desirable to have a realistic projector and backprojector based on the specific scanner that was used. Therefore, the projector and backprojector from STIR reconstruction software specifically designed² for a Siemens PET scanner were used [74].

²The projector and backprojector were not designed by myself.

4.4.5 Translational motion setup

The setup for translational motion consists of an Arduino attached to a stepper motor, which was programmed to perform a sinusoidal motion with a frequency of 0.2 Hz and an amplitude of 10 mm. These are realistic values for liver tumors in adults [44]. The stepper motor was attached to a movable platform that holds the phantom and can perform a translational motion. Figure 4.11 shows this setup. To simulate non-stationary breathing, the baseline of the sinusoidal motion was shifted by two amplitudes after the first cycle.



Figure 4.11: *The setup for translational motion consisted of a stepper motor that drives the motion of the phantom. The stepper motor was connected to an Arduino and was programmed to perform a sinusoidal motion.*

4.4.6 Rotational motion setup

This setup is similar to the translational motion setup. However, the Arduino was now connected to a rotational motor that was programmed to perform a periodic motion of 0.2 Hz with a constant speed. The rotational motor drives a turntable that rotates with an angular deviation of about 20 degrees. The phantom was placed on the turntable, with the smaller sphere closer to the center of the turntable than the larger sphere. Figure 4.12 shows the rotational setup.




Figure 4.12: *In the rotational motion setup, the Arduino was connected to a rotational motor that drives a turntable that holds the phantom.*

4.5 Quantitative analysis

This section describes the quantitative parameters that were used to assess and compare the image quality of the reconstructions of the simulations and experiments.

4.5.1 Mean Standardized Uptake Value

The Standardized Uptake Value (SUV) of a lesion is often used in PET imaging. It describes the ratio between the radiotracer concentration in the tumor and the average radiotracer concentration in the patient. The SUV at time t is given by

$$SUV(t) = \frac{c(t)}{D/m}, \quad (4.7)$$

with $c(t)$ the radioactivity concentration in the tumor tissue (decay corrected to $t = 0$), D the injected dose at $t = 0$ and m the mass of the patient. As we are only interested in the relative SUV, in order to compare different reconstructions, we were able to put $D = m = \text{constant}$. The SUV of a voxel is then just equal to the radiotracer concentration in that voxel, i.e. the voxel intensity. For each tumor, the mean SUV was computed over the tumor volume.

4.5.2 Tumor volume

To compute the volume of a tumor, first the maximum SUV in that tumor is determined. This maximum SUV can then be used to determine the tumor volume using a technique called *thresholding*. All voxels that are above a certain percentage of the maximum SUV are considered to belong to the tumor. In this study, a threshold of 40% was used [75].

The maximum SUV is rather sensitive to random noise, as it is determined by only one voxel. Therefore, rather than just taking one voxel, an average SUV over a few voxels was computed. These voxels are the voxel containing the maximum SUV and its direct neighbors. The resulting mean SUV, called the *SUV peak* in this thesis, was used for thresholding instead of the maximum SUV. This reduces the influence of random noise in the quantitative comparison of different reconstructions.

4.5.3 Contrast-to-Noise Ratio (CNR)

Two important measures for image quality are contrast and noise. *Contrast* is defined as the difference in intensity between parts of the image. *Noise* describes the random statistical fluctuations in the rate of nuclear decay and causes the images to look mottled. A higher contrast relative to the noise variation makes it easier to distinguish an object from the background.

The Contrast-to-Noise Ratio (CNR) quantifies how well an object can be distinguished from the background and can be defined as

$$\text{CNR} = \frac{I_t - I_b}{\sigma_b}, \quad (4.8)$$

with I_t the average voxel intensity in the tumor, I_b the average voxel intensity in the background and σ_b the standard deviation of the voxel intensities in the background due to noise. The Rose-criterion states that the critical value for distinguishing an object from the background is a CNR of 5 [76].

In order to calculate the CNR, a tumor volume and background volume were defined. The previous section described how the tumor volume was segmented. In choosing the background we were presented with several options. The background was chosen to be a ten-voxel band around the segmented tumor, such that the CNR would be calculated locally. This way any smearing of the tumor that was not segmented into the tumor volume, called *spill-out*, would significantly change the mean and standard deviation of the background. The presence of spill-out would therefore show up in the CNR. Several other options were tried for the background volume as well and are described in Appendix A.

4.5.4 Reference values

For the simulations, these quantitative parameters were compared with their reference values obtained from the simulation of a static phantom. For the experiments, the static scan was unavailable due to technical problems with the PET scanner. Instead, the volumes of the spheres were measured by weighing the spheres both unfilled and filled with water. For the CNR, a reference value is not necessary. The optimal CNR, in the case of zero noise, would be infinite. A higher CNR is thus by definition more optimal. The mean SUV was not calculated for the experiments, as there was no reference value available for comparison.

Chapter 5

Results

The performance of model-based motion correction (MBC) was tested in both simulations and experiments. The results were compared with respiratory gating and registration of gated images. This Section will first discuss the results of the simulations and then the results of the experiments.

5.1 Simulations

The first part of this Section shows the results of a very simple simulation that was performed without Poisson noise. This simulation was used to test the method before the actual simulations were performed. The rest of the Section shows the results of the simulations that were listed in Section 4.3.1. In this part, first the image reconstructions are discussed, then the model parameters optimization and it ends with the results of the quantitative analyses that was performed. As mentioned in Section 4.3.1, simulations were performed using both stationary and non-stationary breathing dynamics.

5.1.1 Noise-free simulations

The noise-free simulations were performed with stationary motion in two directions, constant model parameters and without Poisson noise. Figure 5.1 shows the reconstruction that was corrected using model-based motion correction and a reconstruction of the same phantom that contained no motion at all. These reconstructions contain no visible differences.

Figure 5.2 shows the model errors. For each iteration, the model error was grid sampled as a function of the slope A_y . This grid sampling is indicated by the colored lines. The dashed black line indicates the correct value of the slope A_y . The red dots indicate the slope A_y that was estimated in each iteration.

For each iteration, the sampled model error is smooth and convex and has an optimum near the correct value of the slope A_y . For consecutive iterations, the model error decreases and




Figure 5.1: *The final reconstructions for the noise-free simulation that were a) corrected with model-based correction and b) contained no motion.*

its shape becomes more convex. For all iterations, the estimated parameter value was within 0.34% of the correct parameter value.




Figure 5.2: *The normalized model errors for all iterations as a function of the slope A_y (in pixels). The black dashed line indicates the correct slope. The red dots indicate the parameter value that was estimated in each of the iterations. The colored lines represent the grid sampling of the parameter space that was performed for each iteration.*

Figure 5.3 shows the parameter optimization as a function of the iteration number for the x -axis. The dashed black line indicates the correct parameter value. The red dots indicate the parameter value that was found for each iteration. For all iterations, the correct parameter value was found to within 1.02%.




Figure 5.3: *The recovered motion amplitude as a function of iteration number for the noise-free simulation.*

5.1.2 Image reconstructions

Stationary breathing dynamics

Figure 5.4 shows the image reconstructions that were obtained of the liver-like phantom under stationary dynamics, as well as the static reconstruction that serves as a reference for perfect motion correction (i.e. simulations 1 and 2, see Section 4.3.1). Figure 5.4b shows that for the reconstruction that was not motion corrected, the two spheres are clearly enlarged with respect to Figure 5.4a, indicating the presence of residual motion in the reconstruction. Figures 5.4c-e show the reconstructions that were corrected by using either gating, registration of gated images or model-based motion correction (MBC). These reconstructions do not clearly show visible residual motion when compared to Figure 5.4a.

For the reconstructions in Figure 5.4 the assumption was made that the model parameters are constant throughout the simulation. In simulation 4, this assumption was invalid, in order to test what the influence of an invalid assumption would be on the image quality. Simulation 4 was compared with simulation 2, for which the assumption did remain valid throughout the simulation. Figure 5.5 shows the reconstructions that were obtained for simulations 2 and 4, using MBC. No clearly visible residual motion is present in either of the reconstructions.

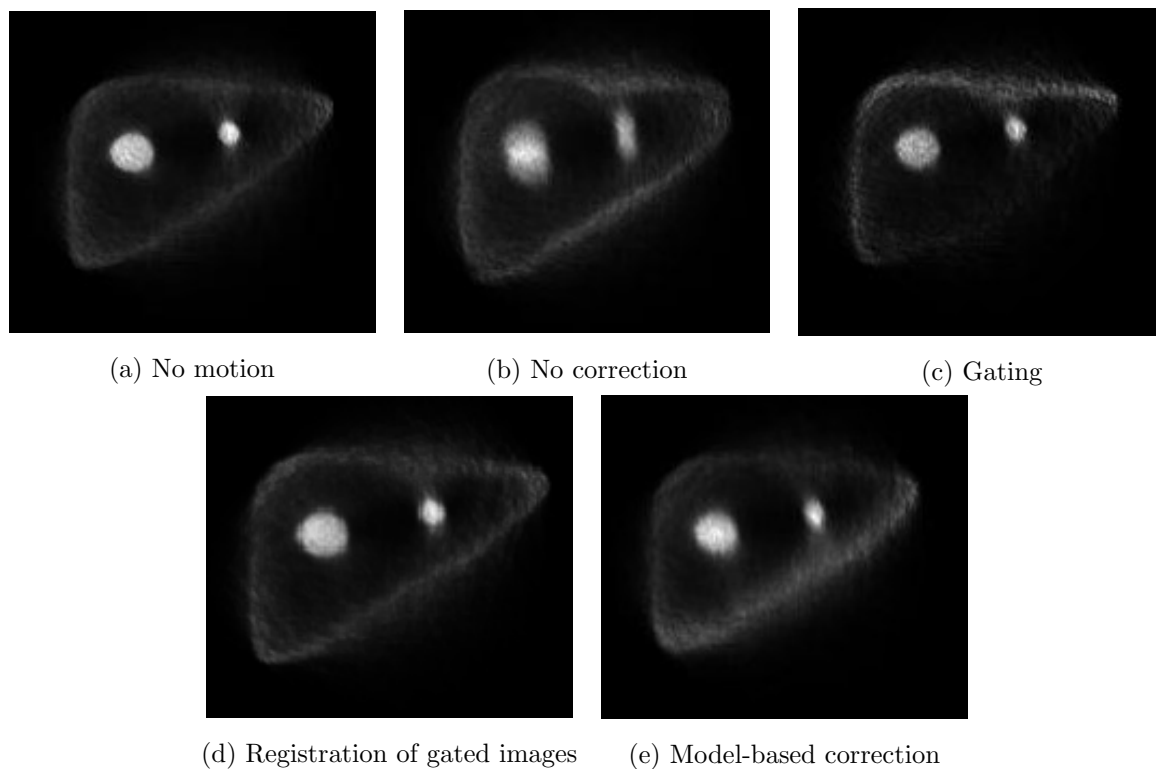


Figure 5.4: *Reconstructions of the liver-like phantom with stationary dynamics that either contained no motion, were not motion corrected or were motion corrected using either gating, registration of gated images or model-based motion correction.*




Figure 5.5: *Reconstructions of the liver-like phantom under stationary dynamics using model-based correction. Left: the model parameters were constant throughout the simulation. Right: the model parameters changed halfway through the simulation.*

Non-stationary breathing dynamics

Figure 5.6 shows the image reconstructions that were obtained of the liver-like phantom under non-stationary dynamics, as well as the reconstruction that contained no motion (i.e. simulations 1 and 3). Figure 5.6b shows the reconstruction that was not motion corrected. The

two spheres are clearly enlarged with respect to Figure 5.6a, indicating the presence of residual motion in the reconstruction. Figures 5.6c-e show that, just as in the case of stationary dynamics, the three reconstructions that used either of the three motion correction techniques do not clearly show visible residual motion.

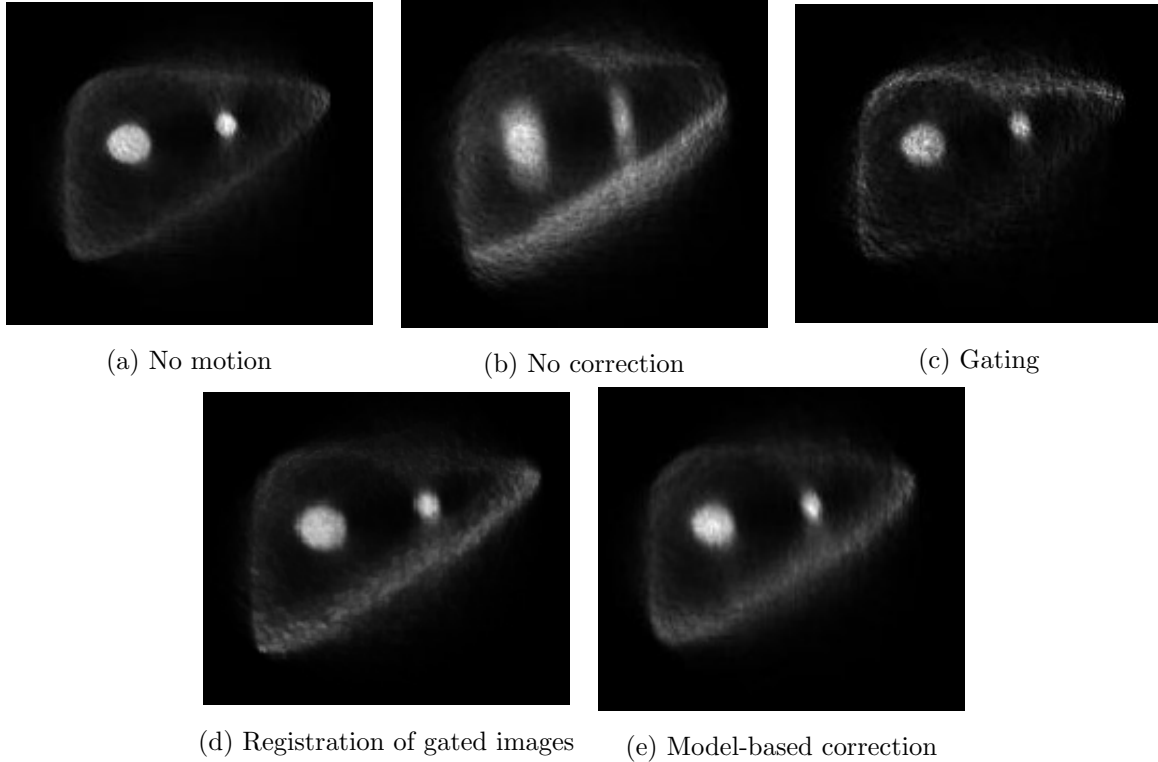


Figure 5.6: *Reconstructions of the liver-like phantom with non-stationary dynamics that either contained no motion, were not motion corrected or were motion corrected using either gating, registration of gated images or model-based motion correction.*

5.1.3 Model parameters optimization

The model parameters optimization of the model-based motion corrected image reconstructions will be discussed in this Section. The parameter space was characterized using grid sampling in order to determine if the model error is suitable for optimization. Figure 5.7 shows an example of a sampling of the parameter space at the start of the 10'th iteration of a reconstruction containing model-based motion correction. This shows that the model error is smooth and convex. The model error also has a clear optimum near the correct value of the model parameter ($A_y = 1.4$).

As mentioned before, simulations 2 and 4 used constant and changing model parameters respectively. Figure 5.8 discusses the obtained optimal model parameters for each iteration for simulations 2 and 4. In this Figure not the model parameters themselves are shown on the




Figure 5.7: *Grid sampling of the normalized model error at the start of the 10'th iteration of a reconstruction containing model-based motion correction. The shape of the model error is smooth and convex and has a clear optimum near the correct parameter value ($A_y = 1.4$).*

vertical axis, but the total motion amplitude $r_t = \sqrt{x_t^2 + y_t^2}$ of the first respiratory cycle that is implied by these parameter values through Equation 3.1. The amplitude of the first cycle is used here, because this amplitude is equal for both stationary and non-stationary surrogate signals. The dashed black line indicates the correct total motion amplitude.

Figure 5.8a shows that in the case of constant parameters, the total motion amplitude is found to within 3 mm at the first iteration. For further iterations, the total estimated motion amplitude converges to within one pixel size (2mm) of the correct value. Figure 5.8b shows that in the case of changing model parameters, the total motion amplitude is estimated to within about one pixel size at the first model optimization and remains approximately constant for future iterations. This was an unexpected result.

Figure 5.9 shows the obtained model parameters for each iteration of a reconstruction under non-stationary dynamics that was corrected using MBC. The results shows that at the first model optimization only half of the total motion amplitude was recovered. For further iterations the total motion amplitude converges to within a pixel size of the correct value.

Figure 5.10 shows the time-resolved model error for simulations 2 and 4. In simulation 2, the model parameters were constant, whereas in simulation 4 the model parameters A_y changed halfway through the simulation. The model errors of simulation 2 and 4 are of comparable size and are both relatively constant as a function of time.




Figure 5.8: *The recovered motion amplitude as a function of iteration for the simulations with a) constant model parameters (simulation 2) and b) one of the parameters changed halfway through the simulation (simulation 4).*




Figure 5.9: *The recovered motion amplitude as a function of iteration for the simulations with non-stationary breathing dynamics.*




Figure 5.10: *The model error for the reconstructions of simulations 2 and 4 using model-based motion correction. Simulations 2 and 4 used constant and changing model parameters respectively.*

5.1.4 Quantitative analyses

Figure 5.11 shows the mean SUV, CNR and segmented volume for both lesions under both stationary and non-stationary breathing dynamics. The black line indicates the reference values as obtained from the reconstruction without motion (see Section 4.3.1). The blue and red bars represent the reconstructions that used stationary and non-stationary motion respectively (i.e. simulations 2 and 3 respectively).

Figures 5.11a and 5.11b show the mean SUV of the large and small lesion respectively for all reconstructions. These results show that the reconstruction without motion correction strongly underestimates the mean SUV, both for the large and the small lesion. The mean SUV values of the reconstructions that were corrected using gating, registration of gated images and MBC are all much closer to the reference value. For non-stationary breathing the mean SUV is lower than for stationary breathing.

Figures 5.11c and 5.11d show the volumes of the large and small lesion respectively for all reconstructions. As expected, the reconstruction without motion correction overestimates the lesion volume both for the small and the large lesion. The lesion volume is overestimated more for non-stationary breathing dynamics than for stationary breathing dynamics. All three motion correction methods reduce the tumor volume to within ten pixels of the reference value. For these three motion correction techniques, the lesion volume does not significantly differ for stationary and non-stationary breathing dynamics.

Figures 5.11e and 5.11f show the contrast-to-noise ratios (CNR) of the large and small lesion respectively for all reconstructions. For the large lesion the CNRs of the reconstructions that were corrected with registration of gated images or MBC were higher than those of the reconstructions that were corrected with gating or not motion corrected. For the small lesion, all CNRs were similar, except that the CNR of the reconstruction without motion correction under non-stationary breathing had a significantly lower value. Both for the large and small lesion the CNR of the reconstructions with non-stationary dynamics is lower than the CNR of the reconstructions with stationary dynamics.

Figure 5.12 shows the mean SUV, segmented volume and CNR for the reconstructions of simulations 2 and 4 corrected with MBC. Simulation 2 had constant motion model parameters, whereas in simulation 4 the model parameter A_y changed halfway through the simulation. No significant differences were observed for these parameters between simulations 2 and 4.

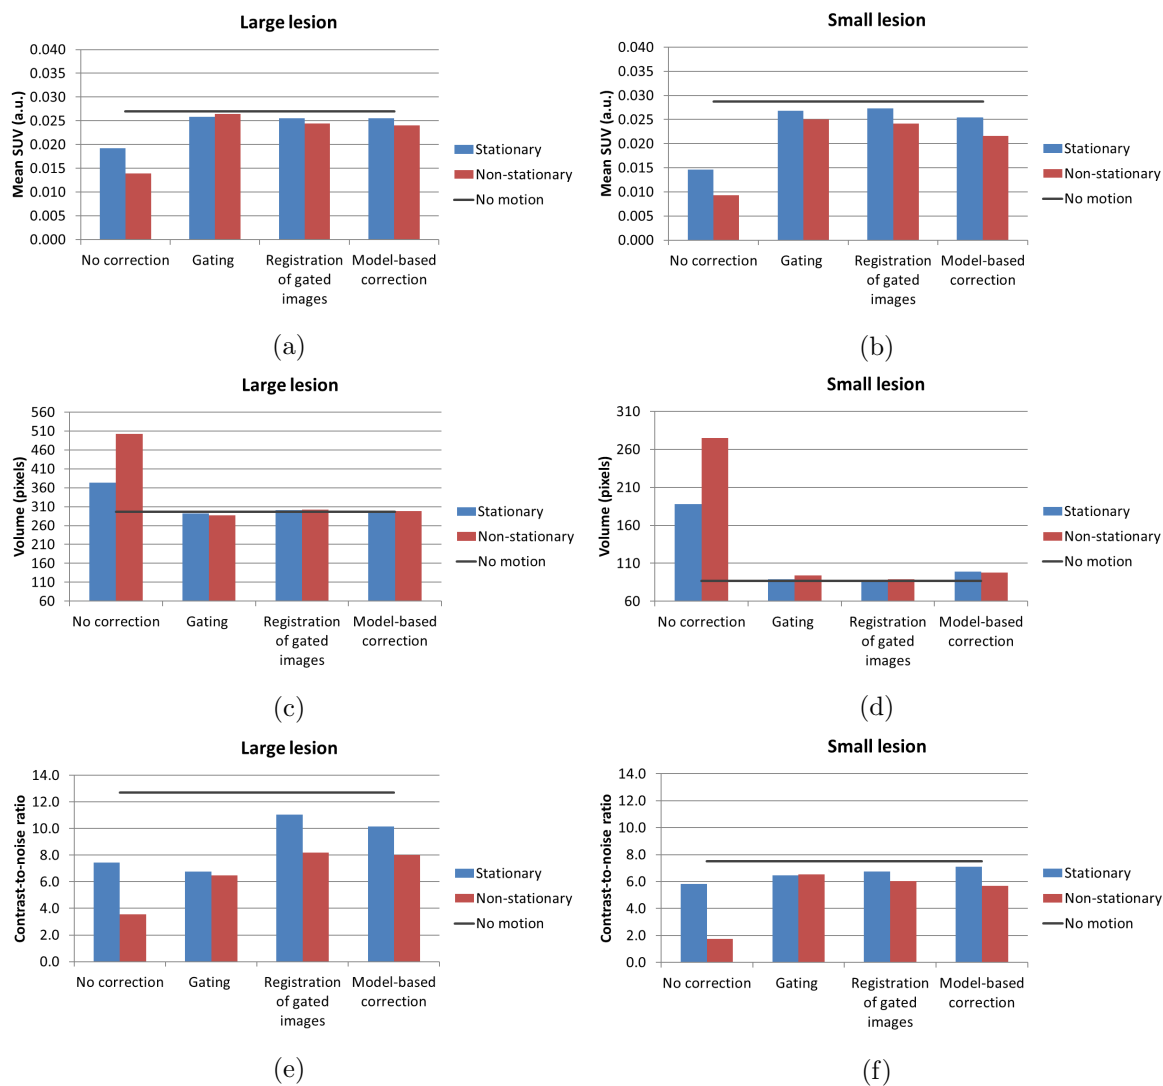


Figure 5.11: The mean SUV, lesion volume and CNR were determined for all of the reconstructions under both stationary (blue) and non-stationary (red) breathing dynamics.

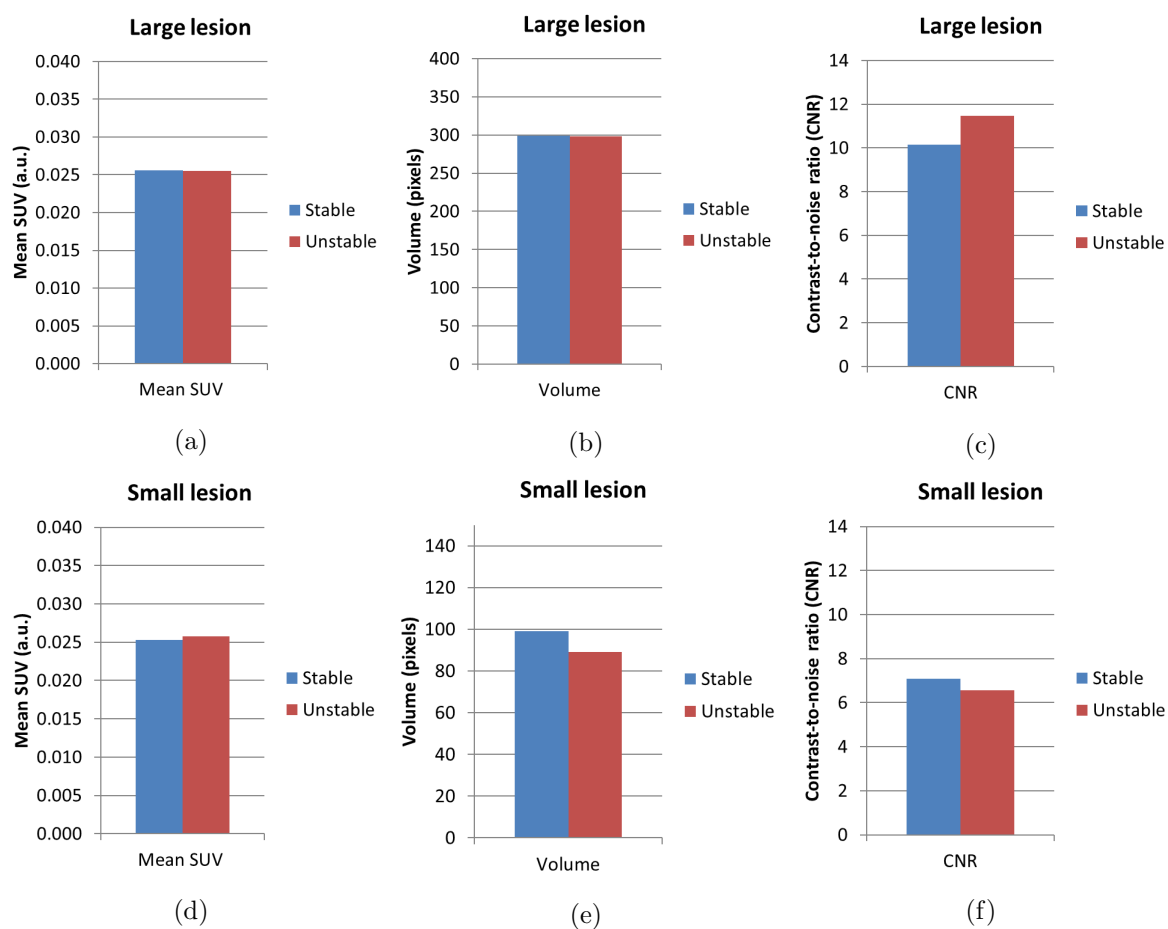




Figure 5.12: Comparison of the quantitative analyses of simulations 2 and 4, which had constant (blue, “stable”) and changing (red, “unstable”) motion model parameters, respectively.

5.2 Experiments

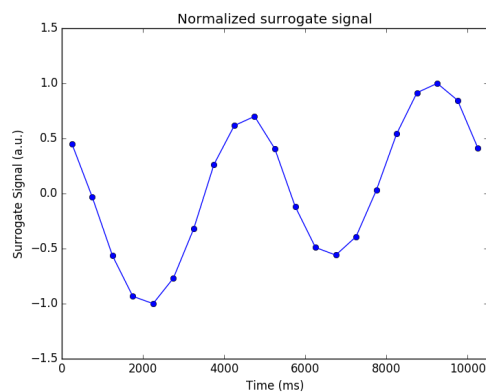

The results of the simulations are divided in four parts. First the measured surrogate signals are shown, then the image reconstructions are discussed, followed by the model optimization and the section ends with discussing the quantitative analyses that was performed.

5.2.1 Surrogate signals


Figure 5.13 shows the surrogate signals that were measured using the center-of-mass method for each of the experimental setups (see Section 4.4.1 for an overview of the setups). Figures 5.13a and 5.13b show the surrogate signals for experimental setup 2 and 3, that used 1D and 2D stationary, translational motion respectively. As expected, both surrogate signals have a sinusoidal shape. Figure 5.13c shows the surrogate signal for setup 4, that used 2D non-stationary, translational motion. This surrogate signal shows that after the first respiratory cycle, the sinusoidal function was shifted down by two amplitudes. Figure 5.13d shows the surrogate signal for setup 5, that used 2D translational motion with changing motion model parameters. Due to the changed motion model parameters for the second respiratory cycle, the surrogate signal was shifted up and the amplitude slightly decreased. This was expected, due to the change in direction of the setup. Figure 5.13e shows the surrogate signal for setup 6, that used rotational motion. As expected, this surrogate signal is periodic, but not sinusoidal.



(a) Setup 2 (1D stationary, translational motion) (b) Setup 3 (2D stationary, translational motion)



(c) Setup 4 (2D non-stationary, translational motion) (d) Setup 5 (2D translational motion with changing coefficients)



(e) Setup 6 (rotational motion)

Figure 5.13: *The surrogate signals that were measured using the center-of-mass method for each of the five experimental setups.*

5.2.2 Image reconstructions

For all reconstructions only a slice parallel to the (x, y) -plane (see Figure 4.3) is shown, as the motion in both the x - and y -direction is visible in these planes. Figures 5.14-5.16, 5.18, 5.20 show the four image reconstructions for experimental setups 2-5. First, the similarities of the reconstructions of different experiments will be discussed. Then each experiment and its unique results will be discussed separately.

Figures 5.14a-5.16a, 5.18a, 5.20a show the reconstructions that were not motion corrected for all experiments. In all reconstructions, enlarged lesion volumes are clearly visible, as expected. Figures 5.14b-5.16b, 5.18b, 5.20b show the reconstructions that were motion corrected using gating. As expected, all of these gated reconstructions are much noisier than the other reconstructions. The reconstructions also clearly contain less motion than the reconstructions that were not motion corrected. Figures 5.14c-5.16c, 5.18c, 5.20c show the reconstructions that were corrected using registration of gated images. These reconstructions also clearly contain less motion than the reconstruction that was not motion corrected. Moreover, the reconstructions contain much less noise than the gated reconstructions, but still slightly noisier than the reconstructions that were not corrected or corrected using MBC. Figures 5.14d-5.16d, 5.18d, 5.20d show the reconstructions that were corrected using MBC. Together with the reconstructions that were not motion corrected, the MBC corrected reconstructions contain the least noise. The MBC corrected reconstructions also clearly contain less motion than the reconstructions that were not motion corrected.

Figure 5.14 shows all reconstructions for experimental setup 2 (see Table 4.2). This experiment was performed using 1D stationary, translational motion. Figure 5.14a shows the reconstruction without motion correction. It clearly shows residual motion in only one direction, as expected.

Figure 5.15 shows all reconstructions for experimental setup 3, that was performed using 2D stationary, translational motion. Figure 5.15a shows the reconstruction without motion correction, which in this case clearly shows residual motion in two directions.

Figure 5.16 shows all reconstructions for experimental setup 4. This experiment was performed using 2D *non-stationary* translational motion. The MBC reconstructions in Figure 5.16d stand out as they seem to contain more residual motion than the MBC reconstructions of the similar experiment using stationary motion in Figure 5.15d. For better comparison, Figure 5.17 shows the MBC reconstructions for stationary and non-stationary motion side-by-side. This reveals that indeed the reconstruction of the experiment with non-stationary motion contains slightly more residual motion than the reconstruction of the experiment with stationary motion. This was an unexpected result.

Figure 5.18 shows the reconstructions for experimental setup 5. This experiment used 2D




Figure 5.14: *Experimental setup 2 (1D stationary, translational motion) was reconstructed without motion correction and using three different motion correction techniques for comparison. For all reconstructions only a slice parallel to the (x, y) -plane is shown. The first row shows the large lesion, the second row shows the small lesion.*




Figure 5.15: *Experimental setup 3 (2D stationary translational motion) was reconstructed without motion correction and using three different motion correction techniques for comparison.*




Figure 5.16: *Experimental setup 4 (2D non-stationary translational motion) was reconstructed without motion correction and using three different motion correction techniques for comparison.*




Figure 5.17: *Reconstructions using model-based motion correction. Left: the reconstructions of experimental setup 3, that used 2D stationary, translational motion. Right: the reconstructions of experimental setup 4, that used 2D non-stationary, translational motion.*

stationary, translational motion, but halfway through the acquisition the value of the model parameter A_y was changed. The MBC corrected reconstruction of the small lesion in Figure

5.18d clearly contains residual motion. The MBC corrected reconstruction of the large lesion also seems to contain residual motion. For a good comparison, Figure 5.19 shows both experimental setups 3 and 5 side-by-side. These experiments only differ in the fact that they used constant and changing model parameters respectively. In this Figure it is clear that not only the small lesion, but also the large lesion contains more residual motion in the case of changing model parameters.

Figure 5.20 shows the reconstructions for experimental setup 6. This experiment used stationary, rotational motion rather than translational. Figure 5.20a shows that the motion amplitude of the small lesion was much larger than the motion amplitude of the larger lesion. The MBC corrected reconstruction of the small lesion in Figure 5.20d clearly contains a lot of visible residual motion. The reconstructions that were corrected using gating and registration of gated images, in Figures 5.20b and 5.20c respectively, contain far less residual motion than the MBC corrected reconstruction.




Figure 5.18: *Experimental setup 5 (2D stationary translational motion with changing coefficients) was reconstructed without motion correction and using three different motion correction techniques for comparison.*




Figure 5.19: *Reconstructions using model-based motion correction. Left: the reconstructions of experimental setup 3, that used 2D stationary, translational motion with constant motion model parameters. Right: the reconstructions of experimental setup 5, that also used 2D stationary, translational motion, but in this setup the value of the model parameter A_y changed half-way through the acquisition.*




Figure 5.20: *Experimental setup 6 (rotational motion) was reconstructed without motion correction and using three different motion correction techniques for comparison.*


5.2.3 Model parameters optimization

Figure 5.21 shows the optimization of the model parameters as a function of iteration for all experiments. Just as for the simulations, not the model parameters themselves are shown on the vertical axis, but the total motion amplitude $r_t = \sqrt{x_t^2 + y_t^2}$ of the first respiratory cycle that is implied by these parameter values through Equation 3.1. The dashed black lines indicate the correct motion amplitude of the phantom. For experimental setups 2-5 that used translational motion (Figures 5.21a-5.21d) the correct motion amplitude was 10 mm. For experimental setup 6 that used rotational motion (Figure 5.21e), the correct motion amplitude was not indicated as the translational amplitude was different for both of the spheres. This is due to the fact that both spheres were located at a different distance from the center of the turntable.


Figures 5.21a-5.21c show the results of experimental setups 2-4. In these experiments we expected that the total motion amplitude would be recovered. The results show that in the first model parameters optimization (during iteration 4) the total motion amplitude is already recovered to within the size of a voxel ($2 \times 2 \times 2 \text{ mm}^3$). For consecutive iterations the estimate of the model parameters converges towards the correct motion amplitude. This is not in agreement with the MBC reconstructions, as these showed more residual motion in the MBC reconstructions of experimental setup 4 (Figure 5.16d) than in the MBC reconstructions of experimental setups 2 and 3 (Figures 5.14d and 5.15d respectively). Figures 5.21a and 5.21c show the presence of discontinuous jumps in the optimization of the model parameters at iterations 5 and 9 respectively. These jumps are significantly larger than the size of a voxel.

Figures 5.21d and 5.21e show the results of experimental setups 5 and 6. In these experiments we did not expect that the total motion amplitude would be recovered. In experimental setup 5 the model coefficients were not constant throughout the acquisition. As the optimization of the motion model parameters assumes that the parameters are constant, it would thus not be possible to find the perfect model parameters and the corresponding correct total amplitude. Figure 5.21d shows that the optimization of the model parameters for this experiment is almost constant for consecutive iterations, just as for experimental setups 2-4. The values of the total motion amplitude are slightly smaller than for experimental setups 2 and 3, which is in agreement with the MBC reconstructions that showed more residual motion for experimental setup 5 (Figure 5.18d) than for experimental setups 2 and 3 (Figures 5.14d and 5.15d respectively). Moreover, in the first model parameters optimization (during iteration 4) the total motion amplitude is already recovered to within one voxel size.

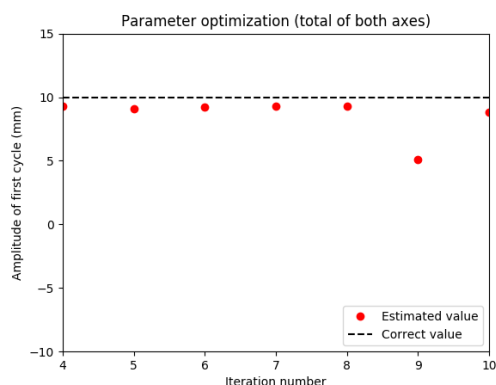
Figure 5.21e shows the recovered total motion amplitude for experimental setup 6 that used rotational motion. The results show the recovered total motion amplitude was approximately constant for consecutive iteration. The results also show that only a very small amount of motion was recovered. This is in agreement with the MBC reconstruction of this experiment that showed a lot of residual motion.




(a) Experimental setup 2 (1D stationary, translational motion)




(b) Experimental setup 3 (2D stationary, translational motion)



(c) Experimental setup 4 (2D non-stationary, translational motion)



(d) Experimental setup 5 (2D stationary, translational motion with changing coefficients)




(e) Experimental setup 6 (rotational motion)


Figure 5.21: *The recovered motion amplitude as a function of iteration of the MBC reconstructions of all experiments.*

Figure 5.22 shows the time-resolved model error for all five experimental setups. The model errors were normalized between 0 and 1. All model errors lie between 0.9 and 1. The model


errors of experimental setups 2,5 and 6, indicated in Figures 5.22a, 5.22d and 5.22e respectively, are constant over time and are all of comparable size. The model errors of experimental setups 3 and 4, indicated in Figures 5.22b and 5.22c respectively, are not constant over time. For experimental setup 4 the model error seems to reach values above 1, which was not supposed to be possible.




(a) Experimental setup 2 (1D stationary, translational motion)




(b) Experimental setup 3 (2D stationary, translational motion)



(c) Experimental setup 4 (2D non-stationary, translational motion)



(d) Experimental setup 5 (2D translational motion with changing coefficients)



(e) Experimental setup 6 (rotational motion)

Figure 5.22: *The time-resolved errors of the model-based corrected reconstructions for all experimental setups.*

5.2.4 Quantitative analyses

Figure 5.23 shows the volumes of the large and small sphere for all experimental setups and all reconstructions. The black lines indicate the reference value for the volume that was obtained by measuring the sphere both without water and filled with water. The results show that in the reconstructions without motion correction, the lesion volume is overestimated for all experimental setups, except for the volume of the small lesion in experimental setup 3 (Figure 5.23b) and the volume of the large lesion in experimental setup 4 (Figure 5.23c). Experimental setup 3 used 2D stationary, translational motion, and experimental setup 4 used 2D *non*-stationary, translational motion. Overestimation of the lesion volume is in agreement with the residual motion that was seen in the reconstructions in Figures 5.14-5.20.

The results for all experimental setups show an overestimation of the volume of the large sphere for all reconstructions, except for the reconstruction without motion correction of experimental setup 4 (Figure 5.23c).

For most setups the volumes of the reconstructions that were corrected with gating, registration of gated images and MBC are of similar size. This is not the case for the volume of the small lesion in the reconstructions of experimental setup 5 (Figure 5.23d), where the volume of the reconstruction that was corrected using registration of gated images is much smaller than the other volumes, but closer to the reference value. Experimental setup 5 used 2D stationary, translational motion with changing model parameters. And it is also not the case for the volumes of both the small and large lesions in the reconstructions of experimental setup 6 (Figure 5.23e), where the volume of the reconstruction that used registrations of gated images is much larger than the reference value for the large lesion and slightly smaller than the reference value for the small lesion. Experimental setup 6 used rotational motion.

Figure 5.24 shows the volumes of all reconstructions of experimental setups 3 and 4 in one Figure, side-by-side, for better comparison. These experiments used both used 2D translational motion, but for experimental setup 3 the motion was stationary whereas for setup 4 the motion was *non*-stationary. The results show that for non-stationary motion, the volumes of the MBC reconstruction and the reconstruction that was corrected using registration of gated images are always higher than for stationary motion. The results for MBC are in agreement with the residual motion in the reconstructions in Figure 5.17. For the gated reconstruction and the reconstruction without motion correction, the volume of the reconstruction of the non-stationary experiment is not always higher than the volume of the stationary experiment.

The quantitative analyses of the MBC reconstruction of experimental setup 5 (Figure 5.23d) shows a large overestimation of the volumes of the small and large lesions. This is in agreement with the residual motion that was visible in the MBC reconstruction in Figure 5.18d.

The quantitative analyses of the MBC reconstruction of experimental setup 6 (Figure 5.23e) shows that both the volumes of the small and large lesion were similar to those of the recon-

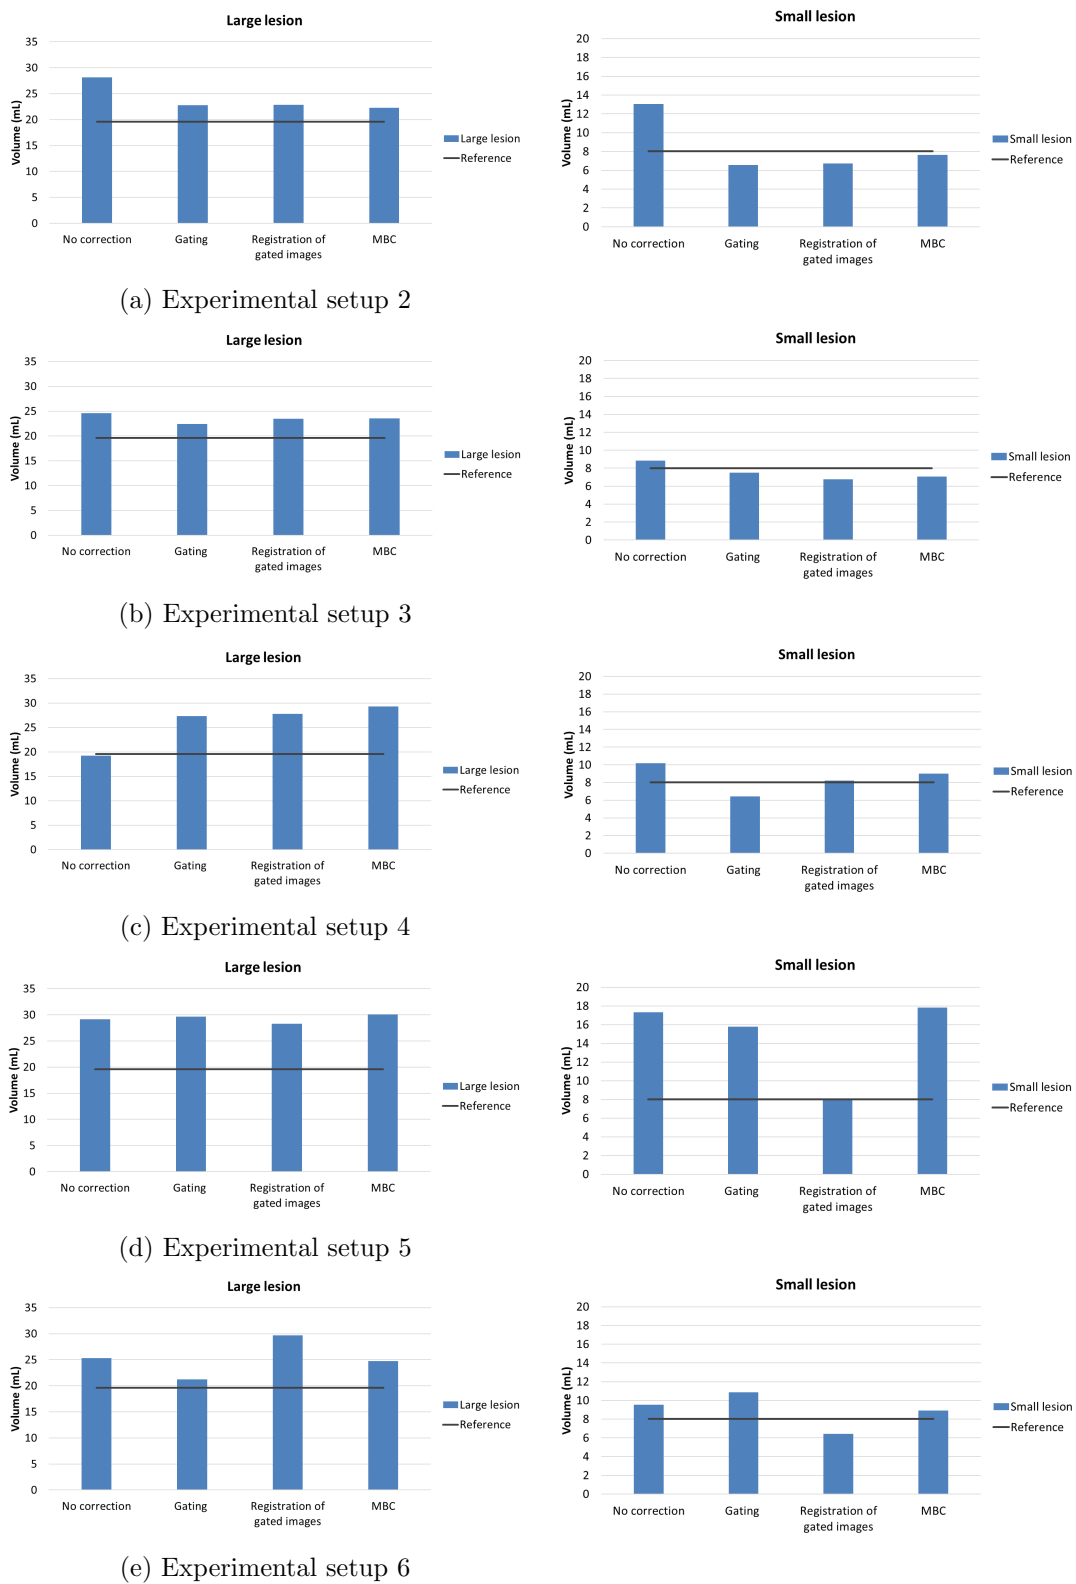


Figure 5.23: The segmented volumes of the large and small lesion for all experimental setups.




Figure 5.24: *The volumes of the large and small lesion for experimental setups 3 (red) and 4 (blue)*

struction that was not motion corrected. This is in agreement with the large amount of residual motion visible in the reconstructions in Figure 5.20d.

Figure 5.25 shows the contrast-to-noise ratios of the large and small lesion for all experimental setups and all reconstructions. Figures 5.25a and 5.25b show the results for experimental setup 2, that used 1D stationary, translational motion, and experimental setup 3, that used 2D stationary, translational motion, respectively. For the large lesion in both experimental setups, the MBC reconstruction had the highest CNR, followed by the reconstruction that was corrected using registration of gated images, the gated reconstruction and the reconstruction that was not motion corrected. For the small lesion, the same pattern applies, except that the gated reconstruction now has a slightly lower CNR than the reconstruction that was not motion corrected.

Figure 5.25c shows the results for experimental setup 4, that used 2D *non*-stationary, translational motion. The same pattern applies for the volumes of the large and small lesions as for experimental setups 2 and 3, except that the MBC reconstruction now has a slightly smaller CNR than the reconstruction that was corrected using registration of gated images. For a better comparison of experiments 3 and 4, Figure 5.26 shows the CNRs of these experiments side-by-side, with experiment 3 in red and experiment 4 in blue. The results show that for the gated reconstruction, the reconstruction that was corrected using registration of gated images and the MBC reconstruction the CNR is always smaller for non-stationary motion than for stationary motion. This is in agreement with the residual motion in the MBC reconstruction for non-stationary motion that was seen in Figure 5.17.

Figure 5.25d shows the CNRs for experimental setup 5, that used 2D stationary, translational motion with changing model parameters. Again, for the large lesion the MBC reconstruction has the highest CNR, followed by the reconstruction that was corrected using registration of gated images, the gated reconstruction and the uncorrected reconstruction. For the small lesion, the reconstruction that was reconstructed using registration of gated images has the highest CNR, followed by the MBC reconstruction, the uncorrected reconstruction and the

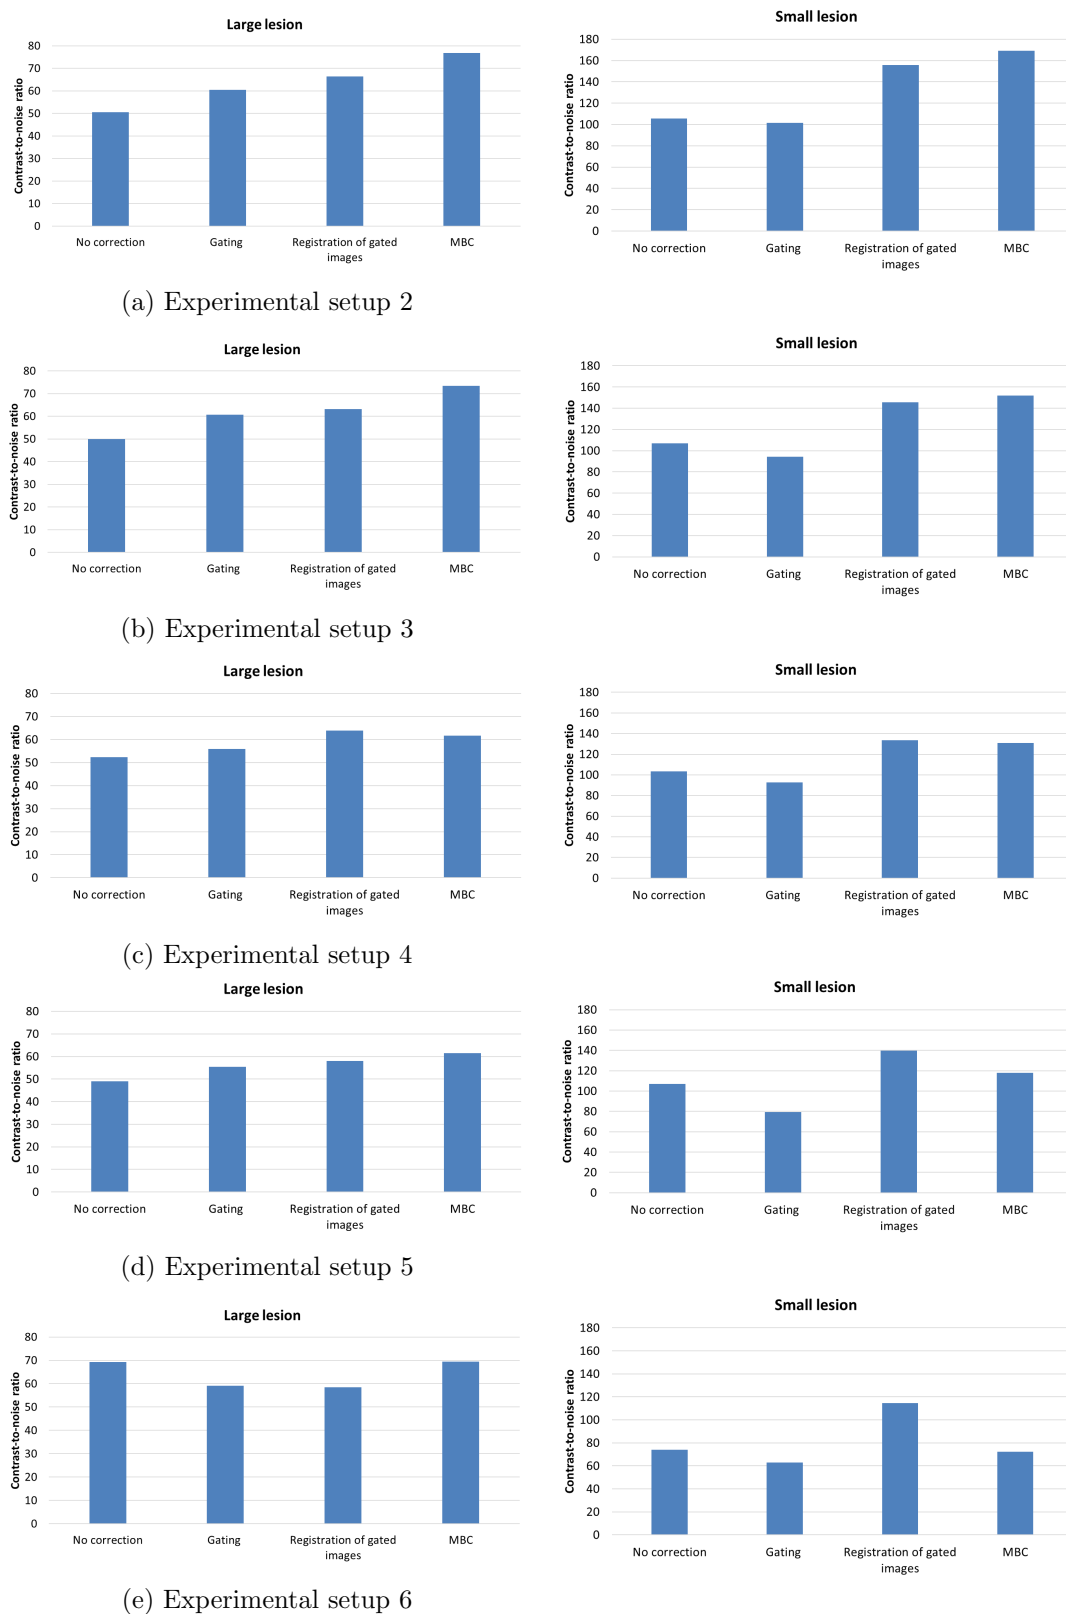


Figure 5.25: The CNRs of the large and small lesion for all experimental setups.




Figure 5.26: *The CNRs of the large and small lesion for experimental setups 3 (red) and 4 (blue), that used stationary and non-stationary motion respectively.*

gated reconstruction. Figure 5.27 compares these results of experimental setup 5 (in blue) with the results of experimental setup 3 (in red). These two setups only differed in the fact that setup 3 used constant motion model parameters, whereas setup 5 used changing motion model parameters. The results show that the CNR is always smaller for the experiment with changing parameters, for all reconstructions. This is an indication of more residual motion in the reconstructions in the case of non-stationary experiments. This is in agreement with the residual motion that is seen in the reconstructions in Figure 5.19.




Figure 5.27: *The CNRs of the large and small lesion for experimental setups 3 (red) and 5 (blue), that used constant and changing motion model parameters respectively.*

Figure 5.25e shows the CNRs for experimental setup 6, that used rotational motion. The CNRs for the large lesion are of similar size to those of the other experimental setups, indicating that there was not much more residual motion for the large lesion in this experiment than in the other experiments. For the small lesion however, the CNRs of all reconstructions are clearly lower than those of the other experiments. This indicates the presence of more residual motion than in the other experiments for the small lesion. Especially the MBC reconstruction has a low CNR. These results are in agreement with the small and large amounts of residual motion that were seen for the large and small lesion, respectively, in Figure 5.20.

Chapter 6

Discussion

The aim of this study was to develop a novel motion correction method for PET acquisitions of the liver that can deal with breathing variability and can be assessed through a quality measure. The performance of this novel method was tested and compared with respiratory gating and registration of gated images and the behavior of the quality measure was investigated.

The results of the simple, noise-free simulation showed that there was no visible difference between the MBC reconstruction and the reconstruction without motion in Figure 5.1. This means that a lot of the motion was corrected for during the reconstruction. The grid sampling of the model error as a function of the slope A_y for each iteration in Figure 5.2 shows that the shape of the model error is smooth and convex and has an optimum at the right parameter value. This model error is therefore very suitable for optimization of the model parameters. It also shows that the model error decreases for successive iterations. This is caused by the fact that each iteration the image guess gets closer to the optimal (motion-free) image guess, therefore decreasing the model error. Figures 5.2 and 5.3 show that for both A_x and A_y the correct model parameters are found with great accuracy in each model iteration. This simple, noise-free simulation thus showed that using the method we propose it is possible to optimize a motion model in these simulations.

The reconstructions without motion correction in Figures 5.4, 5.6 and 5.14-5.16, 5.18 and 5.20 clearly show the presence of residual motion. As expected, this residual motion causes overestimation of the tumor volume (Figure 5.11). Because the activity of the lesions was smeared out over a larger volume, the intensity of the lesions in the reconstructions was reduced. Therefore the mean SUV and CNRs of the lesions were also reduced (Figure 5.11). Both the reconstructions and the quantitative analyses thus show that there is a significant amount of residual motion present in the uncorrected reconstruction, as expected. This is harmful for applications such as radiotherapy, where accurate delineation of the tumor volume is needed to avoid irradiation of healthy tissue and improve tumor coverage.

The results show that in the case of stationary translational breathing (experimental setups 2

and 3 and simulation 2), model-based motion correction (MBC) performs similarly to respiratory gating and registration of gated images in terms of lesion volume and mean SUV, as can be seen in Figures 5.11 and 5.23. Figures 5.11 and 5.25 show that the CNR of MBC is always higher than or similar to those of respiratory gating and registration, indicating less smearing of the lesion in the case of MBC. This means that for stationary motion, model-based motion correction removes more motion from the reconstruction than gating and registration of gated images.

This can be explained by the higher time resolution of MBC with respect to gating and registration of gated images. In MBC, the data was binned into 0.5 seconds time frames. These time frames were used individually, so none of them were averaged. In gating however, the amplitude bins are quite large, causing much more than 0.5 seconds of sequential PET data to end up in a gate. So due to the large width of the bins in gating, there will always be residual motion in a gate, causing the CNR to be smaller than for MBC. So although the motivation in this study to investigate the possibility to apply motion models to PET reconstructions originated from the fact that gating is not suitable for patients with non-stationary breathing dynamics, motion models also proved to be beneficial in the case of stationary dynamics.

The original motivation of this study, to investigate the benefit of model-based motion correction for non-stationary breathing dynamics, was investigated in experimental setups 3 and 4, as well as in simulations 2 and 3. These experiments and simulations differ only in the fact that they used stationary and non-stationary dynamics. Figures 5.11, 5.24 and 5.26 show that in the non-stationary case gating, registration of gated images and model-based correction tend to underestimate the CNR and mean SUV and overestimate tumor volume. These all indicate the presence of more residual motion in the case of non-stationary dynamics. This is in agreement with the results in Figure 5.17, which clearly shows that the residual motion in the MBC reconstruction of the experiment that used non-stationary dynamics (setup 4) is increased with respect to the MBC reconstruction of the experiment that used stationary dynamics (setup 3).

For gating and registration of gated images this was an expected result. The non-stationary dynamics had a total amplitude range of 20 mm, whereas the stationary dynamics had a total amplitude range of 10 mm. In both cases, the surrogate signal was binned into 5 gates based on amplitude. In the case of non-stationary dynamics, these bins were thus twice as large as for the stationary dynamics, causing more residual motion to be present in a single gate. Registration only spatially aligns the individual gated reconstructions, but cannot remove the residual motion that is still in a gated reconstruction.

For MBC however, it was hypothesized that it would perform equally well under stationary and non-stationary breathing dynamics. It was thus not expected that there would be more residual motion in the MBC reconstruction under non-stationary dynamics. Figure 5.21c shows

the total recovered motion amplitude for the MBC reconstruction of the experiment using non-stationary dynamics. The first eight iterations the recovered amplitude is within one voxel size ($2 \times 2 \times 2 \text{ mm}^3$) of the correct value, and comparable to the recovered amplitude in the stationary experiment (Figure 5.21b). At the ninth iteration however, for the non-stationary experiment the recovered motion amplitude suddenly jumps to about half of its value. This causes the motion correction at this iteration to not correct for all of the respiratory motion, and part of the motion is brought back into the reconstruction. This happens one iteration before the reconstruction is finished, and it is thus very likely that this extra motion will not have been removed from the reconstruction in the final iteration.

The residual motion in the experiment using non-stationary breathing is thus probably caused by this discontinuity in the optimization, rather than by the nature of the breathing dynamics. The fact that up until the eighth iteration the recovered motion amplitude of the non-stationary experiment resembles that of the stationary experiment (Figures 5.21c and 5.21b respectively) supports this hypothesis. A more stable optimization method is thus needed to draw a definite conclusion on the performance of model-based motion correction in the case of non-stationary breathing dynamics.

For the motion model it was assumed in the optimization that the model parameters were constant throughout the acquisition. In experimental setup 5 however, this assumption was not valid, and its effect on the image quality of the resulting reconstruction was investigated. Figure 5.23 shows that when comparing the volumes of the model-based reconstructions of experimental setups 3 and 5, that used constant and changing model parameters respectively, the volumes of the lesions were increased in experimental setup 5. This indicates the presence of residual motion. This was expected in the case of non-constant parameters. Because in the optimization method the model parameters were assumed to be constant, the correct parameter values could not be found. Figures 5.21b and 5.21d also show that the total motion amplitude that was recovered was less for setup 5 than for setup 3. This result emphasizes the necessity of a plausible model for the relation between the surrogate signal and incorrect motion. An incorrect model will result in residual motion in the reconstruction, and reduce image quality. The simulations in Figure 5.12 do not show this effect.

In the case of rotational motion (experimental setup 6), residual motion is clearly visible in the reconstruction in Figure 5.20, especially for the small lesion. Less residual motion is visible in the gated and registered reconstructions, indicating that gating and registration of gated images clearly outperform MBC in the case of rotational motion. This was expected, as MBC describes the motion of both lesions with the same linear motion model. Since both spheres were at a different distance from the center of the turntable, they had the same angular motion amplitude but not the same translational motion amplitude. Therefore it is not possible to accurately describe the motion of both spheres by the same motion model parameters for the

motion model that was proposed in this study. Model-based motion correction was therefore very inefficient in the case of rotational motion. Gating and registration of gated images on the other hand are able to correct rotational motion quite well, as the motion was periodic and stationary. Therefore there is relatively less residual motion in the gated PET data.

The model error was hypothesized to quantify the amount of residual motion in the reconstruction, such that it can be used to assess the goodness of fit of the motion model parameters. Figure 5.7 shows that the chosen model error as a function of one of the model parameters is smooth and convex and has a clear optimum, so the model error is very suitable for optimization of the model parameters. Figures 5.10 and 5.22 however do not show a significant increase in error if the model parameters are not constant, which was the case for experimental setup 5 and simulation 4, even though in the MBC reconstruction of experimental setup 5 residual motion was clearly visible in Figure 5.18. Also in the case of rotational motion (experimental setup 6), a significant increase in model error is not shown in Figure 5.22, although significant residual motion was visible in the reconstructions in Figure 5.20. Moreover, the absolute value of the model error in the simulations varies around 0.4 and differs greatly from the model error of the experiments that varies around 0.95, although the simulations and experiments used similar phantoms and similar motion dynamics. All of this suggests that the model error is not very sensitive to the residual motion in the reconstruction, but that its value is mainly driven by other effects. Chapter 7 discusses some effects that might cause this behavior and how these hypotheses might be tested.

However, all of the previous observations rely on the quality of the reconstructions and the quantitative analyses. When comparing the reconstructions of the experiments and simulations, those of the simulations clearly contain less noise although the sinograms of the simulations were scaled such that they contain approximately the same number of counts as a similar sinogram in the experiments. The difference might be that random and scatter coincidences were not modelled in the simulations. In the case of random coincidences, the reconstructed LOR can be placed anywhere. This will decrease contrast and thus image quality. The LORs of scattered coincidences also decrease the image quality.

Furthermore, when comparing simulations and experiments something that stands out is the fact that in the simulations the destructive effect of a model with non-constant parameters is not noticeable. Most likely the model instability in the simulations was too small, as the difference in internal motion amplitude between the first and second half of the simulation was only 2.5 pixels. Due to all of these unintentional differences between the simulations and the experiments, the experiments were performed in a more realistic setting than the simulations and the results of the experiments are therefore more valuable.

Another factor that could limit the reliability of the MBC reconstruction is the optimization method that was used. The results in Figure 5.21 show the performance of the optimization as

a function of iteration number. At the first model iteration the model parameters are already estimated to within one voxel size ($2 \times 2 \times 2 \text{ mm}^3$) of the corrected value. This indicates that it is not necessary to re-optimize the model parameters for every iteration of the reconstruction. If the model parameters are not optimized in an iteration, the model parameters from the previous iteration can still be used for motion correction. It might be sufficient to optimize the model parameters only once or twice during the reconstruction, which would drastically reduce the computation time for MBC reconstructions.

For experimental setups 2 and 4, the value of the parameters at iteration 5 and 9, respectively, deviate by a large amount with respect to the values at the other iterations. For these two experimental setups, the variation is much larger than the size of one voxel and can therefore not just be regarded as random variations. For setup 4 it was already hypothesized that this discontinuity in the optimization caused more residual motion to be present in the final reconstruction. Wrong model parameters cause a wrong amount of motion to be corrected, which directly influences the amount of residual motion in the reconstruction. It is therefore important that the model parameters converge towards the correct values without any discontinuities.

These discontinuities in the optimization could indicate that the optimization method is not functioning properly. In the case of Powell's method, the problem could originate from local minima or discontinuities, although given the shape of the model error as a function of the parameters this might not be very likely. The optimization can be checked by redoing the reconstruction and rather than using Powell's method for optimization, grid sampling the parameter space and just choosing the optimum. This will reveal whether Powell's method was not able to find the optimum, or the optimum was not in the expected place. The problem could also originate from the model error, which might not be chosen in the most effective way. More research is needed to pinpoint the cause of this unexplained behavior in the model parameter optimization.

The quantitative analyses and its reliability also demand some attention. In the analyses, the lesion volumes were segmented using an arbitrary threshold of 40% and then compared with the known, measured volumes of the lesions. Figure 5.23 shows that for the large lesion the volume is always overestimated. It also shows that for experimental setup 4 the volume of the non-corrected reconstruction is smaller than those of the other reconstructions, and even slightly smaller than the measured reference value. This is unexpected, as motion correction should reduce the lesion volume and the presence of motion increases tumor volume. These are all indications that the segmentation method that was used might not be optimal. Moreover, the segmentation method was not validated and it can thus not be guaranteed that even for a static reconstruction this method would yield the correct volume. Therefore, it is not optimal to compare volumes that were segmented with this method with a volume that was actually measured. Unfortunately there was not enough time available in this study to thoroughly vali-

date the segmentation method.

Our study had some limitations. All of the experimental results in this study were obtained using a scan duration of 10 seconds due to the large amount of computation time needed for an MBC reconstruction. In a clinical setting, a scan time of 1-3 minutes per bed position is used. Due to the longer scan time, the reconstructions in a clinical setting would contain less noise which will mainly affect the quality of respiratory gated and registered reconstructions. The qualities of these reconstructions would increase, therefore decreasing the difference in performance between those methods and MBC, although MBC is still expected to perform similarly to or better than both respiratory gating and registration of gated images in terms of lesion volume, mean SUV and CNR.

Another limitation of this study is that no attenuation correction was included in the reconstruction. Attenuation correction is performed by using the attenuation map that is acquired using the CT scan. Attenuation is however spatially dependent. The center of the phantom will therefore appear darker on the PET scans and the edges will appear brighter, as photons coming from the center have a higher chance of being attenuated. As the positions of the different tissues are time-dependent in the case of motion, the attenuation correction should also be adapted to include these dynamics.

Furthermore, the current motion model that is used in MBC was not designed for the correction of rotational motion. The results confirm this model is not able to correct rotations. It is therefore necessary to adjust the motion model accordingly. Likewise, the motion model should also be adjusted to include other types of motion, such as scaling and deformations, that are not included in the current motion model.

It was also shown that the current model error does not behave as an absolute error in the sense that its value does not reflect the amount of residual motion that is present in the reconstruction, limiting the use of the model error. Because the sinograms were binned in 0.5 second time frames, most of the elements in a sinogram contain 0 counts, some may contain 1 or maybe slightly more counts. In the calculation of the model error such a sinogram is subtracted from the forward projection of the image guess, which is often not 0 and therefore contributes to the model error. Our hypothesis is therefore that the current model error mainly reflects these low count statistics and only a very small part of the model error actually reflects the amount of motion in the reconstruction. A possible solution to this problem will be discussed in Chapter 7

Chapter 7

Outlook

For future research, the priority should be in optimizing the model error in order to obtain an absolute quality measure. As mentioned in Chapter 6, the problem might originate from low count statistics. In these low count regimes, noise has a large influence. A possible solution might be to compute the absolute difference of small regions rather than single voxels in the computation of the model error. By averaging over a small region the effect of noise will be reduced. The model error is however computed in sinogram space. By generating a binary mask of such a region in image space and projecting this binary mask to sinogram space it can be determined which parts of the sinogram contribute to that specific image space region. Further simulations and experiments should shed more light on this, perhaps by first testing the current model error in completely noise-free simulations to see if the model error then behaves as expected and noise is indeed the problem.

In order to test such a new error measure an experiment will have to be performed including a quantitative analyses. This experiment could also be used to re-assess the performance of model-based correction under non-stationary breathing dynamics, by using a more suitable optimization method and a more reliable quantitative analyses. The quantitative analyses can be improved upon by first of all adding the data of a static scan as a reference. Variations between different experiments can be reduced by taking an absolute value for thresholding of the tumor volume rather than determining a different threshold for each experiment.

Both in the model optimization and in the motion correction, attention should be paid to the boundaries of the field of view, as activity is moving out there at one end and volume of which no information is acquired is moving in at the other. This affects the computation of the model error and should therefore be corrected for. A possibility is to disregard the part of the sinogram of which no information is available in the computation of the model error. These voxels should then also be disregarded in the calculation of the model error. In other words, in Equation 4.5 the indices i and j should only run over the voxels on which information is available. In this study, the problem was avoided by constraining the motion such that the

entire image remained in the field of view of the detector.

A very important aspect of realistic data is that it comprises more than just a liver, usually (part of) a patient body. But as not all parts of the body move in the same way, the motion model will be spatially dependent. This spatial dependence cannot be modelled with the current motion model. Other classes of models exist that can intrinsically model such a spatial dependence, such as models that use motion vector fields. These motion vector fields describe for each voxel the motion between different time frames [77]. The sinogram of each individual time frame is however rather noisy in PET, which might make it difficult to estimate motion vector fields. A solution to this problem might be to combine voxels into small areas that contain more counts, rather than estimating a motion vector per voxel. This is the same solution as we proposed in this Chapter to possibly reduce the influence of noise on the model error.

Once the method contains an absolute quality measure, the boundaries of the field of view are taken into account and different parts of the patient body can be modeled independently, the method will be suitable to apply on patients that are treated with radioembolization. Patient studies can then be performed and using the absolute quality measure different motion models can be compared for different patients and different (stages of) pathologies. This will give a sense of the robustness of the proposed method. For instance, if a patient has had a liver ablation treatment and consequently misses part of the liver, the motion of the liver will be different. It is still an open question whether this will influence the correctness of the motion model and whether the motion model can be applied to all patients.

The next step could be to extend the method to other organs that undergo more complex motion, such as the lungs. For the lungs specifically this would enable better quantification of the lung shunt fraction in radioembolization. In general, extending the model to different organs will open up the possibility to apply MBC to PET imaging of any organ for which a surrogate signal can be acquired. For these purposes, the motion model should be extended such that it contains rotations and deformations.

This study focused on respiratory motion compensation in PET. The method could however also be used for respiratory motion compensation in SPECT imaging, as the two modalities are very similar. In SPECT imaging, the acquisition of the different projection angles is performed sequentially rather than simultaneously. This poses problems for respiratory gating. If a gating window contains more counts from certain projection angles, the reconstruction will contain artifacts [78].

In SPECT, gating is also applied to cardiac motion, using an electrocardiogram as a surrogate signal. In dual gated SPECT, the data is both cardiac gated and respiratory gated, such that even more data has to be ignored in order to correct for the different number of counts measured under each projection angle. MBC could be applied to cardiac motion once rotations are included in the motion model. Both problems could then be solved at once by optimizing

separate motion model parameters for cardiac and respiratory motion, and all available data could be used for the reconstruction.

Chapter 8

Conclusions

In conclusion, simulations and experiments have shown that it is feasible to optimize a motion model using PET data. Although the current model error is shown to be suitable for optimization, it does not serve as an absolute quality measure to compare the residual motion in different PET acquisitions. More research is needed to redesign the model error such that it can be used as an absolute quality measure. Simulations and experiments suggest that model-based motion correction performs similarly to respiratory gating and registration of gated images for stationary as well as non-stationary breathing dynamics. The latter was not expected, and might be attributed to unexpected behavior in the model parameter optimization. The parameter optimization should be better understood in order to draw a definite conclusion about the performance of model-based motion correction and its benefit for non-stationary breathing dynamics. The current motion model is not able to correct for rotational motion and should be adjusted accordingly for rotational motion.

Acknowledgements

I very much enjoyed doing this project and I would like to thank the people who made it possible. I would like to thank my supervisor Andre Mischke for helping me with the orientation on different projects and introducing me to Hugo de Jong. Andre, thank you for your guidance throughout this project.

I would like to thank Hugo for giving me the opportunity to do this project in his research group and introducing me to Remco Bastiaannet, my daily supervisor on this project. Hugo, I would also like to thank you for all the biweekly meetings that we had. Your feedback was always very helpful and your questions were always very thorough, helping me to both understand and communicate this project better.

Remco, many thanks go to you for your excellent supervision on a daily basis. Thank you for giving me the opportunity to work on this interesting project together. Thank you for always taking the time to answer my questions and helping me understand so many concepts that were completely new to me. Because of this I was able to greatly expand my knowledge and get to now a new field of research this year.

Lastly, I would like to thank the rest of the medical physics group at UMC Utrecht. I really enjoyed all the lunch breaks together and getting to know all of you, and to learn about your work in medical physics through the group meetings.

Appendix A

Unsuccessful techniques

As a documentation for future research, this section will list some of the techniques that were not (yet) successful or for some other reason were not included in the proposed method.

Reconstructions

Before using the Radon and inverse Radon transform as projector and backprojector in the simulations, we also tried to program a simple, one-ring PET detector. This was however rather complicated and takes a lot of time. Given the limited amount of time available for this project, it was considered not worth the time.

The second attempt for the simulations was to use STIR reconstruction software [74]. This however gave a lot of practical problems and unphysical values in the reconstruction. The problems seemed to arise from the motion model and transforming the image in the reconstruction. Although these problems could be worked around, it was not very convenient to have to deal with this while learning to do reconstructions in general. Therefore, the choice was made to continue the simulations in Python using the Radon transforms, to be sure that errors did not arise from software issues. STIR was however used for the reconstructions of the experiments, but only the projector and backprojector were used, as these were specific for the Siemens mCT PET/CT and did not cause the aforementioned problems.

GATE simulations

At the time when we still tried to use the STIR reconstruction software for the simulations, it was also tried to use GATE simulations. These output ROOT files and we attempted to reconstruct the ROOT data in STIR. However, STIR cannot handle ROOT files so these had to be converted to Interfile format. Despite two different attempts, we did not succeed in this necessary step. The GATE_PET_2_STIR platform only produced empty Interfiles, both on Linux and Windows. The `conv_GATE_row_ECAT_projdata_to_interfile` function of STIR was not used successfully

either, due to unclear documentation and confusion about the input parameters. By choosing to use the mathematical Radon transform rather than STIR reconstruction software for the reconstructions, the physical GATE simulations were no longer necessary and simulations could be performed in Python.

Post-radioembolization PET experiment

At the start of the project, an experiment was performed to characterize the image quality of a static post-radioembolization Y-90 PET scan. The image quality of a normal reconstruction and a gated reconstruction were compared. However, in practice there was no time left during this project to test the method under circumstances that are realistic for radioembolization, such as a very low count rate. This experiment was thus no longer of use for this project.

Optimization algorithms

Different optimization algorithms for motion compensation were used during this project. The first was brute force optimization by grid sampling the complete parameter space, which works but is rather slow. This was replaced by the Broyden-Fletcher-Goldfarb-Shanno algorithm, a type of quasi-Newtonian algorithm that estimates the Hessian [79]. This algorithm needs to be passed the Jacobian. However, if the Jacobian is not analytically available, the Jacobian will be calculated numerically.

For our model error it turned out that the Jacobian wasn't estimated correctly, which resulted in suboptimal performance of the algorithm. To circumvent this problem, the algorithm was replaced by Powell's method, which is slightly slower, but does not make use of derivative information and therefore performs better for the optimization of the model error.

Phantoms

Different simulation phantoms were tried throughout the project. We started with a block phantom, but this turned out to be inconvenient. The symmetry and small amount of detail of a block complicated noticing errors in the reconstruction.

Therefore, the block phantom was replaced by a Shepp-Logan phantom (see Section 2.3.2), since this phantom is not spherically symmetric and contains more detail. A drawback of the Shepp-Logan phantom is however that it doesn't resemble a liver. It is the exact opposite: it is a head phantom for CT, so most of the "activity" is located at the border of the phantom, rather than concentrated in a sphere somewhere in the middle.

When we discovered that the model error wasn't as sensitive as expected, we hypothesized that this was due to the fact that motion of 10-20 pixels (which was used at the time) is relatively

small compared to the size of the phantom. Due to the normalization, the model error wouldn't be influenced enough by such a small relative error. Therefore, it was decided to switch to a more realistic liver phantom.

More realistic (moving) phantoms of the human torso, such as the XCAT phantom, exist. This phantom is however only available in 3D and our simulations were in 2D. A new phantom was therefore created in such a way that it resembles a liver.

Hysteresis

Hysteresis is an effect that we looked into when we were looking for a new research topic. In the case of an hysteresis, the breathing motion does not follow the same path during inhale and exhale. We did a literature study on this effect and how it could be incorporated in a motion model.

Before fully implementing the method that we proposed to solve “the problem of hysteresis”, we found some of the few papers that describe the size of this effect in the liver. Hysteresis is on average 1 mm in the liver and is therefore not relevant for PET scanners that have a spatial resolution of a few mm. Instead, we decided to take the model error as a new research topic and investigate its behavior and use the model error to investigate the limitations of our current method.

Contrast-to-Noise Ratio (CNR)

In the case of the CNR, various choices were presented in choosing the background. This choice was complicated by the fact that the background is not homogeneous, due to attenuation and motion. It was tried to chose a fixed volume, at a fixed, homogeneous (for as far as possible) location and use this volume as the background for both spheres in the calculation of the CNR. However, due to the inhomogeneity of the background the CNR strongly varies when a different location for the background volume is chosen. It is thus not clear what volume should be chosen and how the value of the CNR should then be interpreted.

To overcome this problem, the entire liver could be segmented and then all of the liver volume excluding the two segmented tumors could be used as background. The CNR would then primarily indicate whether or not the two tumors can be distinguished from this large background. This is however not yet refined enough for our purpose. We do not just want the tumors to be visible, we also require their volumes to be motion corrected. If part of the smeared tumor would not be in the segmented tumor, it will end up in the background instead. This effect is called spill-out. But if the background is very large, this effect will be small and it will barely be noticeable that there was spill-out.

Therefore, in this case the background was chosen to be a ten-voxel band around the lesion. This way the CNR was calculated locally and local spill-out would be reflected in the CNR.

References

- [1] <http://www.hvrads.com/wp-content/uploads/2015/10/pet-ct.jpg>, . Retrieved: 2017-08-17.
- [2] KA Wood, PJ Hoskin, and MI Saunders. Positron emission tomography in oncology: a review. *Clinical Oncology*, 19(4):237–255, 2007.
- [3] Elena V Dizendorf, Brigitta G Baumert, Gustav K von Schulthess, Urs M Lütolf, and Hans C Steinert. Impact of whole-body 18f-fdg pet on staging and managing patients for radiation therapy. *Journal of Nuclear Medicine*, 44(1):24–29, 2003.
- [4] Angela van Baardwijk, Brigitta G Baumert, Geert Bosmans, Marinus van Kroonenburgh, Sigrid Stroobants, Vincent Gregoire, Philippe Lambin, and Dirk De Ruyscher. The current status of fdg-pet in tumour volume definition in radiotherapy treatment planning. *Cancer treatment reviews*, 32(4):245–260, 2006.
- [5] Andrew S Kennedy and Riad Salem. Radioembolization (yttrium-90 microspheres) for primary and metastatic hepatic malignancies. *The Cancer Journal*, 16(2):163–175, 2010.
- [6] Vanessa L Gates, Abdulredha AH Esmail, Karen Marshall, Stewart Spies, and Riad Salem. Internal pair production of 90y permits hepatic localization of microspheres using routine pet: proof of concept. *Journal of Nuclear Medicine*, 52(1):72–76, 2011.
- [7] Austin C Bourgeois, Ted T Chang, Yong C Bradley, Shelley N Acuff, and Alexander S Pasciak. Intraprocedural yttrium-90 positron emission tomography/ct for treatment optimization of yttrium-90 radioembolization. *Journal of Vascular and Interventional Radiology*, 25(2):271–275, 2014.
- [8] Ted T Chang, Austin C Bourgeois, Anastasia M Balius, and Alexander S Pasciak. Treatment modification of yttrium-90 radioembolization based on quantitative positron emission tomography/ct imaging. *Journal of Vascular and Interventional Radiology*, 24(3):333–337, 2013.
- [9] Andrew J Reader and Jeroen Verhaeghe. 4d image reconstruction for emission tomography. *Physics in medicine and biology*, 59(22):R371, 2014.

-
- [10] Chi Liu, Larry A Pierce II, Adam M Alessio, and Paul E Kinahan. The impact of respiratory motion on tumor quantification and delineation in static pet/ct imaging. *Physics in medicine and biology*, 54(24):7345, 2009.
- [11] Poonam Verma, Huanmei Wu, Mark Langer, Indra Das, and George Sandison. Survey: real-time tumor motion prediction for image-guided radiation treatment. *Computing in Science & Engineering*, 13(5):24–35, 2011.
- [12] Jamie R McClelland, David J Hawkes, Tobias Schaeffter, and Andrew P King. Respiratory motion models: a review. *Medical image analysis*, 17(1):19–42, 2013.
- [13] Simon Hughes, Jamie McClelland, Ségolène Tarte, Jane Blackall, Janet Liong, Shahreen Ahmad, David Hawkes, and David Landau. Assessment of respiratory cycle variability with and without training using a visual guide. *Cancer Therapy*, 6(2), 2008.
- [14] Toni Neicu, Ross Berbeco, John Wolfgang, and Steve B Jiang. Synchronized moving aperture radiation therapy (smart): improvement of breathing pattern reproducibility using respiratory coaching. *Physics in medicine and biology*, 51(3):617, 2006.
- [15] Arthur JAT Braat, Maarten LJ Smits, Manon NGJA Braat, Andor F van den Hoven, Jip F Prince, Hugo WAM de Jong, Maurice AAJ van den Bosch, and Marnix GEH Lam. 90y hepatic radioembolization: an update on current practice and recent developments. *Journal of Nuclear Medicine*, 56(7):1079–1087, 2015.
- [16] <https://www.sirtex.com/us/clinicians/about-sir-spheres-microspheres/>. Accessed: 2017-07-16.
- [17] <https://www.btg-im.com/en-NL/TheraSphere/Home>. Accessed: 2017-07-16.
- [18] Andor F van den Hoven, Maarten LJ Smits, Charlotte ENM Rosenbaum, Helena M Verkooijen, Maurice AAJ van den Bosch, and Marnix GEH Lam. The effect of intra-arterial angiotensin ii on the hepatic tumor to non-tumor blood flow ratio for radioembolization: a systematic review. *PLoS One*, 9(1):e86394, 2014.
- [19] http://images.slideplayer.com/26/8468520/slides/slide_1.jpg. Retrieved: 2017-07-16.
- [20] Lars Jødal. Beta emitters and radiation protection. *Acta Oncologica*, 48(2):308–313, 2009.
- [21] Maurits Wondergem, Maarten LJ Smits, Mattijs Elschot, Hugo WAM de Jong, Helena M Verkooijen, Maurice AAJ van den Bosch, Johannes FW Nijsen, and Marnix GEH Lam. 99mTc-macroaggregated albumin poorly predicts the intrahepatic distribution of 90y resin microspheres in hepatic radioembolization. *Journal of Nuclear Medicine*, 54(8):1294–1301, 2013.

- [22] Hojjat Ahmadzadehfar, Amir Sabet, Kim Biermann, Marianne Muckle, Holger Brockmann, Christiane Kuhl, Kai Wilhelm, Hans-Jürgen Biersack, and Samer Ezziddin. The significance of ^{99m}Tc -*maab* *SPECT/CT* liver perfusion imaging in treatment planning for ^{90}Y -microsphere selective internal radiation treatment. *Journal of Nuclear Medicine*, 51(8):1206–1212, 2010.
- [23] Monia E Hamami, Thorsten D Poeppel, Stephan Müller, Till Heusner, Andreas Bockisch, Philipp Hilgard, and Gerald Antoch. *SPECT/CT* with ^{99m}Tc -*maab* in radioembolization with ^{90}Y microspheres in patients with hepatocellular cancer. *Journal of Nuclear Medicine*, 50(5):688–692, 2009.
- [24] Yung-Hsiang Kao, Jeffrey D Steinberg, Young-Soon Tay, Gabriel KY Lim, Jianhua Yan, David W Townsend, Charley A Budgeon, Jan A Boucek, Roslyn J Francis, Timothy ST Cheo, et al. Post-radioembolization yttrium-90 *PET/CT*-part 2: dose-response and tumor predictive dosimetry for resin microspheres. *EJNMMI research*, 3(1):57, 2013.
- [25] Shyam M Srinivas, Navin Natarajan, Joshua Kuroiwa, Sean Gallagher, Elie Nasr, Shetal N Shah, Frank P DiFilippo, Nancy Obuchowski, Bana Bazerbashi, Naichang Yu, et al. Determination of radiation absorbed dose to primary liver tumors and normal liver tissue using post-radioembolization ^{90}Y *PET*. *Frontiers in oncology*, 4, 2014.
- [26] RG Selwyn, RJ Nickles, BR Thomadsen, LA DeWerd, and JA Micka. A new internal pair production branching ratio of ^{90}Y : the development of a non-destructive assay for ^{90}Y and ^{90}Sr . *Applied radiation and isotopes*, 65(3):318–327, 2007.
- [27] C Chiesa, M Maccauro, R Romito, C Spreafico, S Pellizzari, A Negri, C Sposito, C Morosi, E Civelli, R Lanocita, et al. Need, feasibility and convenience of dosimetric treatment planning in liver selective internal radiation therapy with (^{90}Y) microspheres: the experience of the national tumor institute of milan. *The quarterly journal of nuclear medicine and molecular imaging: official publication of the Italian Association of Nuclear Medicine (AIMN)[and] the International Association of Radiopharmacology (IAR),[and] Section of the Society of..*, 55(2):168–197, 2011.
- [28] Simon R Cherry, James A Sorenson, and Michael E Phelps. *Physics in nuclear medicine*. Elsevier Health Sciences, 2012.
- [29] Arthur H Compton. A quantum theory of the scattering of x-rays by light elements. *Physical review*, 21(5):483, 1923.
- [30] Paul A Llewellyn Ralph A Tipler. *Modern physics international edition*. John Wiley & Sons, 6th edition, 2012.

- [31] <http://www.nuclear-power.net/wp-content/uploads/2015/03/attenuation.png>. Retrieved: 2017-07-20.
- [32] <https://www.healthcare.siemens.nl/molecular-imaging/pet-ct/biograph-truepoint-petct>. Retrieved: 2017-07-18.
- [33] E.M. Henley and A. Garcia. *Subatomic Physics*. World Scientific, 2007. ISBN 9789813101678. URL <https://books.google.nl/books?id=j847DQAAQBAJ>.
- [34] <https://en.wikipedia.org/wiki/Photomultiplier>. Retrieved: 2017-07-20.
- [35] <https://www.radiologycafe.com/images/physics/nm-pet-lineofresponse.png>. Retrieved: 2017-07-23.
- [36] http://rochrad.com/images/pet_ctscan.jpg, .
- [37] Gengsheng Lawrence Zeng. *Medical image reconstruction*. Springer, 2010.
- [38] Warren Campbell. Sensitivity improvements in tof pet given 20 ps resolution. 2009. URL https://www.researchgate.net/publication/281296244_Sensitivity_improvements_in_TOF_PET_given_20_ps_resolution.
- [39] Lawrence A Shepp and Yehuda Vardi. Maximum likelihood reconstruction for emission tomography. *IEEE transactions on medical imaging*, 1(2):113–122, 1982.
- [40] Alvaro R De Pierro. A modified expectation maximization algorithm for penalized likelihood estimation in emission tomography. *IEEE transactions on medical imaging*, 14(1):132–137, 1995.
- [41] Frederic H Fahey. Data acquisition in pet imaging. *Journal of nuclear medicine technology*, 30(2):39–49, 2002.
- [42] Mónica Vieira Martins. Positron emission mammography. In *Mammography Techniques and Review*. InTech, 2015. URL <https://www.intechopen.com/books/mammography-techniques-and-review/positron-emission-mammography>.
- [43] Doumit Daou. Respiratory motion handling is mandatory to accomplish the high-resolution pet destiny. *European journal of nuclear medicine and molecular imaging*, 35(11):1961–1970, 2008.
- [44] Paul J Keall, Gig S Mageras, James M Balter, Richard S Emery, Kenneth M Forster, Steve B Jiang, Jeffrey M Kapatoes, Daniel A Low, Martin J Murphy, Brad R Murray, et al. The management of respiratory motion in radiation oncology report of aapm task group 76. *Medical physics*, 33(10):3874–3900, 2006.

- [45] Ismet Sarikaya, Henry WD Yeung, Yusuf Erdi, and Steven M Larson. Respiratory artefact causing malpositioning of liver dome lesion in right lower lung. *Clinical nuclear medicine*, 28(11):943–944, 2003.
- [46] Ruijiang Li, John H Lewis, Xun Jia, Tianyu Zhao, Weifeng Liu, Sara Wuenschel, James Lamb, Deshan Yang, Daniel A Low, and Steve B Jiang. On a pca-based lung motion model. *Physics in medicine and biology*, 56(18):6009, 2011.
- [47] Luca Guerra, Elena De Ponti, Federica Elisei, Valentino Bettinardi, Claudio Landoni, Maria Picchio, Maria Carla Gilardi, Annibale Versari, Federica Fioroni, Mirosław Dziuk, et al. Respiratory gated pet/ct in a european multicentre retrospective study: added diagnostic value in detection and characterization of lung lesions. *European journal of nuclear medicine and molecular imaging*, 39(9):1381–1390, 2012.
- [48] Andrea Lupi, Marta Zaroccolo, Matteo Salgarello, Veronica Malfatti, and Pierluigi Zanco. The effect of 18f-fdg-pet/ct respiratory gating on detected metabolic activity in lung lesions. *Annals of nuclear medicine*, 23(2):191–196, 2009.
- [49] AM García Vicente, AM Soriano Castrejón, MP Talavera Rubio, AA León Martín, AM Palomar Muñoz, JP Pilkington Woll, and VM Poblete García. 18f-fdg pet-ct respiratory gating in characterization of pulmonary lesions: approximation towards clinical indications. *Annals of nuclear medicine*, 24(3):207–214, 2010.
- [50] Matthias K Werner, J Anthony Parker, Gerald M Kolodny, Jeffrey R English, and Matthew R Palmer. Respiratory gating enhances imaging of pulmonary nodules and measurement of tracer uptake in fdg pet/ct. *American Journal of Roentgenology*, 193(6):1640–1645, 2009.
- [51] B Thorndyke, E Schreibmann, P Maxim, B Loo, A Boyer, A Koong, and L Xing. Tu-d-j-6c-08: Enhancing 4d pet through retrospective stacking. *Medical Physics*, 32(6):2096–2096, 2005.
- [52] Tobias Klinder, Cristian Lorenz, and Jörn Ostermann. Prediction framework for statistical respiratory motion modeling. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010*, pages 327–334, 2010.
- [53] Florian Büther, Mohammad Dawood, Lars Stegger, Frank Wübbeling, Michael Schäfers, Otmar Schober, and Klaus P Schäfers. List mode–driven cardiac and respiratory gating in pet. *Journal of Nuclear Medicine*, 50(5):674–681, 2009.
- [54] Christian Buerger, Rachel E Clough, Andrew P King, Tobias Schaeffter, and Claudia Prieto. Nonrigid motion modeling of the liver from 3-d undersampled self-gated golden-radial phase encoded mri. *IEEE transactions on medical imaging*, 31(3):805–815, 2012.

- [55] Tobias Klinder and Cristian Lorenz. Respiratory motion compensation for image-guided bronchoscopy using a general motion model. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 960–963. IEEE, 2012.
- [56] Andrew P King, Christian Buerger, Charalampos Tsoumpas, Paul K Marsden, and Tobias Schaeffter. Thoracic respiratory motion estimation from mri using a statistical model and a 2-d image navigator. *Medical image analysis*, 16(1):252–264, 2012.
- [57] Qinghui Zhang, Alex Pevsner, Agung Hertanto, Yu-Chi Hu, Kenneth E Rosenzweig, C Clifton Ling, and Gig S Mageras. A patient-specific respiratory model of anatomical motion for radiation treatment planning. *Medical Physics*, 34(12):4772–4781, 2007.
- [58] M Filipovic, P-A Vuissoz, A Codreanu, M Claudon, and J Felblinger. Motion compensated generalized reconstruction for free-breathing dynamic contrast-enhanced mri. *Magnetic resonance in medicine*, 65(3):812–822, 2011.
- [59] Floris Ernst, Volker Martens, Stefan Schlichting, Armin Beširević, Markus Kleemann, Christoph Koch, Dirk Petersen, and Achim Schweikard. Correlating chest surface motion to motion of the liver using ε -svr—a porcine study. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2009*, pages 356–364, 2009.
- [60] Jan Ehrhardt, René Werner, Dennis Säring, Thorsten Frenzel, Wei Lu, Daniel Low, and Heinz Handels. An optical flow based method for improved reconstruction of 4d ct data sets acquired during free breathing. *Medical Physics*, 34(2):711–721, 2007.
- [61] Wei Lu, Daniel A Low, Parag J Parikh, Michelle M Nystrom, Issam M El Naqa, Sasha H Wahab, Maureen Handoko, David Fooshee, and Jeffrey D Bradley. Comparison of spirometry and abdominal height as four-dimensional computed tomography metrics in lung. *Medical physics*, 32(7):2351–2357, 2005.
- [62] Yvette Seppenwoolde, Ross I Berbeco, Seiko Nishioka, Hiroki Shirato, and Ben Heijmen. Accuracy of tumor motion compensation algorithm from a robotic respiratory tracking system: a simulation study. *Medical physics*, 34(7):2774–2784, 2007.
- [63] D Ruan, Jeffrey A Fessler, JM Balter, RI Berbeco, S Nishioka, and H Shirato. Inference of hysteretic respiratory tumor motion from external surrogates: a state augmentation approach. *Physics in medicine and biology*, 53(11):2923, 2008.
- [64] JR McClelland, S Hughes, M Modat, A Qureshi, S Ahmad, DB Landau, S Ourselin, and DJ Hawkes. Inter-fraction variations in respiratory motion models. *Physics in medicine and biology*, 56(1):251, 2010.

- [65] Floris Ernst, Ralf Bruder, Alexander Schlaefer, and Achim Schweikard. Correlation between external and internal respiratory motion: a validation study. *International journal of computer assisted radiology and surgery*, 7(3):483–492, 2012.
- [66] JR McClelland, AG Chandler, JM Blackall, S Ahmad, DB Landau, and DJ Hawkes. 4d motion models over the respiratory cycle for use in lung cancer radiotherapy planning. In *Medical Imaging*, pages 173–183. International Society for Optics and Photonics, 2005.
- [67] AE Torshabi, Andrea Pella, Marco Riboldi, and Guido Baroni. Targeting accuracy in real-time tumor tracking via external surrogates: a comparative study. *Technology in cancer research & treatment*, 9(6):551–561, 2010.
- [68] Andrew Peter King, Kawal S Rhode, Reza S Razavi, and Tobias R Schaeffter. An adaptive and predictive respiratory motion model for image-guided interventions: theory and first clinical application. *IEEE transactions on medical imaging*, 28(12):2020–2032, 2009.
- [69] AP King, C Tsoumpas, C Buerger, V Schulz, P Marsden, and T Schaeffter. Real-time respiratory motion correction for simultaneous pet-mr using an mr-derived motion model. In *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE*, pages 3589–3594. IEEE, 2011.
- [70] Sarah J McQuaid, Tryphon Lambrou, Vincent J Cunningham, Valentino Bettinardi, Maria Carla Gilardi, and Brian F Hutton. The application of a statistical shape model to diaphragm tracking in respiratory-gated cardiac pet images. *Proceedings of the IEEE*, 97(12):2039–2052, 2009.
- [71] Florian Büther, Thomas Vehren, Klaus P Schäfers, and Michael Schäfers. Impact of data-driven respiratory gating in clinical pet. *Radiology*, 281(1):229–238, 2016.
- [72] Florian Büther, Iris Ernst, Mohammad Dawood, Peter Kraxner, Michael Schäfers, Otmar Schober, and Klaus P Schäfers. Detection of respiratory tumour motion using intrinsic list mode-driven gating in positron emission tomography. *European journal of nuclear medicine and molecular imaging*, 37(12):2315–2327, 2010.
- [73] William H Press. *Numerical recipes in C: The art of scientific computing*. Cambridge university press, 2nd edition, 1992.
- [74] Kris Thielemans, Charalampos Tsoumpas, Sanida Mustafovic, Tobias Beisel, Pablo Aguiar, Nikolaos Dikaios, and Matthew W Jacobson. Stir: software for tomographic image reconstruction release 2. *Physics in medicine and biology*, 57(4):867, 2012.

-
- [75] Yusuf E Erdi, Osama Mawlawi, Steven M Larson, Massimo Imbriaco, Henry Yeung, R Finn, and John L Humm. Segmentation of lung lesion volume by adaptive positron emission tomography image thresholding. *Cancer*, 80(S12):2505–2509, 1997.
- [76] Albert Rose. *Vision: Human and Electronic*. Plenum Press, 1973.
- [77] Deshan Yang, Wei Lu, Daniel A Low, Joseph O Deasy, Andrew J Hope, and Issam El Naqa. 4d-ct motion estimation using deformable image registration and 5d respiratory motion modeling. *Medical physics*, 35(10):4577–4590, 2008.
- [78] Joyoni Dey, William P Segars, P Hendrik Pretorius, Ronn P Walvick, Philippe P Bruyant, Seth Dahlberg, and Michael A King. Estimation and correction of cardiac respiratory motion in spect in the presence of limited-angle effects due to irregular respiration. *Medical physics*, 37(12):6453–6465, 2010.
- [79] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.