

Event-Driven Motion Compensation in Positron Emission Tomography: Development of a Clinically Applicable Method

Inaugural dissertation submitted in partial
fulfillment of the requirements for the degree
DOCTOR RERUM MEDICINALIUM (DR. RER. MEDIC.)



**TECHNISCHE
UNIVERSITÄT
DRESDEN**



Faculty of Medicine Carl Gustav Carus
University of Technology Dresden, Germany

submitted by

Jens Langner

from Stuttgart a.N.

Dresden 2008

Copyright © Jens Langner, MMVIII. All rights reserved.

Jens.Langner@light-speed.de
<http://www.jens-langner.de/>

The author hereby grants full permission to the *Technische Universität Dresden, Germany* (TUD) and the *Forschungszentrum Dresden-Rossendorf, Dresden, Germany* (FZD) to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

1. Gutachter: Prof. Dr. hum. biol. habil. Dr. rer. nat. Jörg van den Hoff

2. Gutachter: Prof. Dr.-Ing. Hans Herzog

Tag der mündlichen Prüfung: 28. Juli 2009

gez: Prof. Dr. med. G. Wozel
Vorsitzender der Promotionskommission

„Wenn es zutreffen sollte, daß ich nicht nur weiß, was schlimm und häßlich, sondern auch, was schön ist, so verdanke ich diese Gabe dem Glück, in Dresden aufgewachsen zu sein. Ich mußte, was schön sei, nicht erst aus Büchern lernen. Nicht in der Schule, und nicht auf der Universität. Ich durfte die Schönheit einatmen wie Försterkinder die Waldluft.“

Erich Kästner
(Als ich ein kleiner Junge war)

Acknowledgments

Many people have contributed to the success of this thesis. Therefore, I would like to take the opportunity to express my deep gratitude for the tremendous amount of support I have received.

First of all, I would like to thank the head of the PET centre of the Forschungszentrum Dresden-Rossendorf, Prof. Dr. Jörg van den Hoff for the superb supervision of this thesis. His profound knowledge of PET and physics, the numerous constructive discussions as well as the scope for development have been the most important reasons for the success of this work. Furthermore, I would like to especially thank Prof. Dr. Hans Herzog of the Forschungszentrum Jülich for his spontaneous willingness for an expertise on this thesis.

Thanks also go to the former head of the Institute of Bioanorganic and Radiopharmaceutical Chemistry, Prof. Dr. Bernd Johannsen as well as to the head of the Institute of Radiopharmacy, Prof. Dr. Jörg Steinbach for the possibility to perform the thesis at the Forschungszentrum Dresden-Rossendorf and thus for the necessary funding. Furthermore, I would like to thank the head of the Clinic of Nuclear Medicine of the Technical University Dresden, Prof. Dr. Jörg Kotzerke for his supplemental approval of this thesis.

Special thanks go to Dr. Frank Hofheinz and Dr. Paul Bühler who have supported this thesis greatly with their knowledge and who encouraged me to work on motion compensation initially. In addition, I would like to thank Dr. Bettina Beuthien-Baumann and Liane Oehme of the Clinic of Nuclear Medicine for the willingness to answer my medical questions and for supporting the clinical evaluation part of this thesis.

For the professional discussions and for the continued technical support I would like to thank Dr. Christian Pöttsch, Hagen Mölle, Dr. Sören Dittrich, Dr. Uwe Just and Dr. Edmund Will of the PET department. Thanks also go to Tobias Seiler for his support during hardware development of the adapter card. Furthermore, I would like to thank Stefan Ritter and Dr. Jens Tröger for their numerous corrections of grammar and spelling.

My parents, Marlies and Anton Zlatoper, my brother Heiko Langner and my father Norbert Langner have to be acknowledged for their continued support and belief in me. Finally I would like to thank Sylvia for her endless patience, encouragement and breathtaking love during this important step of my life.

Contents

Introduction	1
Current State of Research	5
1 Fundamentals	7
1.1 Positron Emission Tomography	7
1.1.1 Data Registration Process	9
1.1.2 Image Reconstruction	18
1.2 Motion Susceptibility	21
1.2.1 Motion Types	23
1.2.2 Motion Quantification	24
1.2.3 Motion Compensation	29
1.3 Medical Nomenclature	34
2 Material and Methods	37
2.1 Material	37
2.1.1 PET Scanner: ECAT Exact HR ⁺	37
2.1.2 Motion Tracking: ARTtrack	41
2.2 Optimised List-Mode Data Acquisition	42
2.2.1 Shared Storage Device Acquisition	43
2.2.2 DAQ-based Data Acquisition	46
2.2.3 Event-based Attenuation Measurement	49
2.3 Motion Quantification	52
2.3.1 Cross-Calibration	52
2.3.2 Motion Target	56
2.3.3 Motion Analysis	59
2.4 Event-Driven Motion Compensation	64
2.4.1 Normalisation Correction	65
2.4.2 LOR Discretisation Correction	66
2.4.3 Out-of-FOV Correction	70
2.5 Clinical Integration	77
2.6 Clinical Evaluation	81
3 Results	87
3.1 Optimised List-Mode Data Acquisition	87
3.1.1 Shared Storage Device Acquisition	87

3.1.2	DAQ-based Data Acquisition	88
3.1.3	Event-based Attenuation Measurement	89
3.2	Motion Quantification	91
3.2.1	Cross-Calibration	91
3.2.2	Motion Target	93
3.2.3	Motion Analysis	98
3.3	Event-Driven Motion Compensation	103
3.3.1	Normalisation Correction	103
3.3.2	LOR Discretisation Correction	104
3.3.3	Out-of-FOV Correction	105
3.4	Clinical Integration	108
3.5	Clinical Evaluation	110
4	Discussion	117
4.1	Optimised List-Mode Data Acquisition	117
4.1.1	Shared Storage Device Acquisition	117
4.1.2	DAQ-based Data Acquisition	118
4.1.3	Event-based Attenuation Measurement	119
4.2	Motion Quantification	120
4.2.1	Cross-Calibration	120
4.2.2	Motion Target	122
4.2.3	Motion Analysis	123
4.3	Event-Driven Motion Compensation	124
4.3.1	Normalisation Correction	125
4.3.2	LOR Discretisation Correction	125
4.3.3	Out-of-FOV Correction	126
4.4	Clinical Integration	127
4.5	Clinical Evaluation	127
	Summary	131
	Theses	135
	List of Figures	137
	List of Tables	141
	Abbreviations	143
	Bibliography	145

Introduction

There is nothing more difficult to take in hand, more perilous to conduct or more uncertain in its success than to take the lead in the introduction of a new order of things.

(Niccolo Machiavelli)

During the last decades, *positron emission tomography* (PET) has emerged as an important technique for non-invasive imaging of metabolic processes in living organisms. In oncology it has become an indispensable medical imaging method for the diagnosis and monitoring of cancer. Furthermore, PET allows to quantify e.g. transport processes in the brain. Therefore, it serves as an excellent assistant technique for the diagnosis and monitoring of frequent neurological questions such as *Parkinson's disease* or *Alzheimer's disease*.

The imaging process of PET is based on the radioactive tracer principle developed by *George de Hevesy* which he was honoured for with the *Nobel Prize* in 1943. In PET, a radiopharmaceutical labelled with a positron-emitting isotope (tracer) is brought into the metabolism of the examined organism. In oncology the most common radiopharmaceutical is the ^{18}F labelled *fluorodeoxyglucose* (FDG). This tracer shares common metabolic pathways with glucose. The distribution of FDG and thus the increased glucose metabolism of most tumours can be visualised with a good sensitivity and spatial resolution in PET [1].

The possibility to visualise metabolic processes in living organisms makes PET a *functional* imaging method. This is an important contrast to other tomography techniques such as *computed tomography* (CT) or *magnetic resonance imaging* (MRI) which are primarily morphological oriented methods. The increasing importance of functional imaging – and thus PET – can also be seen by recent developments in the field of multimodality imaging. In the last few years combined PET/CT systems have become widely available and by now have superseded the existing sales of pure PET systems. In addition, several research groups have started to study the possibilities and potential advantages of combining PET and MRI systems. The first promising results of those studies suggest that combined PET/MR systems will become available in the next years [2, 3].

Another key strength of PET is the possibility to conduct quantitative analyses of the image data. This is possible because the number of registered radioactive decays per voxel corresponds directly to the local concentration of the tracer in the tissue. This allows to compare the injected activity to the imaged activity concentration in a specific *region-of-interest* (ROI). Thus conclusions concerning the physiology of the examined organism can be drawn. In addition, a PET system allows to investigate the time course of tracer accumulation in a ROI thanks to *dynamic acquisitions*. This is achieved by the calculation of *time-activity-curves* (TAC). These

curves are then, for example, used within a *tracer kinetics analysis* for the modelling of the transport processes in the brain by application of *compartment models* [4].

The three-dimensional delineation and determination of volumes is another essential technique required in PET. The exact delineation of the volume of a structure in the examined subject is important for various medical questions. In oncology, for example, the alteration of the vital tumour volume during *follow-up* examinations is an important information for assessing the success or failure of a cancer therapy [5–7]. This is another reason why hybrid imaging is becoming more and more important. In a PET/CT scan, for example, the morphology of a lesion seen in the CT scan can be combined with the functional information seen in the PET scan. This allows to optimise the available therapy strategies because conclusions on the vitality of a tumour can be drawn with the additional functional information.

A precise delineation of the ROIs by a physician is a fundamental necessity for utilising the aforementioned advantages of functional imaging. Apart from factors that are related to the subjective assessment of the observer who is analysing the PET data (e.g. by use of threshold-based ROI delineation methods), several factors of the imaging technique itself may add to a potentially incorrect delineation of ROIs or assessment of functional processes. Common among those factors is the presence of patient motion during an examination. As PET is based on the continuous observation of functional processes over a long period of time (up to 20 to 90 minutes), the occurrence of patient motion is very likely during a typical examination. In addition, in clinical PET the amount of applied radioactivity is generally kept low because of radiation protection regulations.

Since the susceptibility to motion is a well-known limitation of PET [8], various approaches for the avoidance and compensation of patient motion have been developed in the past. These include, *inter alia*

- patient immobilisation or fixation solutions [9–11],
- gated acquisition techniques [12],
- multiple acquisition framing techniques [13, 14],
- event-driven motion compensation techniques [15, 16].

Among these techniques the *event-driven motion compensation* approach has the highest potential for a precise compensation of the patient motion in PET data. In this context ‘event-driven’ refers to the processing of each coincidence information (event). In fact, an event-driven motion compensation operates on the raw data of a PET scanner because of spatial reorientation of the *line-of-response* (LOR) of each coincidence event prior to the image reconstruction. The *gated acquisition*, *multiple acquisition framing* and most other motion compensation methods (apart from the immobilisation and fixation solutions) try to compensate the motion in the sinogram or image space instead. While this approach is easier to apply, all of these techniques are, however, more limited in their ability to compensate patient motion than the event-driven methods. A major drawback is, for example, that these image-based methods often do not allow to correct for motion throughout the complete acquisition. Instead, they compensate patient motion by storing the acquisition data into separate time intervals according to the state or

magnitude of motion (frames/gates). This has several disadvantages such as a lower amount of counts per frame or gate. In addition, a spatial alignment of each frame/gate to the others is often not easily possible or leads to additional inaccuracies because of required voxel interpolations. Therefore, not all motion which occurs within a frame or gate is fully compensated.

In comparison, the event-driven compensation technique has a clear advantage over the sino-gram or image-based approaches. The ability to operate on the raw data does not only suggest a high potential for accurate motion compensation, it also offers greater flexibility. This is the main reason for this work to concentrate on the development of an event-driven compensation technique. However, one drawback is the fact that the complexity of event-driven methods is considerably higher. For an event-driven motion compensation, for example, several post-corrections have to be applied. The three most important post-corrections, which are applied after spatial reorientation of the LORs, are correcting for

- differences in detector sensitivity – (*Normalisation correction*),
- the ambiguity of the LOR to detector mapping – (*LOR discretisation correction*),
- LORs falling outside the *field-of-view* (FOV) – (*Out-of-FOV correction*).

Amongst these methods, the Out-of-FOV correction is the most complex one and requires a large amount of processing resources. Moreover, this post-correction often causes artefacts in the resulting image data. Thus, the Out-of-FOV post-correction is one of the main reasons that an event-driven motion compensation has not yet been integrated into clinical PET.

Other aspects of the instrumentation for acquiring data with a PET scanner or the routine acquisition of motion data have hindered a clinical application as well. To be able to use an event-driven method the acquisition of *list-mode* data is a necessity. However, many PET scanners, like the ECAT Exact HR⁺ PET scanner (Siemens/CTI, Knoxville, Tennessee), do not allow for the processing of a clinical acquisition with a direct or fast access to list-mode data. This is unfortunately also the case for many newer PET scanner models and is often the reason that those event-driven methods cannot easily be applied.

In addition, if an optical tracking system is used for the tracking of patient motion, it is usually not available with a clinically usable motion target. As the main market for such optical tracking systems is the motion animations market, tracking targets for a clinical application are rarely available or not at all.

A detailed analysis of the patient motion prior to motion compensation is another important requirement for a potential integration into clinical use. Such a motion analysis allows to distinguish ‘significant’ from ‘non-significant’ motion and thus applies the motion compensation only to relevant movements of the patient. Otherwise a motion compensation would have to operate on all registered motion samples and would require an unacceptable amount of time. A popular approach for such motion filtering is to apply a threshold to each parameter (rotations, translations) of the motion data. However, the assessment of the magnitude of three-dimensional motion via such an evaluation of each parameter is critical because of the combinatorial relation of the raw motion parameters.

The main goal of the present thesis has been to develop an event-driven motion compensation method that resolves the existing shortcomings. One important focus of our work has been on

the integration of this method into clinical routine. For this purpose we concentrated on the following key aspects:

- to develop hardware and software methods for a clinical acquisition of list-mode data (for an ECAT Exact HR⁺),
- to develop and evaluate a head motion tracking target for clinical use,
- to develop a method for a quantitative analysis and evaluation of the motion data,
- to develop and resolve the aforementioned shortcomings in the event-driven motion compensation method described by Bühler et al. [15],
- to provide the necessary software tools for an integration of the motion compensation into clinical routine,
- to evaluate the final implementation of our event-driven motion compensation with clinical data.

According to these points the present thesis has been structured as follows: After a short outline of the current state of research, the fundamentals of PET and its general susceptibility to patient motion are outlined in chapter 1.

In section 2.2 (*Optimised List-Mode Data Acquisition*) our methods to improve the list-mode capabilities of an ECAT Exact HR⁺ PET scanner are presented. Furthermore, section 2.3 (*Motion Quantification*) describes the development and evaluation of a clinically usable head motion target as well as our methods for a quantitative analysis of patient motion. In section 2.4 (*Event-Driven Motion Compensation*) the implementation and details of the post-correction improvements of our event-driven motion compensation are explained. Our software tools for an integration of the motion compensation in clinical routine are presented in section 2.5 (*Clinical Integration*). Finally, section 2.6 (*Clinical Evaluation*) concludes chapter 2 with a quantitative evaluation of our event-driven motion compensation method applied to clinical data.

The results of all aforementioned methods and evaluations are presented in chapter 3. Furthermore, the discussion on the results can be found in chapter 4 and the thesis closes with a summary.

Current State of Research

It takes a great deal of history to produce a little literature

(Henry James)

The problem of patient motion during a medical imaging examination as well as its compensation is a topic widely covered in the literature. In Bettinardi et al. [9], Shrawder et al. [10] and Sweeney et al. [11] it has been shown that the early focus has been mainly on trying to avoid patient motion, e.g., by application of immobilisation or fixation systems such as thermoplastic masks or vacuum cushions. This has changed rather dramatically since the mid 1990s because of improvements in computer processing speed and increasing data storage capacities. In 1994 Murray et al. [17] proposed the use of an optical camera system for the tracking of motion with *optical flow algorithms* extracting the motion parameters by constant inspection of the pixel intensities in captured two-dimensional image data. In parallel specialised motion tracking devices began to be used in medicine. Strasburger and Klenk [18] proposed the use of an opto-electronic belt for recording the respiration phases in psychophysiological experimentation and therapy as early as 1983. This technique has been adopted by a number of companies in the 1990s, which started to sell similar belt systems for the general tracking of respiratory motion. In addition, optical tracking devices which initially had been developed for computerised animations, found their way into medical applications. In Picard and Thompson [19] and Lopresti et al. [20], various optical tracking solutions have been proposed for tracking of the patient's head motion. Similar to the devices for tracking of respiratory motion, a number of companies took up the development of those kinds of optical tracking systems.

At the same time, several research groups started working on methods for compensation of patient motion in PET. Daube-Witherspoon et al. [21] first proposed a method to rebin the sinogram data of a PET study in 1990 by acquiring the data in *list-mode* along with the head position of the patient. Similarly, Menke et al. [22] proposed a method in 1996 which uses a *look-up table* to select the appropriate sinogram location according to the head position. As these methods, however, required much disk space, Picard and Thompson [13] proposed a simplified motion compensation method by using *multiple acquisition frames* so that a new image is generated for each relevant patient motion and later aligned to the others via rigid-body transformations. While all of these methods already used list-mode data for sorting the acquisition data dynamically into separate frames according to the patient motion, Bühler et al. [15] and Bloomfield et al. [16] first proposed new methods in 2004 for taking advantage of the full potential of the list-mode data of a PET scanner. Instead of applying the motion compensation *after* binning each line-of-response (LOR) into a sinogram, these methods introduced the principle of spatially reorienting each LOR *before* the actual sinogram binning process. This is a new

approach to a potentially most accurate correction of patient motion in PET. It became feasible because of the recent improvements in information technology which allowed the processing of multiple gigabyte list-mode data within a manageable time.

The present thesis builds upon these methods and offers an optimised version of an event-driven motion compensation method for brain PET imaging. Moreover, it enhances earlier compensation approaches by resolving existing shortcomings and by introducing optimisations to the motion compensation algorithms. Furthermore, our methods focus on the aspects required for integration into clinical routine. This includes methods for the clinical acquisition of list-mode which were published and presented as Langner et al. [23–28] in 2005 and 2006. Furthermore, we integrated methods for a quantitative analysis of patient motion to identify ‘significant’ motion, and presented these results in 2007 and 2008 as Langner et al. [29, 30] at several conferences. Finally, to achieve a fast processing of the motion compensation and to resolve the existing shortcomings in the compensation methods of Bühler et al. and Bloomfield et al., the present work combines several methods which optimise and complement those motion compensations. These methods have been developed between 2003 and 2008 and were also presented as Langner et al. [31–34].

In parallel to our work, several other groups are also working on the optimisation of motion compensation techniques for PET imaging. In recent publications, e.g., by Thielemans et al. [35], Lamare et al. [36] and Rahmim et al. [37, 38] motion compensation methods have been proposed to be potentially integrated within the image reconstruction process. In addition, other research groups such as Tellmann et al. [14] and Woo et al. [39] are working on optimisations of motion compensation methods that make use of a compensation in the sinogram space mainly. In contrast to these experimental event-driven motion compensation techniques, the methods described in the present thesis mainly concentrate on the development and enhancement of a motion compensation for clinical PET.

1 Fundamentals

Learn the fundamentals of the game and stick to them. Band-Aid remedies never last.

(Jack Nicklaus)

1.1 Positron Emission Tomography

Based on the discovery of the positron in 1932 by C. D. Anderson [40], *positron emission tomography* (PET) is in principle an imaging technique to detect and quantify the spatial distribution of positrons. This is realised by the registration of positron-electron annihilations within a monitored area. However, only further developments have enabled PET to become the established medical imaging method that it is today. The findings of O. Warburg in 1924 [41, 42], showing that a cancer cell has an increased glucose metabolism due to its high energy demands, and the first experiments with radioactive labelled fluoride in 1940 [43] have to be mentioned in this context. Nevertheless, it took almost 40 years for the first patient to be routinely examined with PET. The mathematical basics for filtered-backprojection (FBP) image reconstruction in 1976 [44], but especially the successful synthesis of 2- ^{18}F fluoro-2-deoxy-D-glucose (FDG) in 1977 [45] were the main reasons for the broad introduction of PET into medicine.

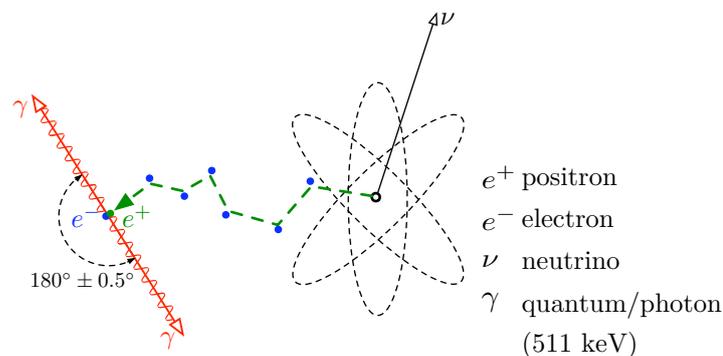


Figure 1.1:

β^+ decay with the flight path of an emitted positron (e^+) and neutrino (ν): After the positron loses enough energy by interaction with electrons (e^-), it will annihilate with a single electron. Due to the antimatter reaction, two γ -quanta with energies of 511 keV are generated and emitted diametrically opposed to each other.

Since then, PET has become an indispensable method for non-invasive imaging and quantification of metabolic processes of living organisms. A radioactively labelled substance (*tracer*), whose concentration is within a range of a few picomolar or nanomolar, is brought into the metabolism of the examined organism, usually by injection. Compared to other tomography imaging techniques like *computed tomography* (CT) or *magnetic resonance imaging* (MRI), where a con-

trast medium within a macroscopic range is applied, PET has the advantage that the metabolism of the organism itself is not affected by the applied substance. Due to the accumulation of the tracer within a target region, a spatial image is calculated (*reconstructed*) from the registered γ -quanta that are emitted by the positron-electron annihilation (cf. figure 1.1). According to the intrinsic tracer accumulation over time, conclusions on the physiology and pathophysiology within the target region can be drawn. These factors make PET a *functional imaging* method and represent a clear advantage over other medical imaging techniques.

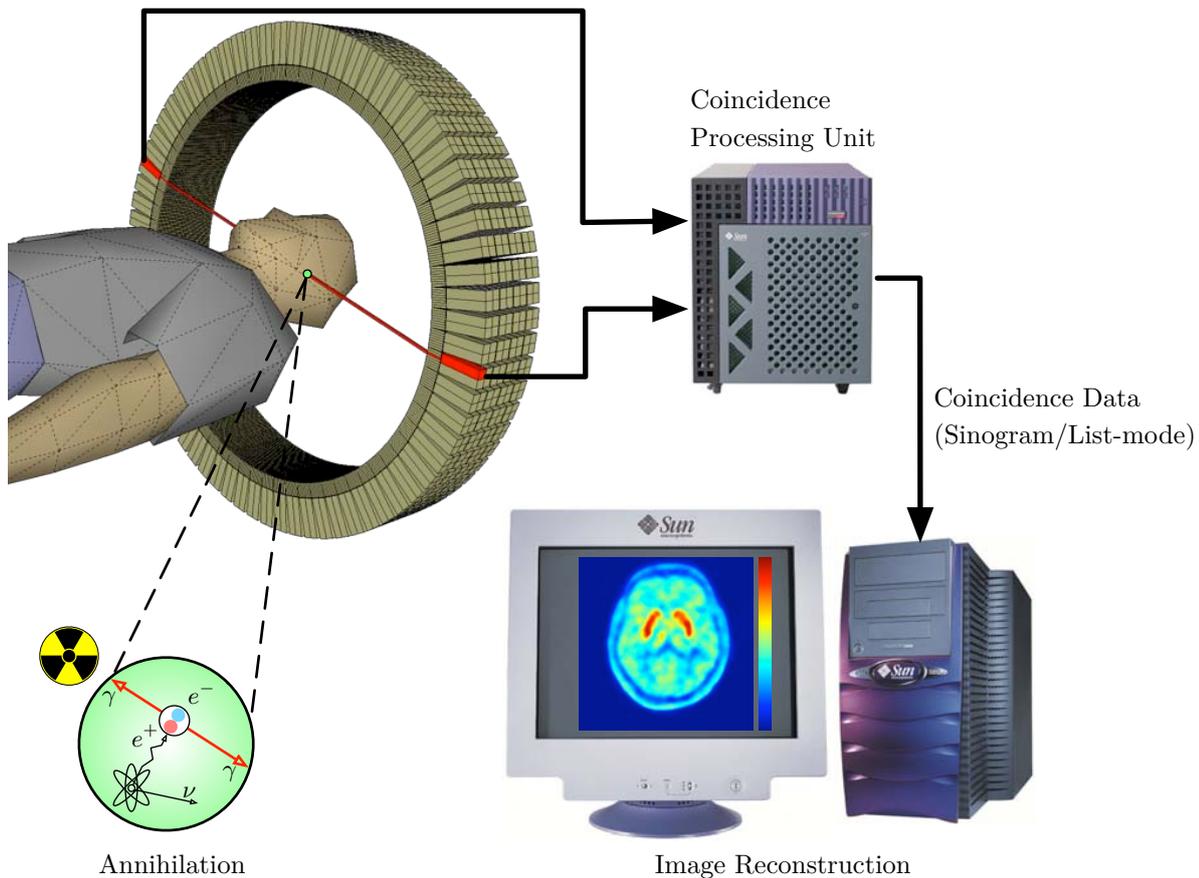


Figure 1.2:

*A flow-chart of the data registration process in PET: After injection of the radioactive tracer substance, the emitted γ -quanta are registered via scintillator crystals. Coincidence electronics filter these events and define which of the so-called *singles* are to be considered part of an actual coincidence event. After forwarding the filtered coincidences to an acquisition workstation, the final image is *reconstructed* [31].*

Since its first application, PET has undergone a constant change regarding its main focus of medical examinations. Whilst during its first years, the focus has been on neurological and cardiological examinations only, most of the PET examinations today focus on oncology. In the new field of *radiotherapy planning* (RTP), for example, the metabolic information retrieved from PET is becoming more and more important [46–48]. In addition today, PET is commonly used for the diagnosis and monitoring of diseases that directly affect the metabolism of the brain (e.g. *Parkinson's disease*, *Alzheimer's disease*, etc.). This topic has become quite important these days, especially due to the constantly extending life expectancy of humans and because of

the demographic changes in Western industrialised nations with an increasing number of elderly people.

1.1.1 Data Registration Process

To register emitted γ -quanta produced during positron-electron annihilations, scintillator crystals are commonly used in physics. These crystals have the physical attribute of generating secondary photons of a different energy range and wavelength as soon as the energy of a striking γ -quantum is absorbed. To detect the generated light, the crystals are connected to *photomultiplier tubes* (PMT) [49], *avalanche photodiodes* (APD) [50, 51] or just recently even to *silicon photomultipliers* (SiPM) [52].

In case of a PMT (which most conventional PET scanners use) the secondary photons hit a photo-cathode wherein electrons are then emitted (cf. figure 1.3). Next, the generated signal is amplified via multiple dynode transpositions in the PMT, resulting in a measurable signal reaching the output connectors. These electrical signals are then registered by electronics, which decide if a signal is considered a valid registration event (*singles*) or not.

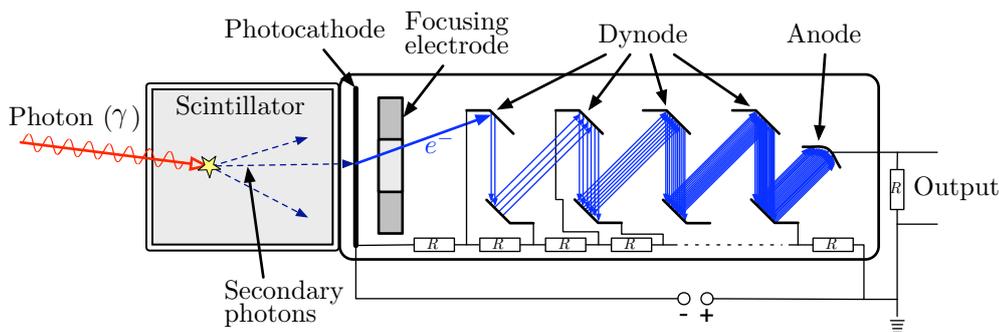


Figure 1.3:

A photomultiplier tube (PMT) connected to scintillator crystals as usually used in a PET scanner: The emitted γ -quantum is absorbed by the scintillator and transformed into secondary photons of a different energy range. These photons hit a photocathode where electrons (e^-) are emitted and subsequently amplified via multiple dynodes such that a signal at the output connectors of the PMT is measurable.

By positioning two γ -detector devices in diametrically opposed locations, both emitted γ -quanta of a single positron-electron annihilation can be registered simultaneously. For obtaining the spatial location where the quantum was transformed into secondary photons, however, multiple PMTs are usually grouped together and connected to an array of scintillator crystals (*detector block*) [53, 54]. Using the differences in intensity of the electrical signal on each of the PMTs in such a combined scintillator array, the area wherein a γ -quantum hits the scintillator is identified (cf. figure 1.4). If several detector blocks are put in a circular arrangement (a ring of detectors), annihilation events can be caught in one complete plane. Furthermore, by stacking multiple *detector rings* in axial direction, a cylindrical arrangement of detectors is usually achieved. This allows a complete spatial detection of events within the *field-of-view* (FOV) of the arranged detector blocks.

For detecting which of the registered events (*singles*) belong to the same positron-electron annihilation, electronics apply a limiting time window of several nanoseconds to the registration

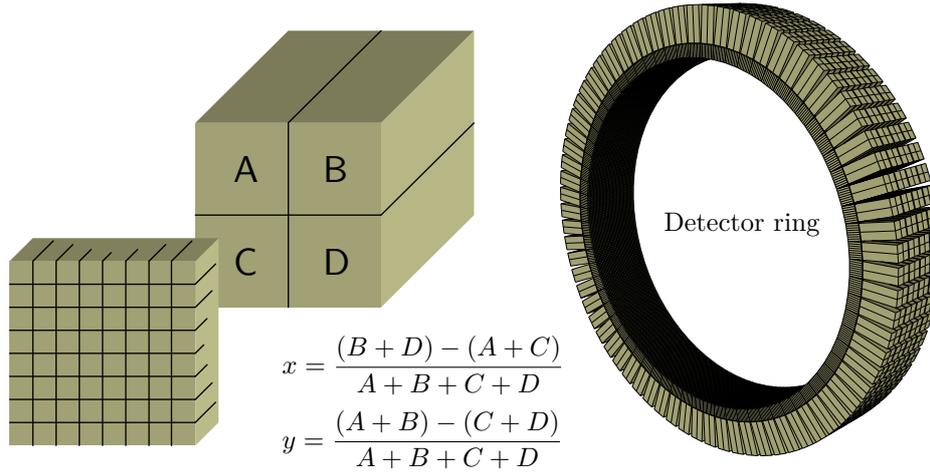


Figure 1.4:

*Scintillator crystals in PET are grouped together as a single detector block by combining an even number of crystals: To calculate the crystal coordinate (x, y) where a γ -quanta was transformed into a secondary photon, the signals at the PMTs are put in relation to each other [55]. Furthermore, several hundred detector blocks are put together in a circular orientation to each other to form the typical *detector ring* used in PET.*

process. If two events of an allowed detector combination are registered during the limiting time frame, they are considered to be of the same annihilation (a so-called *coincidence*). This allows to identify the two detectors that registered the γ -quanta. However, it usually does not allow to identify where between those two detectors the annihilation has taken place. A standard PET system therefore registers a coincidence event with information about the two involved detectors only. Thus, a coincidence is also called a *line-of-response* (LOR), because the actual point where the annihilation took place has to be on the imaginary line that can be drawn between the two opposed detectors.

Limitations of the coincidence detection process

In an ideal situation, a spherical arrangement of γ -detectors around the examined object would allow to register all of the annihilations resulting from the β^+ -decay of the applied tracer substance. In practice, however, the circular arrangement of the detectors in PET is only one reason why not every annihilation event is registered during an examination. The technical limitations of the PET system itself, such as the dead time of electronic components and the PMTs, are also reasons for the *loss* of a certain amount of events. Moreover, there are other physical attributes of the β^+ -decay and its registration in PET, which have an influence on the efficiency of PET in general. In the following paragraphs, these effects and common compensation methods are explained.

Sensitivity differences and dead time: Because of differences in structure and setup of scintillator crystals, and due to differences in the electronics of the scanner, the registration sensitivity varies between individual γ -detectors. In addition, the PMT electronics also imply a certain dead time when registering a γ -quantum. To compensate for these sensitivity gaps in the data registration process, a *normalisation* between detectors is required for

a uniform statistical distribution. This sensitivity normalisation is achieved by performing a separate data acquisition with a radioactive source using a long half-life isotope. After the normalisation acquisition, a coefficient $\eta_{(i,j)}$ is calculated for each detector tuple (i, j) by

$$\eta_{(i,j)} = \eta_i \cdot \eta_j \cdot \phi \quad (1.1)$$

where ϕ is a geometric factor describing the influence of the angle of incidence of a γ -quantum on the detection efficiency [56]. During image reconstruction, this normalisation factor is used to compensate for the different sensitivities of crystals, and for dead time differences of the involved components.

Attenuation: On their way from the point of annihilation through the FOV, the emitted γ -quanta lose energy until they reach a γ -detector. A γ -quantum may even be absorbed completely by the surrounding matter so that it does not reach a scintillator crystal at all. The magnitude of attenuation, however, varies with the density of matter a γ -quantum has to pass. For the data acquisition in PET, this has the effect that for anatomical areas with high density (e.g. bones), the calculated images end up with underestimated areas in the image and thus with inaccurate values in the volume data. This is problematical, especially for a precise quantification of the tracer concentration within those areas. For a compensation of the influence of attenuation on the PET data, the attenuation is commonly measured with a *transmission* scan. Coefficients on the magnitude of the attenuation are calculated either by using a long-living radioactive source (e.g. ^{68}Ge) rotating around the respective object, or by deriving the attenuation from a CT or MRI image in case of multi-modality systems. Due to the measurement of the spatial attenuation, a three-dimensional *attenuation map* is calculated and used during the image reconstruction to compensate for the influences of the attenuation.

Positron range: As previously discussed, a positron interacts with several electrons until it loses enough energy to annihilate with a single electron. During that process, each interaction with an electron causes the positron to slightly change its flight path (cf. figure 1.1). In practice, this leads to a misinterpretation of the real spatial origin of the positron, which e.g., for ^{18}F is about $\approx 0.2\text{ mm}$ within water [57]. In addition, the two γ -quanta are not emitted exactly 180° opposed to each other. Because a small amount of energy is still left at the time of the positron-electron annihilation and because of the principle of energy conservation, the two γ -quanta are emitted at an angle of 180° with an actual deviation of $\approx \pm 0.5^\circ$. Both effects are unavoidable and mark the physical limitation in terms of expectable spatial resolution of a PET scanner.

Radioactive decay: While the radioactive decay itself is not a limitation per se, it is a factor that needs to be corrected to produce correct results during a quantitative analysis of PET data. During the data acquisition, the radioactive isotope decays exponentially over time according to its half-life. To calculate the amount of radioactivity A at a specific time t , the equation

$$A(t) = A_0 e^{-\lambda t} \text{ with } \lambda = \frac{\ln(2)}{T_{\frac{1}{2}}}, \quad (1.2)$$

is used where A_0 is the applied/injected start activity and $T_{\frac{1}{2}}$ the specific half-life of the isotope. During image reconstruction, a correction factor f is calculated, which corresponds to

$$f = \frac{N_c}{N_m} = \frac{A_0}{\frac{1}{t_\Delta} \int_{t_0}^{t_0+t_\Delta} A_0 e^{-\lambda t} dt} = \frac{t_\Delta \lambda e^{\lambda t_0}}{1 - e^{-\lambda t_\Delta}}, \quad (1.3)$$

where N_c is the number of corrected counts, N_m the number of the measured counts, t_Δ the acquisition duration and t_0 the start time of the acquisition. The resulting correction factor is used to normalise the registered counts in the resulting sinogram during the image reconstruction phase.

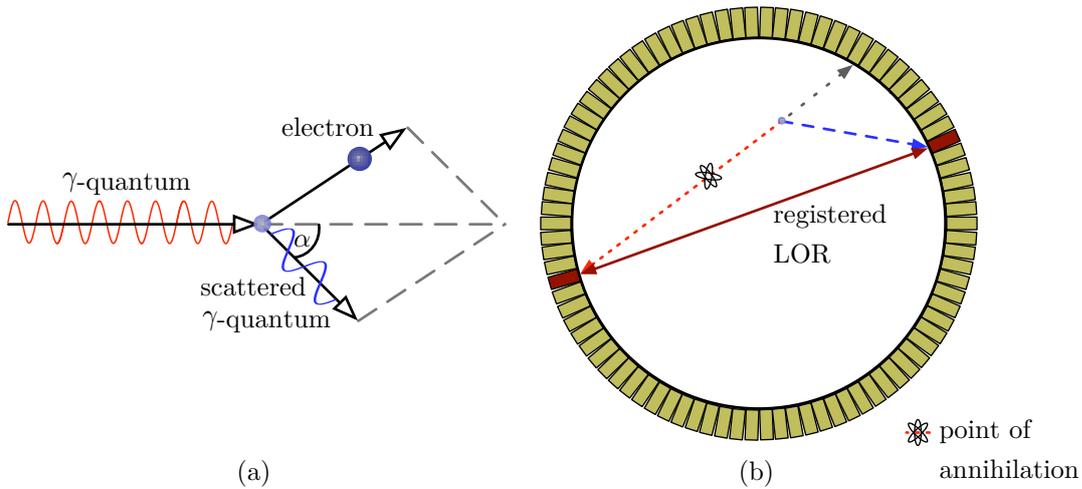


Figure 1.5:

Scattering due to Compton effect: Due to Compton scattering (a), a γ -quantum changes its direction and loses energy when interacting with an electron. For PET, this causes a registration of incorrect LORs (i.e. scattered events) as shown in (b).

Scattering: A γ -quantum, like any other photon, interacts with the electrons of the surrounding matter. However, compared to photons of a different energy, an interaction with a γ -quantum of 511 keV results in *Compton scattering*. Because of this type of scattering, a γ -quantum not only changes its direction but also loses energy. This has the effect that either the γ -quantum may be scattered so that it leaves the FOV without being registered, or that it arrives at a detector in a non-diametrically direction opposite to the other quantum emitted by the annihilation (cf. figure 1.5). Such incorrectly registered coincidences are commonly called *scattered events*. As scattered γ -quanta lose some of their initial energy, one way of reducing the registration of scattered events is by increasing the lower energy (LE) at which a γ -quantum is still being accepted by the electronics. In practice, increasing the LE from 250 keV to 350 keV reduces the fraction of scattered events from 16 % to 13 % [58]. However, as the LE cannot be further increased without dropping too many true events, several different methods exist to further reduce the amount of scattered events [59]. One common way, for example, is an attenuation-based calculation of a spatial scatter distribution map. According to this scatter map, the image reconstruction

calculates a scanner dependent scatter model, which diminishes the amount of registered scatter events due to a complex scatter simulation [60, 61].

Random coincidences: Two gamma rays originating from two unrelated positron-electron annihilations may be registered as a random coincidence (*randoms*) in case the gamma rays are simultaneously detected within the coincidence resolving time of the scanner (cf. figure 1.6). While the electronics cannot directly identify such random events, the probability of random events can be statistically determined by

$$N_{rand} = 2\tau \cdot N_i \cdot N_j \quad (1.4)$$

where N_{rand} is the random rate, N_i and N_j are the singles rates of the two involved detectors i and j , and 2τ the coincidence resolving time of the scanner [62]. As random coincidences provide a generally non-uniform background to the reconstructed image, a way of compensating for the effects of random events is essential. This is especially important since the amount of true coincidences is usually only a very small fraction of the singles rates, e.g. about 0.1% for a brain scan. Therefore, methods have been developed to compensate for the influences of random coincidences. One method often used is measuring the rates of the *singles* registered at each γ -detector unit. According to this information, the randoms are calculated by equation 1.4 and subtracted accordingly. Another common method for estimating the randoms is using a second, delayed coincidence window. This allows to measure the accidental coincidences and to subtract them directly from the registered coincidences [63].

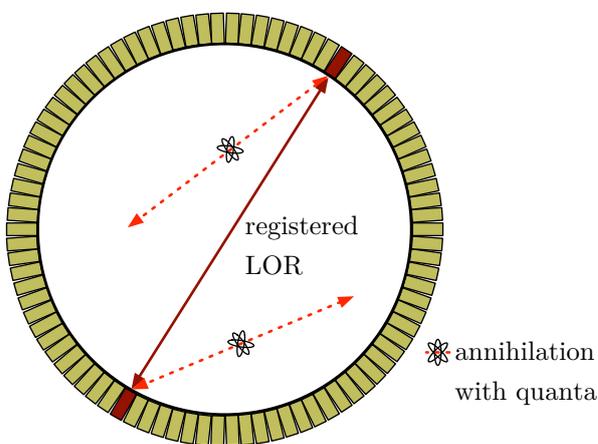


Figure 1.6:

Random coincidences: Two γ -quanta originating from two unrelated annihilations may be registered as a *random coincidence* if the two single quantum hit a detector within the coincidence resolving time.

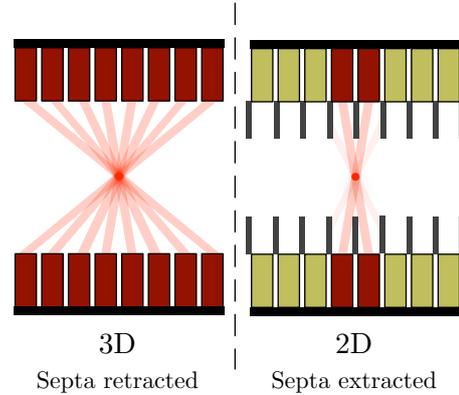
Sensitivity of the sampling process

For the imaging process in PET, a certain amount of coincidences is required to allow the image reconstruction to calculate appropriate images. Moreover, a high amount of registered coincidences is generally desirable because of the way the image reconstruction process works. In practice, however, the amount of applied radioactivity is mostly limited by radiation protection regulations and technical implications of the PET scanner. This, and the commonly high amount of events lost due to accidental coincidences (e.g. scatter, randoms), is the reason for only a limited amount of true events being available for the image reconstruction. To compensate for

this statistical limitation, various technical adaptations of the data registration process have been introduced in the past. They either try to increase the count rate, or to reduce the amount of accidental coincidences by technical means.

Figure 1.7:

Illustration of 2D and 3D mode processing: In 2D mode, septa are put between each crystal of a detector. These physical barriers limit the amount of possible axial detector combinations. As a result, the amount of accidental coincidences (scatter, randoms) is reduced – at the cost of a generally lower statistics outcome.

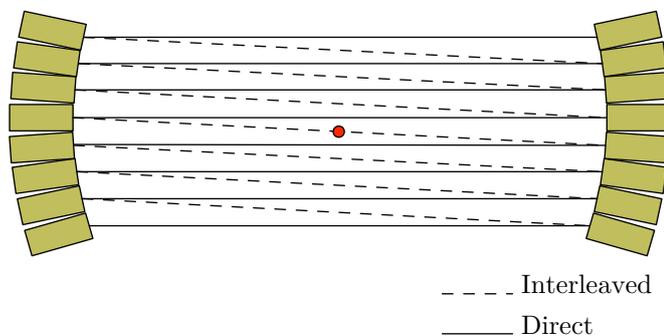


Traditionally, a PET scanner is running in *2D mode*. In this mode, the scanner accepts coincidences only from a limited number of axial ring combinations (*span*). This is done by masking out non-allowed coincidences during the image reconstruction, but also via physical barriers in the scanner itself. The so-called *septa* rings are made of 0.8 mm thick tungsten and are located between each of the detectors in axial direction. Due to the physical characteristics of tungsten and the axial location of the septa, gamma rays entering from flat angles are absorbed as illustrated in figure 1.7. While this generally reduces the rate of scatter from 30% – 40% to 10% – 15%, [59] – the initial reason for the development of this method – it also reduces the amount of true events, i.e. the overall count rate.

Therefore, modern PET scanners allow to run an acquisition in *3D mode* as well. During this mode, the septa are removed, allowing for a utilisation of a wider ring combination. This mode has become available and well accepted mainly due to development of scatter correction algorithms, which compensate the additional amount of scatter in 3D as mentioned earlier. As a result, the simulated scatter distribution is calculated during the image reconstruction so that a compensation for the higher fraction of scatter in 3D mode is automatically performed.

Figure 1.8:

Interleaving technique commonly used in PET: Here a point source (red circle) is shown with a cut-out of the possible direct and interleaved LOR combinations of one projection angle. This interleaving is repeated for every subsequent projection angle so that due to the comprised information of the interleaved combinations, the number of detector combination is doubled and thus the sensitivity is increased.



Another technique to increase the sensitivity of a PET scanner is *interleaving* [64]. By assuming, that coincidences coming from a point source, which is located exactly between neighbouring detectors, LORs between the cross detectors are normally not identified. Interleaving describes

a method to increase the available angular detector combinations by introducing an additional cross combination LOR. These *interleaved LOR* combinations are illustrated in figure 1.8. Due to this technique, the resolution, especially near the centre of the FOV is increased because of the additional information in one projection angle.

Data storage and data formats

One of the most commonly used methods to store the raw coincidence information in PET is the storage of the coincidence data into large three-dimensional data arrays. Such a data array is, in simplified form, a set of histograms where each entry (*bin*) of a histogram corresponds to a specific possible LOR. Thus, for each coincidence event detected, the bin corresponding to the LOR is incremented by one. In practice, a histogram bin corresponds to a line-integral through the original object; the collection of all line-integrals parallel to a given angular direction is called a *projection*. Figure 1.9 illustrates this for a 2D object.

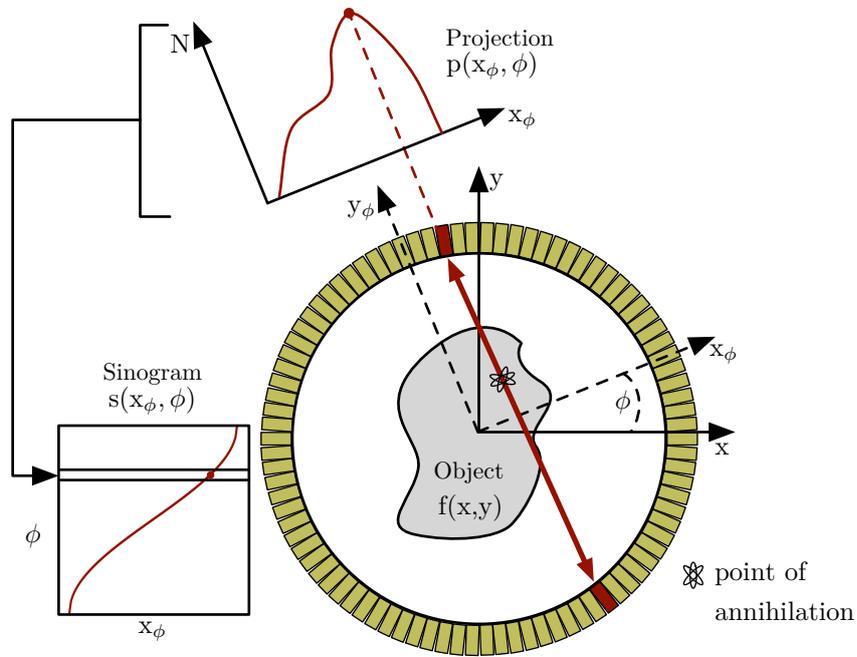


Figure 1.9:

One-dimensional projection of a 2D object and its storage: The relation between a one-dimensional projection $p(x_\phi, \phi)$ of a 2D object as the integral of all parallel LORs for an angle ϕ relative to the scanner ring and its storage in a histogram-like *sinogram* $s(x_\phi, \phi)$ is illustrated.

Such a projection, $p(x_\phi, \phi)$, is described by the integral,

$$p(x_\phi, \phi) = \int_{-\infty}^{\infty} f(x, y) dy_\phi, \quad (1.5)$$

where (x_ϕ, y_ϕ) are the rotation offsets of (x, y) with an azimuthal angle ϕ relative to the scanner axis. As shown in figure 1.9, these two-dimensional data are generally stored in a histogram indexed by x_ϕ and ϕ . If stored in such a histogram, the projection data $p(x_\phi, \phi)$ are commonly referred to as a *sinogram* with indices (x_ϕ, ϕ) . This term was introduced by Edholm et al. [65]

in 1978 because the sets of LORs containing a fixed point (x_0, y_0) are located along a sinusoid in the (x_ϕ, ϕ) plane of the histogram.

For a 3D object, $f(x, y, z)$, equation 1.5 can be generalised by considering a projection to be a two-dimensional set of parallel line-integrals, not only with a given azimuthal angle ϕ , but also with a polar angle θ . The resulting 2D projection, $p(x_\phi, y_\phi, \phi, \theta)$, as generally present in PET can be expressed by

$$p(x_\phi, y_\phi, \phi, \theta) = \int_{-\infty}^{\infty} f(x, y, z) dz_\theta, \quad (1.6)$$

where the coordinates for the 3D rotations are defined by

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} -\sin \phi & -\cos \phi \sin \theta & \cos \phi \cos \theta \\ \cos \phi & -\sin \phi \sin \theta & \sin \phi \cos \theta \\ 0 & \cos \phi & \sin \theta \end{bmatrix} \cdot \begin{bmatrix} x_\phi \\ y_\phi \\ z_\theta \end{bmatrix} \quad (1.7)$$

In case $\theta = 0$, this is reduced to the 2D case as shown earlier in equation 1.5. The relationship between such two-dimensional projections and their corresponding locations in a sinogram with a fixed polar angle θ is illustrated in figure 1.10.

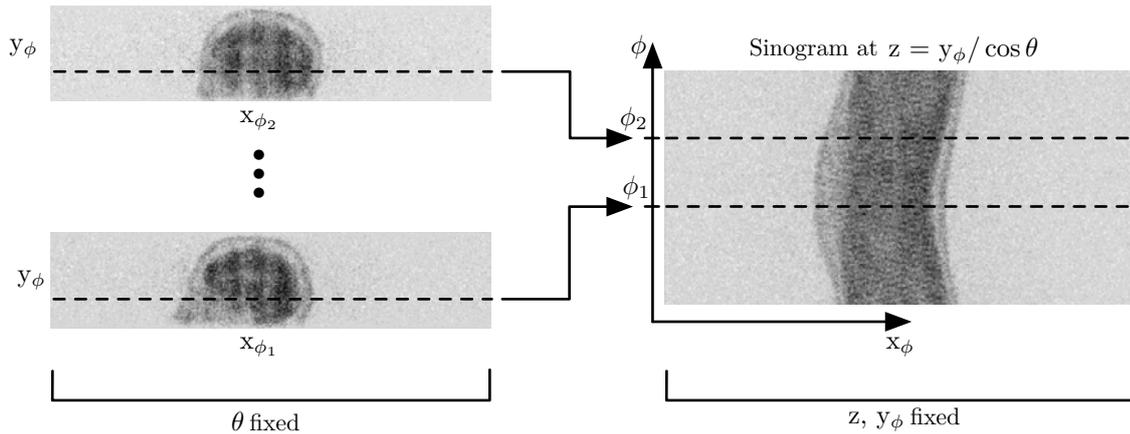


Figure 1.10:

The correlation between two-dimensional projections and oblique sinograms illustrated on real data of a PET brain study: For a fixed polar angle θ , all projection data x_ϕ at a fixed y_ϕ (left-hand side) is summed up into a single sinogram at location z (right-hand side).

In addition to the storage in the compressed sinogram format, undersampling is often applied in order to further reduce the storage and thus the computing requirements. In case the average number of detected coincidences per sinogram sample is small, angular undersampling ($\Delta\phi > 1$) allows for a relevant reduction of required space and computing time while losing only a minimum amount of relevant data. This process is also called *transaxial mashing*, with a mashing factor m defined by

$$m = \frac{\Delta\phi \cdot N_d}{2\pi}, m \in \mathbb{N}_{>0}, \quad (1.8)$$

where N_d is the amount of available detectors. This factor is usually an integer so that undersampling amounts to summing up groups of m consecutive rows in the sinogram. While transaxial mashing produces relatively stable results – especially for studies with a small reconstructed FOV (R_{FOV}) – radial undersampling ($\Delta x_\phi > 1$) tends to generate more severe artefacts.

A rule of thumb to match the radial and angular sampling can be derived by using Shannon's sampling theorem via the relation $\Delta\phi \simeq \frac{\Delta x_\phi}{R_{FOV}}$ [66].

Another data format for storing coincidence data is the collection of the LOR information in a list-type format. Within the *list-mode* format [67, 68], each entry in the list contains information about an individual coincidence event (cf. figure 1.11). The information collected can be considered the raw sampling data of a PET scanner as it could even contain information about the specific energy levels present on the PMTs when the coincidence event was detected. While the detail of information varies between scanner models, the amount of data is significantly larger than in *histogram-mode*. Whereas the list-mode data require several gigabyte of space, the compressed sinogram data require a few hundred megabytes only [69]. For traditional PET imaging, the implied loss of information by the sinogram compression is often negligible as the common characteristics are typically chosen to be the pair of detectors involved in a coincidence event. Other information, such as the amount of energy of each singles and the order of occurrence are normally not of interest. Therefore, the additional memory space required for list-mode was the main reason for PET to generally store its data directly in sinograms only, especially in the early days of PET.

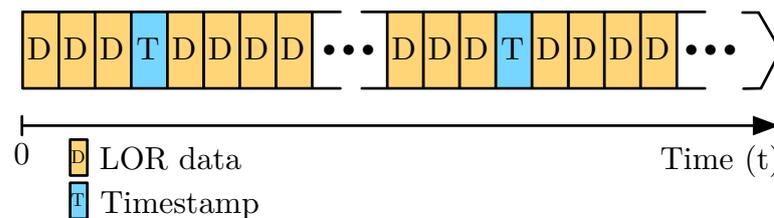


Figure 1.11:

Schematic description of the list-type data format (list-mode) which is available on modern PET scanners: This data format contains the raw coincidence information as well as timing information. For some scanner types the list-mode format even contains the registered energy levels at both PMTs of the involved detector blocks.

Today, however, and thanks to the improved computing capabilities, list-mode is becoming more and more attractive. The additional information preserved in list-mode allows for new data processing procedures that may improve the imaging process of PET. Such procedures are e.g. improved image reconstruction methods such as a list-mode-based image reconstruction [36, 70, 71]. In addition, methods which generally need the coincidence information of a single LOR, require a PET scan to be performed in list-mode. One of these methods is the *event-driven motion compensation* which will be discussed in section 1.2.3.

Another important advantage of list-mode acquisition is the possibility of changing the coincidence data sorting scheme even after the data acquisition is completed. This is particularly attractive for dynamic PET acquisitions, where the coincidence data are commonly split into different time slices (*frames*). Distinct frames are often used to perform a quantitative analysis on the tracer kinetics. By comparing the tracer distribution of different time frames, the dynamics of the accumulation of the tracer in the analysed areas can be visualised via *time-activity-curves* (TACs) and thus further analysed. This allows to draw conclusions about the physiology of the examined organism. In histogram-mode, this framing scheme has to be defined before the scan takes place, whereas list-mode processing allows to put the raw data in different frames even

after acquisition. This enables the physician to apply different analyses to the data without having to perform additional scans.

In addition, list-mode data are also used to separate coincidence data based on certain other criteria. It can, for example, be split into *gates*, where one gate contains only coincidence data which were collected during a time where the analysed object was in a specific periodic state. This method is commonly used throughout different imaging modalities (MRI, CT, etc.) to compensate for the influence of certain types of periodic motions, such as respiratory and cardiac motion and will be further discussed in section 1.2.3.

1.1.2 Image Reconstruction

One of the main steps in the imaging process of PET is the *image reconstruction* phase where the final volume image is calculated from the coincidence data sampled during acquisition. This process is the inverse of the line integral summation process in section 1.1.1 where the tracer distribution $f(x, y, z)$ is in relation to the 2D parallel projections $p(x_\phi, y_\phi, \phi, \theta)$ as defined in equation 1.6. To calculate the spatial tracer distribution from the line integral information gained in PET, two types of reconstruction algorithms are available, namely analytic and iterative reconstruction methods. While analytic methods are ‘one-step’ solutions, that is, the image estimate is directly calculated from the coincidence data, iterative methods require a subsequent calculation of a set of image estimations. These are based on a statistical model of the data acquisition process and derive subsequent estimations until an image estimation matches a certain quality criterion.

The basis for all analytic reconstruction methods is the projection-slice theorem, better known as the *central-section theorem*. This theorem states that for a three-dimensional object the 2D Fourier transform (FT) of a two-dimensional line integral projection profile is equal to a central section (values on a line through the centre) of the 3D Fourier transform of the very object [59, 72]. This relation is expressed by

$$P(v_{x\phi}, v_{y\phi}, \phi, \theta) = F(v_x, v_y, v_z), \text{ with } v_z\theta = 0, \quad (1.9)$$

where $F(v_x, v_y, v_z)$ is the 3D Fourier transform of an object $f(x, y, z)$ and $P(v_{x\phi}, v_{y\phi}, \phi, \theta)$ the 2D Fourier transform of the two-dimensional line integral at a fixed azimuthal angle ϕ . Figure 1.12 illustrates this relationship on a spheric object.

This means that each two-dimensional parallel ray projection $p(\cdot)$ contains the data for one line through the 3D Fourier transform of the object. By taking the 2D Fourier transform of the views at different azimuthal angles ϕ around the object, one can obtain the values for the complete 3D Fourier transform of the object. Thus, a simple analytical reconstruction algorithm would then be to regrid these values onto a three-dimensional array and perform an inverse 3D Fourier transformation. One of the most common analytical reconstruction algorithms in PET based on that process is the *filtered-backprojection* (FBP) algorithm [44].

In contrast to analytic reconstruction, iterative methods progressively refine estimates of $f(x, y, z)$, rather than by direct inversion of the image transform (cf. equation 1.9). While these methods imply a higher computational cost, they allow for a more accurate modelling

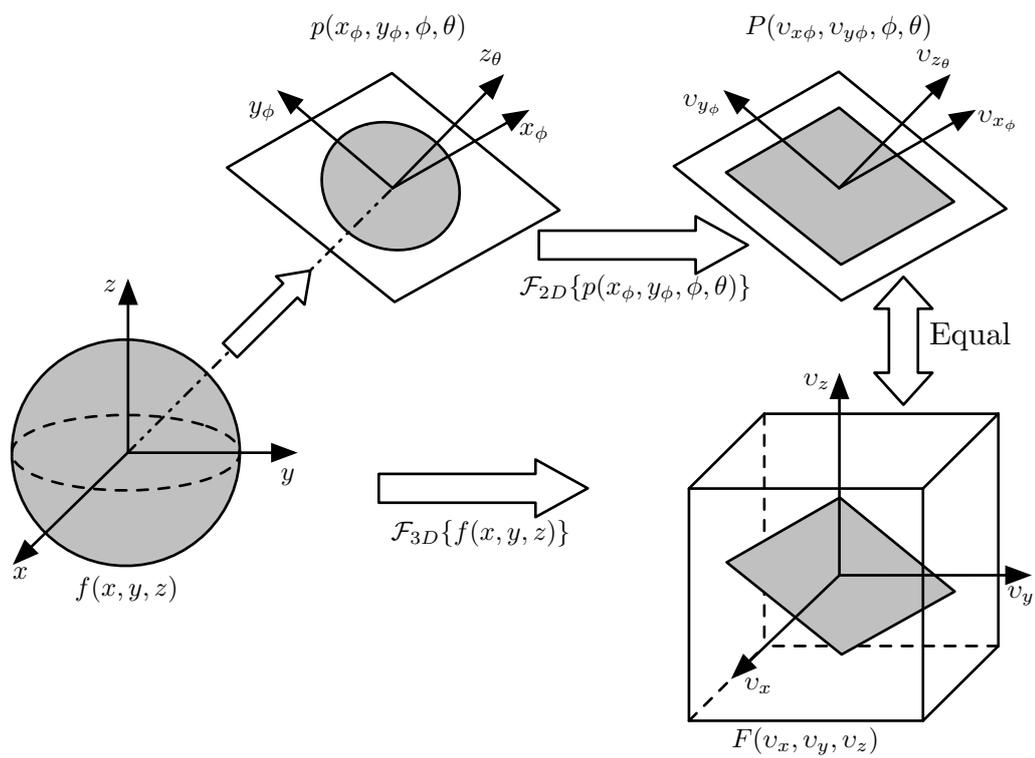


Figure 1.12:

Illustration of the 3D central-section theorem as the basis of analytic image reconstruction used in PET: It shows that the 3D Fourier transform $F(v_x, v_y, v_z)$ of an object $f(x, y, z)$ is equal to the 2D Fourier transform $P(v_{x\phi}, v_{y\phi}, \phi, \theta)$ of the line integral projection at a given azimuthal angle ϕ [59].

of the data acquisition process. The possibility of incorporating *a priori* information into the image reconstruction algorithms allows for a more acceptable solution. Therefore, iterative reconstruction methods prove to produce more accurate results, especially for acquisitions with a low amount of registered coincidences or in cases of a low *signal-to-noise ratio* (SNR). One such *a priori* information is the non-negativity nature of the tracer concentration and therefore the non-negativity when calculating image estimates of $f(x, y, z)$. In contrast, analytic algorithms, such as FBP, cannot take non-negativity into account so that voxels in images reconstructed with FBP may even carry negative values.

The *Maximum-Likelihood by Expectation-Maximization* (ML-EM) algorithm as well as its accelerated variation, the *Ordered Subset Expectation-Maximization* (OSEM) algorithm [73–75], are the most well-known iterative reconstruction methods in PET. Both algorithms seek to maximise the logarithm of the *Poisson likelihood objective function*, as defined by

$$L(\vec{f}, \vec{p}) = \sum_{j=1}^m \left\{ p_j \cdot \log \left(\sum_{i=1}^n A_{j,i} f_i \right) - \sum_{i=1}^n A_{j,i} f_i \right\}, \quad (1.10)$$

where \vec{f} is a vector with all voxels values $\{i = 1, \dots, n\}$ in the three-dimensional image; \vec{p} are the values of the projections $\{j = 1, \dots, m\}$ where $m = N_{x_\phi} \times N_\phi$ and $A_{j,i}$ is the $m \times n$ system matrix. Following this reasoning, the system matrix holds the probabilities of a γ -quantum emitted at voxel i which is being detected in a projection bin j . In addition to the maximisation of the logarithm of the likelihood objective function, the ML-EM algorithm updates voxel values by using the EM algorithm

$$\vec{f}_i^{n+1} = f_i^n \frac{1}{\sum_{j'=1}^m A_{j',i}} \sum_{j=1}^m \left\{ \frac{p_j}{\sum_{i'=1}^n A_{j,i'} f_{i'}^n} \right\} \mid i = 1, \dots, n, \quad (1.11)$$

where \vec{f}^n is the n -th estimate of the image. This iterative algorithm has some remarkable properties: it implies non-negativity of the subsequent image estimates \vec{f}^n and therefore includes the general *a priori* information of non-negativity of the tracer distribution. In addition, the image iterations converge for $n \rightarrow \infty$ to an image that maximises the loglikelihood and therefore result in an image that ‘best’ fits the supplied data. Here, ‘best’ is understood in the sense of the *Poisson likelihood objective function*. This means that the ML-EM algorithm, as a regular iterative process, needs an abort criterion because, depending on the input data, subsequent image estimates introduce instabilities that may lead to image artefacts when the number of iterations exceeds a given threshold [76]. Even if there have been suggested methods to automatically specify an abort criterion [77], in practice, the number of iterations n_{max} are determined empirically. In addition, the use of post-reconstruction filters, e.g. a 3D Gaussian filter, are used to ‘smooth out’ certain inhomogeneities caused by the image estimation nature of the algorithm.

While the plain ML-EM algorithm takes a considerable amount of time to converge to an acceptable image estimate, the OSEM algorithm proves to be more practical for routine use. It is based on the ML-EM algorithm and simplifies the EM equation 1.11 considerably. By separating the iterative calculations in disjoint subsets with a fixed amount of LORs, the required time to

reconstruct an acceptable image is reduced. In addition to that, each subset is processed in a well-defined order. In this order, a certain subset is usually used at a specific iteration step. If applied to equation 1.11, an OSEM iteration is described by

$$\vec{f}_i^{n+1} = f_i^n \frac{1}{\sum_{j \in J_n \bmod s} A_{j,i}} \sum_{j \in J_n \bmod s} \left\{ \frac{p_j}{\sum_{i'=1}^n A_{j,i'} f_{i'}^n} \right\} \mid i = 1, \dots, n, \quad (1.12)$$

where J_1, \dots, J_S are the disjoint subsets S used in the calculation. In contrast to the ML-EM algorithm, the OSEM algorithm, as a rule of thumb, tends to converge about S times faster. A negative effect of the split computation is, however, that the asymptotic convergence to the maximum-likelihood estimator can no longer be guaranteed. Thus, OSEM tends to cycle between S different image estimates instead of taking a single one until the final estimate. To compensate for these adverse effects, the number of subsets as well as the number of projections per subset are often identified empirically when using OSEM reconstruction [55].

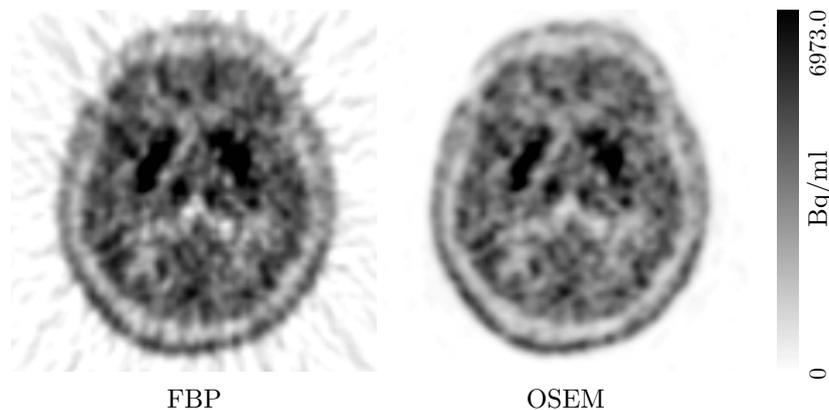


Figure 1.13:

Qualitative comparison between the analytic FBP and the iterative OSEM image reconstruction: Here, a single transaxial plane of a $[^{18}\text{F}]\text{DOPA}$ brain study is shown (3D list-mode, 184 MBq, 90 min). On the left-hand side the FBP reconstructed image is shown. On the right-hand side the same data but reconstructed with OSEM (16 iterations, 6 subsets) is shown. A clear indicator of the better performance of the iterative OSEM reconstruction is the reduced amount of streak artefacts mostly visible in outer areas of the FBP reconstructed image.

Compared to analytic reconstruction methods, especially FBP reconstruction, iterative methods produce generally less ‘streaky’ artefacts in final images as can be seen in figure 1.13. While these artefacts are especially visible along the outer edges of the image, they also have an impact on the inner area and therefore also on the quantification of a *region-of-interest* (ROI). In addition, a generally better SNR in low activity regions is a result of iterative methods which supports e.g. a better visibility of the contours of the examined object.

1.2 Motion Susceptibility

By considering the data sampling process only, a PET scanner works very similar to an optical camera system. Although it does not sample visible photons within the energy range of

1.91 eV (red) to 3.26 eV (violet), some technical parameters of the sampling, such as an *aperture* and *exposure time* can be mapped in the way a PET scanner works. While for an optical camera system the settings of the aperture in relation to the exposure time defines how much light can enter the camera system at a given time, for a PET scanner the term *aperture* may be used for the sampling efficiency of the scanner. To continue this analogy, the term *exposure time* used with an optical camera system can be mapped to the examination time where LORs are summarised in PET, thus corresponding to the term *frame* or *acquisition duration*. Apart from this theoretical mapping of terms, the sampling process of the photons is also partly prone to the same type of negative influences. The most common ones are scattering and absorption/attenuation, as already discussed in section 1.1.1 for the 511 keV γ -photons sampled in PET.

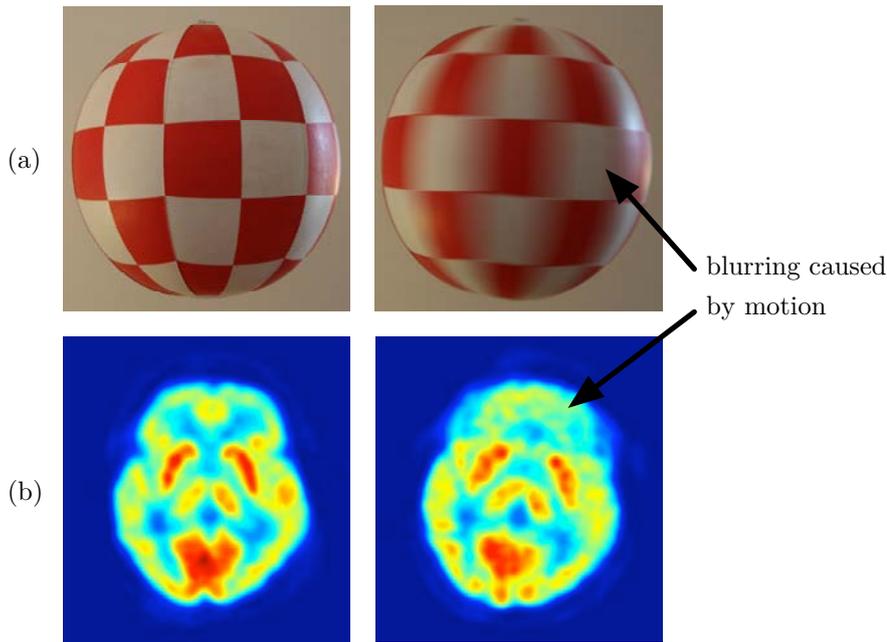


Figure 1.14:

Illustration of similarity of motion influences on different imaging systems: In the top row (a), two images of a standard photography of a ball are shown. While both images have been taken with the same exposure time (0.8s), the right one shows motion artefacts due to the ball rotating during the take. Similar effects can be seen on PET images where two transaxial images of the same brain study are shown in the bottom row (b). On the right-hand side, the blurring due to a rotation around the Z-axis is clearly visible in comparison to the motionless image on the left-hand side.

Another negative influence which both camera systems are susceptible to is the case when the subject which is being examined, moves during data acquisition. As shown in figure 1.14, blurring and image artefacts are the results for both systems in case of such motion. Thus, the object needs to remain still to produce best results. For an optical camera system, the aperture and exposure time or even the existence of a flash light are parameters that can be tuned so that their optimal configuration still allows to produce clear and sharp pictures in case of motion. For a PET system, however, the ‘aperture’ is a technical limitation of the PET scanner in general. In addition, a ‘flash light’ would imply a higher applied radioactivity which is mostly restricted by radiation regulations and examination requirements. Also, the regular ‘exposure time’ of 20 to 90 minutes is considerably longer in PET than in an optical camera. This is due to the

fact that the amount of emitted 511 keV photons in PET is only a fraction of what an optical camera system generally has available for the sampling of visible photons and that, due to a metabolic analysis, a longer acquisition time is often desired. As a net effect, sampling the tracer distribution with PET immediately leads to artefacts in the resulting image if the patient moves during the examination. Aside from image blurring, this also has the effect that a quantitative analysis on such motion-affected data automatically results in an incorrect evaluation of the accumulated tracer.

1.2.1 Motion Types

Depending on the examination, different types of patient motion may or may not have a parallel impact on the imaging process of PET. During a *whole body* scan, for example, the constant chest movement caused by inhalation and exhalation does not only result in the movement of the lung [78]. It also results in a movement of some of the other inner organs such as the liver (i.e. *indirect organ motion*). In addition, the constant cardiac motion and a potential outer body movement of the patient add to the resulting image artefacts caused by motion. In contrast, during neurological examinations only the head movement and the body motion of a patient have an impact on the image data. Table 1.1 summarises common types of motion that may occur during the different types of PET examinations and classifies the motion according to its nature.

Table 1.1:

Motion Type Classification: Different types of motions may occur during a PET acquisition with a negative influence on the resulting image quality. The possible types of motion vary with the type of examination where different areas of a patient are analysed. Depending on the motion type, the patient movement can be classified as *voluntary* or *involuntary* with direct influence on a possible motion compensation method.

Examination Type	Motion Type	Motion Class
Whole Body	Body Motion	Voluntary
	Head Motion	Voluntary
	Indirect Organ Motion	Voluntary/Involuntary
	Respiratory Motion	Involuntary
	Heart Motion	Involuntary
Cardiac	Heart Motion	Involuntary
	Respiratory Motion	Involuntary
	Body Motion	Voluntary
	Indirect Organ Motion	Involuntary
Neurological	Head Motion	Voluntary
	Body Motion	Voluntary

In accordance to the type of motion, patient movement can generally be classified as *voluntary* or *involuntary* motion. Moreover, motion of a periodic nature (e.g. respiratory motion, cardiac motion) can be classified as involuntary motion, whereas motion of a non-periodic nature (e.g. body motion, head motion) can be classified as voluntary motion. This also includes indirect motion of inner organs which mostly occurs due to motion of other areas, e.g. the liver movement caused by respiration.

As mentioned, during a typical whole body examination all different types of motion may occur simultaneously. A typical neurological examination, however, is generally only prone to the outer head motion and possible body motion of the patient. This observation is based on the assumption that the inner organs of a human skull (brain, eyes, etc.) are moving in correspondence to the outer limits; thus a human head can be considered a rigid object where one point on the skull moves equivalently to any other point on the head.

1.2.2 Motion Quantification

In three-dimensional space (\mathbb{R}_3), the position and orientation of a body relative to a reference point (commonly the origin of ordinates) is described by the *six degrees of freedom*. Within the Cartesian coordinate system the position of an object is fully specified by a vector \vec{P} with its three coordinates (P_x, P_y, P_z). In addition, a rotation matrix \hat{R} , defined by three angles (η, θ, ϕ) relative to the respective coordinate axes (x, y, z), is used to express the orientation of the body within the coordinate system. While the vector \vec{P} is relative to the point of origin and only a three-dimensional vector, the rotation matrix consists of three consecutive rotation parameters $R_i(\alpha)$ whereas the application of these rotations is non-commutative and defined by $i : x \rightarrow y \rightarrow z$ with an angle α around the axis i . Furthermore, the rotation is commonly defined to be relative to the centre of mass of the body itself and expressed by

$$\hat{R} = R_x(\eta) \otimes R_y(\theta) \otimes R_z(\phi) , \quad (1.13)$$

which results in the rotation matrix \hat{R} , which is defined by

$$\hat{R} = \begin{pmatrix} \cos \phi \cos \theta & -\sin \phi \cos \theta & \sin \theta \\ \sin \phi \cos \eta + \cos \phi \sin \theta \sin \eta & \cos \phi \cos \eta - \sin \phi \sin \theta \sin \eta & -\cos \theta \sin \eta \\ \sin \phi \sin \eta - \cos \phi \sin \theta \cos \eta & \cos \phi \sin \eta + \sin \phi \sin \theta \cos \eta & \cos \theta \cos \eta \end{pmatrix} \quad (1.14)$$

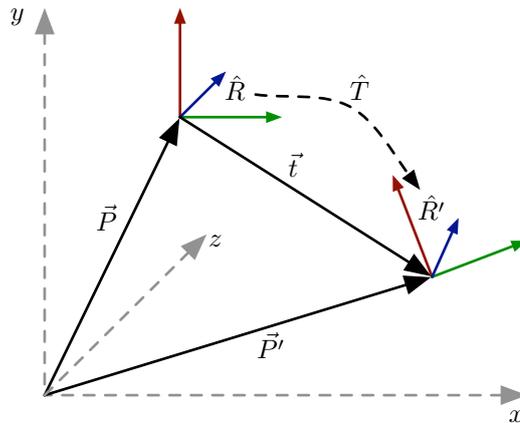


Figure 1.15:

Illustration of the nomenclature of motion quantification: An object (small coloured coordinate axes) is presented by a distance vector \vec{P} and a rotation matrix \hat{R} within \mathbb{R}_3 . In case of movement the new position and orientation is expressed by (\vec{P}', \hat{R}') . This change is also described by a translation vector \vec{t} and a rotation matrix \hat{T} .

If the body moves within the same coordinate system, the new position and orientation is defined by \vec{P}' and \hat{R}' which corresponds to a new position vector and rotation matrix relative to the point of origin (cf. figure 1.15). To retrieve the movement of a point on the object relative to a reference point (e.g. initial position at acquisition start), a translation vector \vec{t} and rotation matrix \hat{T} can be calculated and used to express the patient motion by

$$\vec{x}' = \hat{T} \cdot \vec{x} + \vec{t}, \quad (1.15)$$

which defines that a body at a source location \vec{x} is transformed into \vec{x}' by first rotating it via \hat{T} around the coordinate axes and then translating it via \vec{t} . The order of rotation and translation is non-commutative in this case. Similar to the non-commutative order of the application of the angles of rotation (cf. equation 1.13), a different destination location would otherwise be the result, thus $\vec{x}' \neq \hat{T}(\vec{x} + \vec{t})$.

To quantify the motion of a patient externally, different technical solutions are available and have already been discussed in the literature (cf. table 4.1). The most common methods to capture motion information are shortly outlined in the following section.

Radioactive Markers (Landmarks)

Using external radioactive sources to identify and quantify a movement of the examined object during a PET acquisition is probably one of the oldest and the most intuitive methods [79].

By attaching a radioactively labelled *marker* (e.g. a radioactive sphere) to the examined object, motion can be identified via analysis of the reconstructed image data. The marker's proportions are measured and compared to its known values in order to identify if motion has occurred and if it has had a negative influence on the image data. Furthermore, in case of dynamic acquisition where data are split into several data frames, the motion is quantified by comparing the identified proportions of the marker in one frame image to the proportions found in another (reference) frame.

Video-based Methods

A different way to identify and quantify motion during a PET acquisition is by using a video camera system during the examination.

During PET acquisition, a camera system records a video of the whole FOV of the scanner where optically distinguishable markers are placed on the circumference of the object of interest [17]. Motion can be identified by image analysis of the reconstructed images. Furthermore, the motion can be quantified by deriving the proportions of the used markers from the almost 25 images the camera captures per second. *Optical-flow* algorithms, common in image data processing, are used to extract the motion parameters from the two-dimensional images captured by the system and to transfer them into 3D space [80]. After an image segmentation, the location of the markers (usually spherical markers) are extracted and analysed, which in turn requires that the marker-to-background contrast is sufficiently high for their identification.

Image-based Methods

The image data of a dynamic PET acquisition is commonly interpreted via quantitative image analysis. Such an analysis can, for example, be used to deduce information on the tracer accumulation within a specific target region in the PET data. For motion quantification, similar statistical image analyses are available but allow to derive information on the spatial orientation of the scanned object instead.

Statistical Parameter Mapping (SPM) is one method to achieve this task. SPM is used to test for a specific null hypothesis within the image data, usually an equivalent or regional physiology or absence of correlation in different volume data sets. SPM calculates sub-images of change or correlation significance which can be used to derive motion information from [81]. By fitting the individual time-relative images of a dynamic PET study into SPM, the six parameters describing the motion are calculated. Another method with a similar approach of analysing differences in the image data is the use of three dimensional optical-flow analysis methods [82]. These methods calculate the motion for every pixel in two images taken at different times. By analysing each voxel at a specific location and by comparing it to a later image, motion information can be derived from inspecting the intensities of a small window of neighbouring pixels. This process is repeated for a defined sub-area of the whole image and is based on the assumption that the selected neighbouring pixels in that window have a cohesive relationship regarding their pixel intensities.

Analogue Motion Sensors

If absolute coordinates to quantify motion are not required, analogue sensor-based motion capturing is often used. To identify and capture the different phases of the respiration or the cardiac movement of a patient, different types of analogue sensor-based solutions are available (cf. figure 1.16).

For respiration, all these solutions have in common that they are able to output an analogue signal on the different phases (*inhalation, expiration*) of respiration. This is *inter alia* used to derive the breathing rate. They only differ in how this information is sampled from the patient. Some solutions, for example, capture the motion via gas pressure sensors [83] which induce a small electric current with increasing pressure on a worn respiration belt. Other solutions use piezo-electric devices that respond linearly to changes in length of the belt. If worn by a patient, the belt expands due to the extension of the chest which causes the output of an analogue signal. This signal is then digitalised via an analogue-digital (A/D) converter [84].

A similar technique can register and quantify the cardiac motion during a PET examination. In analogy to the motion quantification of respiration, an *electrocardiogram* (ECG) can register the electrical activity of the heart. This device allows to derive and capture the overall rhythm of the heart over time by means of attached electrodes (electrical contacts). A heart rhythm curve can be recorded by direct measurement of the activity of the ECG via the commonly available analogue connections on such a device. By analysis and classification of the different heart rhythm phases, a statement on the current position, and therefore on the motion of the heart is generated.

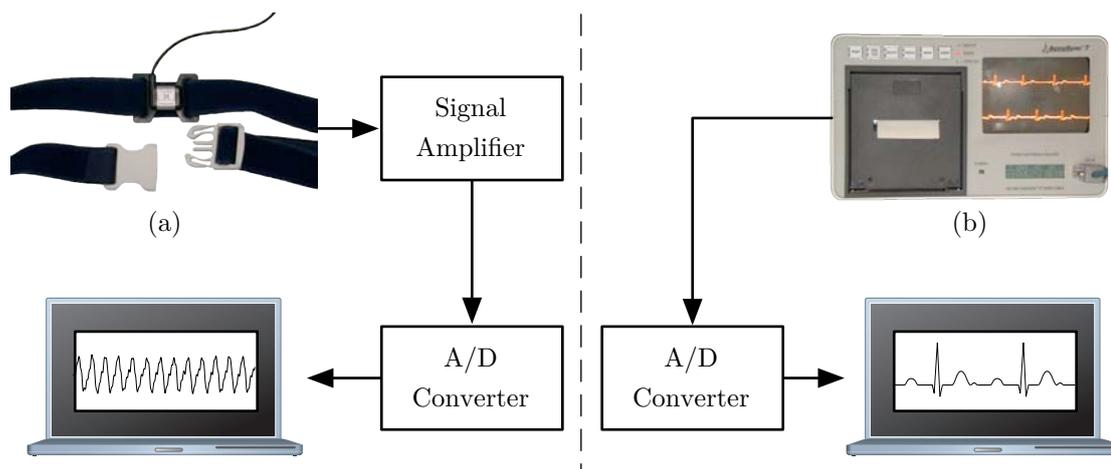


Figure 1.16:

Schematic illustration of the use and data processing of analogue signal-based motion sensors: On the left-hand side a respiratory belt (a) is used to sample the motion of the chest that results in an induction of a small electric current due to the use of a piezo-electric component in the belt. On the right-hand side, a common ECG device (b) is used to sample and quantify the electrical activity of the heart and thus the different cyclic phases of the heart motion. [12]

Optical Tracking Devices

Thanks to developments in the field of computer graphics and computer animations (motion capturing), different kinds of commercial motion tracking devices for quantification of external body motion are available today. Optical systems have improved tremendously over the years and allow to track motion not only with a high spatial resolution of ≈ 1 mm but also with a high sampling rate of ≈ 60 Hz. In addition, most of these optical systems work in the infrared light spectrum to overcome potential visible light influences, such as changing ambient light conditions [85]. Similar to video camera-based motion quantification, *markers* are placed within the FOV of the PET system. However, due to its operation in the infrared light spectrum, the emitting (*active tracking* [86]) or reflecting (*passive tracking* [87, 88]) markers are more easily distinguishable than the markers in the video camera-based solutions (cf. figure 1.17). In addition, optical tracking devices are often stereoscopic setups. This means, two or more cameras are set up to cover the FOV under different perspectives. Because of the individual perspective of each camera – similar to the stereoscopic nature of human eyes – the spatial coordinates of the marker positions are calculated from the two-dimensional images each single camera generates.

Active motion tracking systems allow for more difficult setup scenarios thanks to the possibility to place the active markers at more variable positions. This is possible because active infrared light flashes located on each single marker usually emit a higher signal than the signal passive markers reflect to the cameras. This allows for a setup where the marker is farther away from the camera system. In contrast, passive systems have the advantage that they are less expensive and easier to maintain. This is because for active systems each marker requires an electric source to produce the infrared light and a way of synchronising their flash frequency, which is often achieved via cables. If a large set of such markers is placed on a trackable object, the amount of cables is often a clear handicap for clinical application. Therefore, passive systems where the

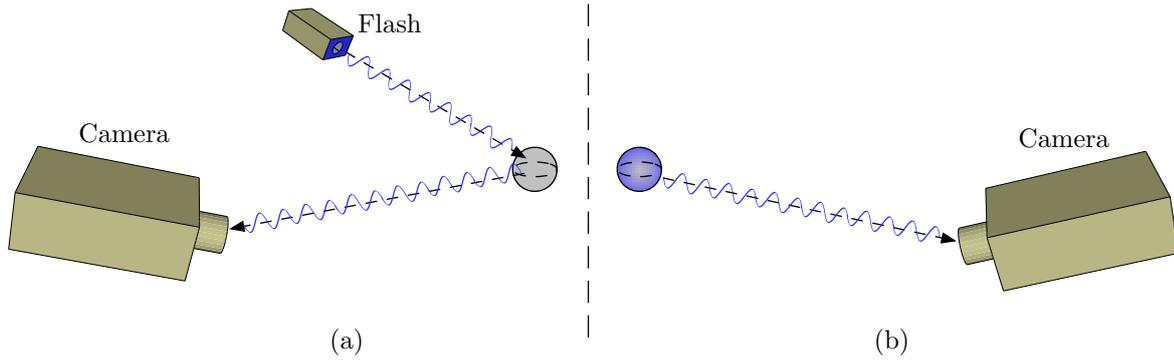


Figure 1.17:

Illustration of the difference between 'active' motion tracking and 'passive' motion tracking: On the left-hand side (a) passive tracking is shown with a reflective spherical marker. This *passive* marker has a reflective surface and thus reflects the light (infrared light) emitted by a flash (either external or directly embedded in the tracking camera). In contrast, the right-hand side (b) shows an *active* marker emitting light itself which is then sampled by the respective motion tracking camera.

infrared light is generated by a strobe, either embedded in the camera itself or external, are often the preferred choice.

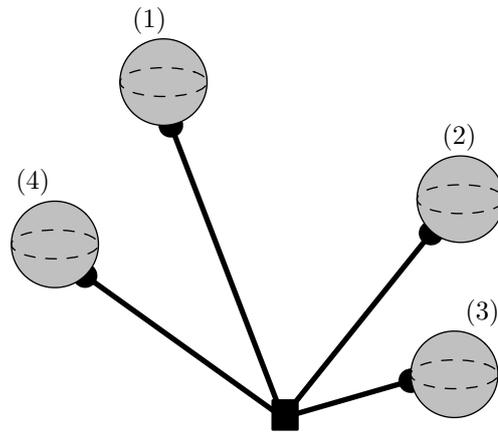


Figure 1.18:

Illustration of the combination of a set of individual tracking markers (grey spheres) and a single tracking body (connected spheres): Due to the unique respective distances between all involved markers in a tracking body, some optical tracking systems [87] are able to identify and quantify motion of the body relative to the centre of gravity of the tracking body construct.

In addition to the individual marker positions, optical tracking systems commonly allow a set of three or more markers, with a unique distance to each other, to be grouped together as a single logical *body* (cf. figure 1.18). These spatial bodies are identified by a tracking system that is able to distinguish if a certain *tracking body* is currently within the FOV or not. Furthermore, optical tracking systems allow to output the captured motion data directly with the six degrees of freedom of each body position. In addition to the three rotation angles (η, θ, ϕ), these systems also usually provide a time stamp information as well as the three translation parameters (P_x, P_y, P_z) relative to the centre of gravity of each tracking body.

1.2.3 Motion Compensation

Different approaches and algorithms exist to perform a motion compensation on PET volume data [89]. As table 1.1 shows, the most appropriate method for a correction mainly depends on the type and class of motion that is likely to occur during the PET examination. For some examinations, where only a certain area of the body is analysed, a *gated acquisition* may be the preferred choice. For others, however, methods to realign either the image data or to realign each single coincidence based on external motion data are necessary. All of the commonly used and well-documented methods are outlined throughout the next few sections giving an overview over the possibilities to perform a motion compensation in PET.

Gated Acquisition

A common and widely used method for a compensation of the influences of motion of a periodic nature is the splitting (*gating*) of acquisition data according to the periodic motion phases, a so-called *gated acquisition* [90]. During such an acquisition, the dynamics of motion are analysed either online or after the acquisition, and different recurring periodic phases within the motion are identified. According to these phases/intervals, the acquisition data (in PET the coincidence data) are sorted into different *gates*. Each gate is handled separately which results in multiple images of the same object but ‘taken’ at different motion states.

For a medical examination method like PET, periodic movement of the heart and the periodic nature of breathing are common types of motion which can be compensated for by gated acquisition. During a cardiac PET examination, the patient is constantly monitored via an ECG device. The different heart phases are then identified via an analysis of the electro-cardiac information recorded on the ECG (cf. figure 1.19). For an examination of the chest where the breathing of a patient has a direct influence on the image quality, the same process of gated acquisition can be used. After acquisition of the chest motion which results from the patient’s constant breathing, the PET data are split into different gates according to the determined phases of the breathing. Separate images in one of the main phases (inhalation, exhalation) are generated via gated acquisition – a process that is commonly used in different medical imaging techniques, not only PET.

Multiple Acquisition Framing

Multiple acquisition framing (MAF) is a motion compensation method that allows to split acquisition data into separate units (*frames*) based on the magnitude of the motion [13]. The acquired motion information is analysed immediately after data acquisition and in case it reaches a relevant level, subsequent coincidence data from the PET acquisition are put into a new frame until the next ‘significant’ motion is identified. Later, these MAF frames are reconstructed and the resulting images are spatially aligned to each other via single rigid volume transformation. As a result, a single data set is generated where each frame in the data set has the same spatial orientation, thus the whole data end up being motion compensated (cf. figure 1.20).

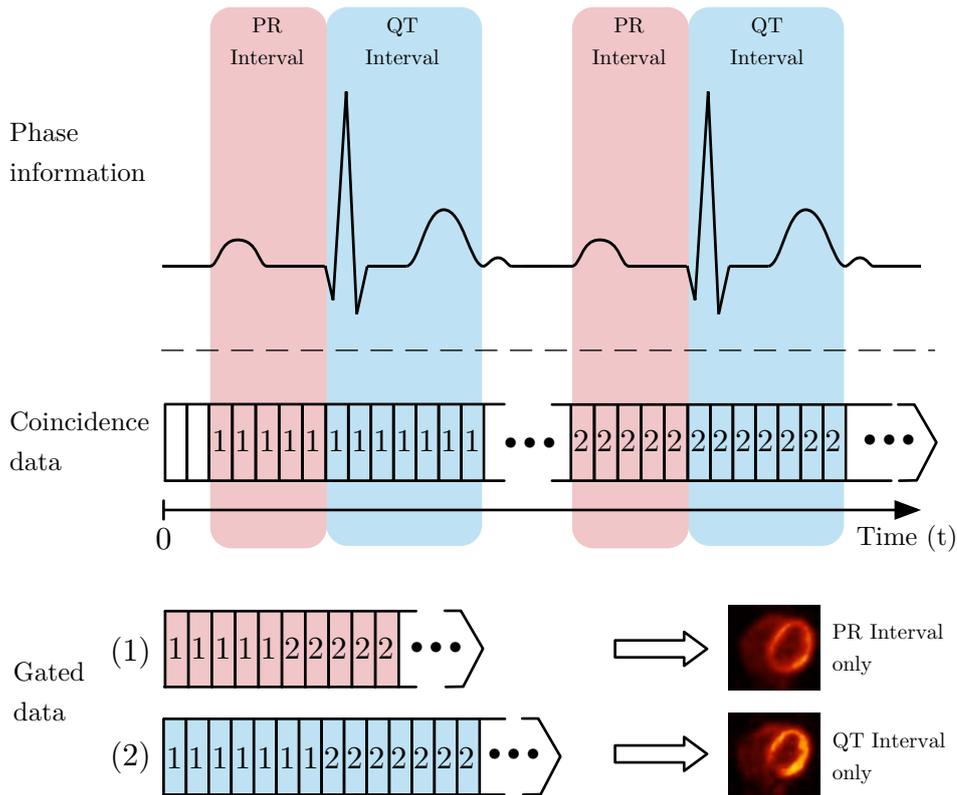


Figure 1.19:

*Schematic illustration of the principles of gated acquisition – here on a hypothetical example of a gated heart acquisition: Based on different heart phases – here PR interval and QT interval – the acquisition data are split into different *gates* so that separate images are reconstructed and therefore show the heart in a different state. This procedure can be, in principle, transferred to respiratory gating where the periodicity of breathing is used to separate the data according to, e.g. inhalation and exhalation phases.*

Coregistration

Image coregistration is used to spatially align two or more images relative to each other via statistical analysis of the data of volume images. This process is mainly used to align PET images from the same subject which have been taken at different examinations (i.e. *follow-up* studies). For PET images only, the scans of the same subject generally look quite similar and so it is often sufficient to find the rigid-body transformation parameters by minimisation of the sum of squares difference from the PET image data. However, for a coregistration of images from different modalities (e.g. PET and MRI), a registration is less straight-forward and usually requires more advanced coregistration algorithms [91]. If no external alignment information, such as visible landmarks within the data, is available a manual editing or alignment is often required for a proper coregistration of images from different modalities. Only for modern hybrid-scanners, like PET/CT scanners, a direct overlay and alignment of the PET and CT image is available due to the known translation between the used scanner ring in each modality.

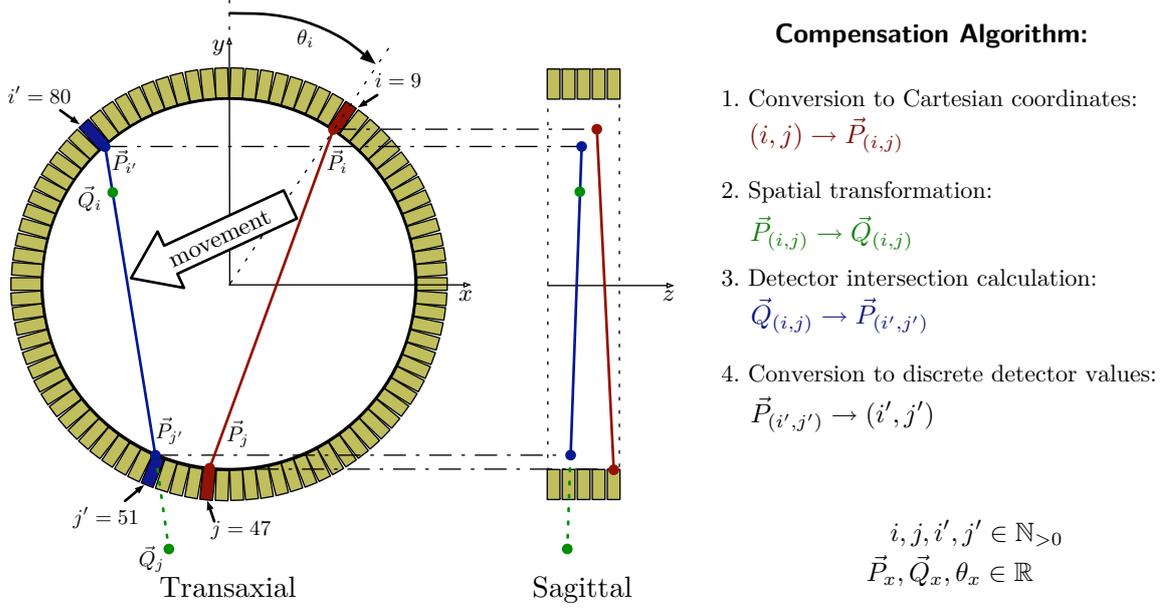
For dynamic PET studies, where the coincidence data are split into different logical units (*frames*), a coregistration can also be used to identify and compensate for motion. Instead of aligning volume data from different modalities or follow-up studies, a coregistration can be used to spatially align volume data of individual frames. By choosing a reference frame with a high count rate, a coregistration allows to spatially realign all other frames to the same orientation as the chosen reference frame. Similar to the problems of gated acquisition a coregistration cannot compensate for the motion that occurred within a frame. Furthermore, due to the rigid alignment of the volume data, the image resolution of the moved data is slightly reduced. This is caused by the required voxel interpolations which are necessary for an image coregistration.

Event-driven Motion Compensation

A motion compensation method that utilises the raw data of a PET scanner is the *event-driven motion compensation* [15, 92]. In principle, this correction method takes the raw coincidence information of the involved detectors of an LOR and reassigns it to a new detector tuple according to the available motion information of the analysed subject. After reorientation of all LORs, a motion compensated image is then reconstructed from the corrected data set.

While this approach, due to the use of each single coincidence event, has the highest potential for the most accurate motion compensation technique in PET, it also has the highest demands on the data processing itself. Event-driven processing not only requires the PET system to run in list-mode so that each single LOR is registered instead of binning (compressing) them into a sinogram. It also requires certain post-correction methods [31], e.g. an additional *normalisation correction*, *LOR discretisation correction* and an *Out-of-FOV correction*. While this section concentrates on discussing the principles of the motion compensation algorithm only, the details of these post-correction methods will be explained in section 2.4.

As discussed previously, an LOR is properly defined by the index numbers of the two involved γ -detectors (i, j) whereas the indexing starts at the top centre of each detector ring and increases clockwise (cf. figure 1.21). The actual spatial transformation of an LOR operates, however,

**Figure 1.21:**

Schematic illustration of the event-driven motion compensation algorithm: Each involved detector number (i, j) of an LOR is converted into Cartesian coordinates (1). Afterwards, these coordinates are spatially transformed (2) and new detector intersections (3) are derived. Finally, the transformed coordinate points are converted into new discrete detector combinations (i', j') (4), i.e. they are put into a sinogram.

within the Cartesian coordinate system. Therefore, (i, j) have to be converted via cylindrical coordinates to the Cartesian coordinates $\vec{P}_{(i,j)}$ with

$$\vec{P}_{(i,j)} = \begin{pmatrix} r \cdot \cos \theta_{(i,j)} \\ r \cdot \sin \theta_{(i,j)} \\ z_{(i,j)} \end{pmatrix}, \quad (1.16)$$

where r is the radius of the detector ring, $z_{(i,j)}$ the coordinate of the centre of the axial detector ring along the Z-axis and $\theta_{(i,j)}$ the angle at which the centre of a detector is located on the ring circumference and which is given by

$$\theta_{(i,j)} = \frac{N_{(i,j)}}{2\pi} \cdot (i, j), \quad (1.17)$$

with $N_{(i,j)}$ as the total number of detectors available in one ring plane.

After the conversion into Cartesian coordinates, the actual spatial transformation takes place where $\vec{P}_{(i,j)}$ is transformed into $\vec{Q}_{(i,j)}$ via

$$\vec{Q}_{(i,j)} = \hat{R} \cdot \vec{P}_{(i,j)} + \vec{t}, \quad (1.18)$$

where \hat{R} is the rotation matrix and \vec{t} the translation vector of the motion information gathered during the motion quantification (cf. section 1.2.2). As a result of this transformation, the new calculated LOR endpoints do not necessarily correspond to a new set of valid detector coordinates on the ring circumference. Therefore, to assign a new detector pair (i', j') , the

intersection points $\vec{P}_{(i',j')}$ of the transformed LOR coordinates $\vec{Q}_{(i,j)}$ with the detector ring are calculated by

$$\vec{P}_{(i',j')} = \vec{Q}_j + q_{(i,j)} \cdot (\vec{Q}_i - \vec{Q}_j) , \quad (1.19)$$

with $q_{(i,j)}$ as the two possible solutions of the quadratic equation

$$0 = q^2 \cdot (\Delta x^2 + \Delta y^2) + 2 \cdot q \cdot (Q_{j_x} \cdot \Delta x + Q_{j_y} \cdot \Delta y) + (Q_{j_x}^2 + Q_{j_y}^2 - R^2), \quad (1.20)$$

where

$$\begin{aligned} \Delta x &= Q_{i_x} - Q_{j_x}, \\ \Delta y &= Q_{i_y} - Q_{j_y}, \\ R^2 &= P_x^2 + P_y^2, \end{aligned}$$

and which is equivalent to a line extension of $\vec{Q}_{(i,j)}$ with the two points of intersection $\vec{P}_{(i',j')}$ on the ring circumference.

As the last step of the correction algorithm, the new LOR intersection coordinates $\vec{P}_{(i',j')}$ are converted into the discrete detector tuple (i', j') by the inverse operation of equation 1.16. During a motion compensation process, these new detector pair indices are then used instead of the original ones, and thus the spatially reoriented LOR is used during the image reconstruction.

In practice, each single LOR of a list-mode data stream is spatially reoriented according to the available motion information. Before this realignment is performed, all detector index tuples encoded in the list-mode data are converted into Cartesian coordinates. Then they are spatially reoriented. Afterwards, the new detector tuples are stored into a sinogram which then consists of the reoriented LORs only. This sinogram is then used during image reconstruction. As a consequence, all LORs are in fact reoriented to a position where they would have been registered in case the examined subject remained still.

1.3 Medical Nomenclature

In medicine standard definitions exist for the localisation of a certain area or structure of the human body. In the present work, where we concentrate on the anatomy of the human skull and brain mainly, important definitions should be given in the following, as they are used throughout this thesis.

In the literature, the mean perimeter of a human head is documented to vary between ≈ 54 cm for female adults and ≈ 56 cm for male adults, respectively [94, 95]. In humans, the adult skull consists of 22 bones which are, except for the *Mandibula*, joined together by small *sutures*. Eight of these bones form the *neurocranium* which includes the *Os frontale*, two *Os parietale*, *Os occipitale*, *Os sphenoidale*, two *Os temporale* and *Os ethmoidale*. Fourteen bones form the *splanchnocranium* which comprises the bones supporting the face. In addition to the definition of the human skull structure, the bone names are usually used as a means of providing a localisation information when referring to a human head. Figure 1.22 illustrates the localisation of the most prominent bones with their respective names.

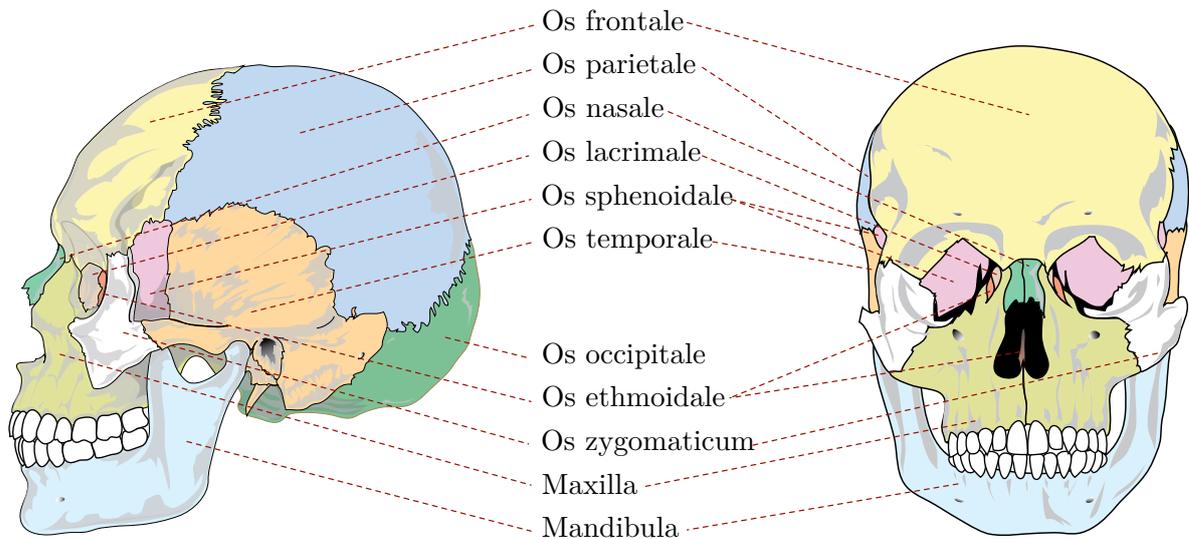


Figure 1.22:

Human Skull Anatomy [93]: The human skull of an adult consists of 22 bones. Here the names of the most prominent ones are given.

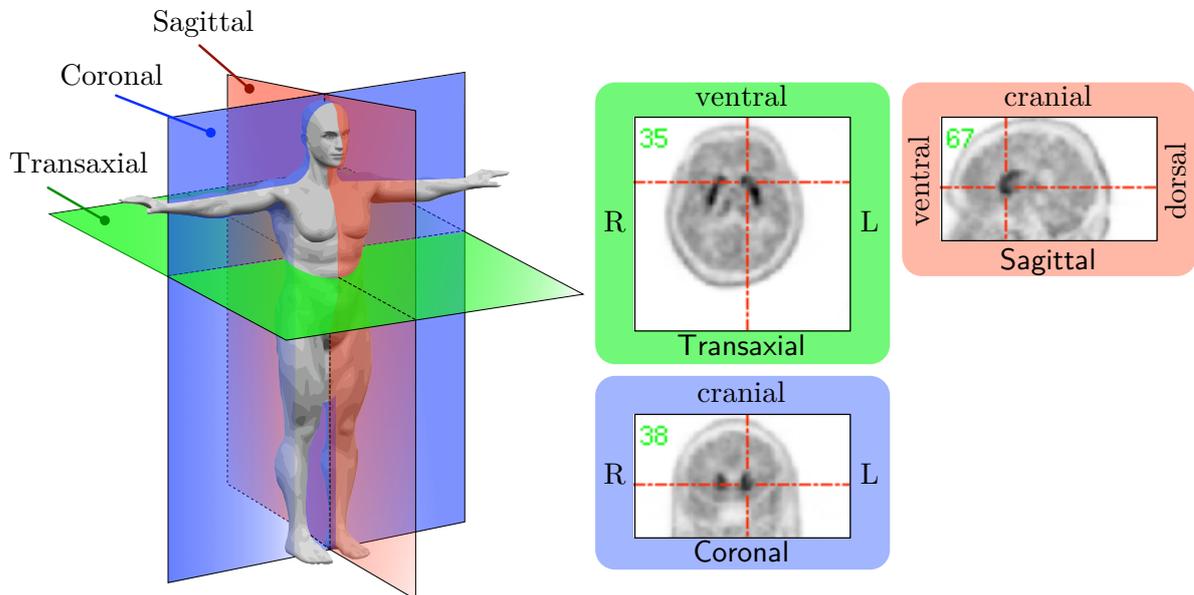


Figure 1.23:

Medical Imaging Plane Nomenclature [96]: In medical imaging three main planes exist. The *transaxial* (often *transverse*) plane divides the body into superior (cranial) and inferior (caudal) parts. The *coronal* plane divides the body into anterior (ventral) and posterior (dorsal) parts. The last one is the *sagittal* plane which divides the body into left and right sections. On the right-hand side of the figure the common orientations for each plane are given for a cross-section through the human brain.

Furthermore, in medical imaging the location of a body part in relation to another part is given by three planes. As shown in figure 1.23, the *transaxial* (also often *transverse*) plane is an imaginary plane dividing the body into superior and inferior parts (top and bottom). In contrast, the *coronal* plane divides a body into anterior and posterior (front and back) sections. For a human, the mid-coronal plane would transect a standing body into two halves in an imaginary line that cuts through both shoulders. Finally, the *sagittal* plane is an imaginary plane dividing the body into left and right sections.

As for the standard orientation of these planes, a transaxial image is by definition to be displayed with the anterior surface (ventral) up and the right side of the body on the left. When a series of transaxial sections are presented, they are shown from above down. Coronal images are usually presented with the upper section (cranial) on top and the right body side on the left, as though one were facing a standing patient. A series of coronal images are presented from front (ventral) to back (dorsal). Sagittal images are usually visualised from the right with the upper section (cranial) on top and the anterior surface (ventral) on the left. A series of sagittal images usually goes from right to left.

Based on those definitions medical software applications which allow to visualise the three-dimensional PET data generally use these planes for a well-defined display on a computer monitor. In addition, medical report documents usually contain images of single slices of the three-dimensional data which are based on this nomenclature.

2 Material and Methods

When I examine myself and my methods of thought, I come to the conclusion that the gift of fantasy has meant more to me than any talent for abstract, positive thinking.

(Albert Einstein)

2.1 Material

For the development and evaluation of the methods presented in this chapter, an ECAT Exact HR⁺ PET scanner has been used (cf. figure 2.1). In addition, a commercially available motion tracking system (ARTtrack) has been used to develop and evaluate the motion tracking and motion analysis enhancements which are an essential part of this thesis (cf. figure 2.6).

2.1.1 PET Scanner: ECAT Exact HR⁺

The ECAT Exact HR⁺ PET scanner series was developed in 1993 by Siemens/CTI, Knoxville, Tennessee and the model used throughout this thesis was installed at the PET centre Dresden-Rossendorf, Germany in 1996. It uses Bismuth germanate (BGO – Bi₄Ge₃O₁₁) as scintillator crystals. One crystal (*detector*) has a size of 4×4.5×20 mm (W×H×L) and is combined in a 8×8 array with other crystals as a *detector block*. Each detector block is connected to a 2×2 photomultiplier array similar to the example shown in figure 1.4. The arrangement of 72 of such detector blocks in a *detector ring* with a diameter of 80 cm, and the stacking of four detector rings in axial direction result in a total of 18432 γ -detectors. This represents an actual axial extension of 15 cm.

However, because of technical limitations and geometrical considerations, not all of the possible detectors combinations among the 18432 detectors are available. Some detector combinations are, for example, excluded because of the steep angle of LORs between two adjacent detectors. This, in fact, defines the maximum extension of the FOV in radial direction as can be seen in the illustration in figure 2.5. Also, to simplify the electronics and to reduce the influence of the electronics' dead time, 12 detector blocks are combined to a *detector bucket* which is connected to a single controller board. Each of the 96 buckets include 192 detectors not only from one detector ring, but from adjacent detectors in axial direction as well. In addition



Figure 2.1:
ECAT Exact HR⁺ PET scanner

to the standard 2D processing with extracted septa rings minimising the amount of scattered events, the scanner also allows to process an acquisition in 3D mode where the septa rings are retracted (cf. section 1.1.1).

Figure 2.2:

View into the inner part of the ECAT Exact HR⁺ scanner with its covers removed: The right half of the first detector ring with the circular arrangement of the detector blocks is shown. In addition, the combination of multiple detector blocks in one *bucket* with the corresponding electronics is visible on the right-hand side of the figure.



Furthermore, the HR⁺ is equipped with a relocatable bed which allows whole-body examinations to be performed in addition to single position acquisition for, e.g., neurological or cardiac examinations.

Coordinate System

The coincidence data sampled by the HR⁺ are either stored in sinogram data format or in list-mode data format which has been outlined in section 1.1.1. However, for both these data formats the same type of coordinate system is used to encode the LOR information, as illustrated in figure 2.3.

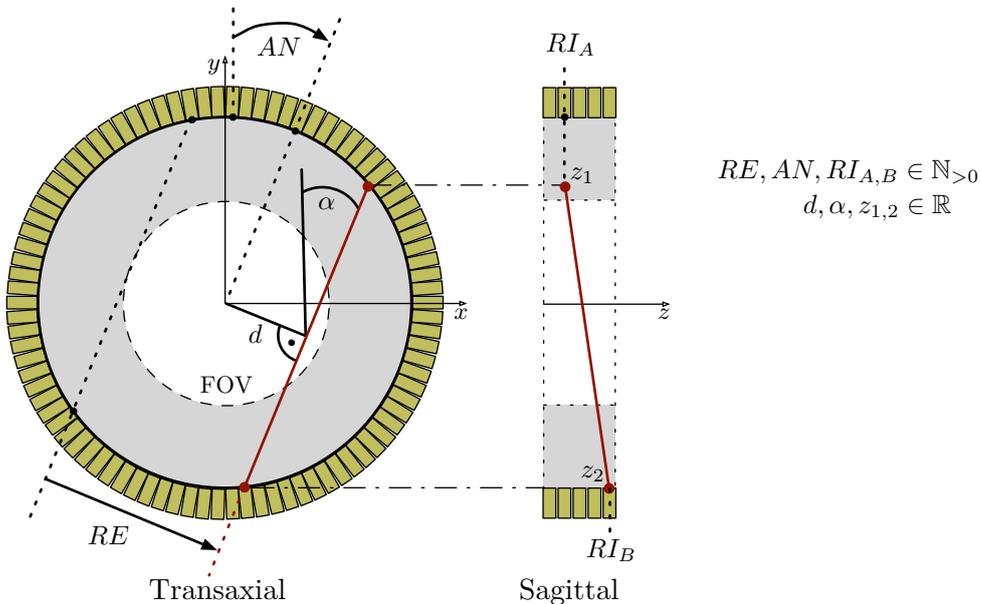


Figure 2.3:

Illustration of the standard coordinate system of an ECAT Exact HR⁺ PET scanner: Here, an LOR is defined by discrete numbers ($\mathbb{N}_{>0}$) for the radial element (RE), view angle (AN) and the ring combination ($RI_{A,B}$). In addition, an LOR is also often defined within a polar coordinate system via real numbers (\mathbb{R}) by a radial distance (d), a view angle (α), and axial distances for both end points (z_1, z_2).

Within the HR^+ coordinate system, an LOR is defined and encoded by integer values ($\mathbb{N}_{>0}$) based on the physical detector geometry. Therefore, in the coincidence data encoding of an HR^+ , an LOR is fully described by the *radial element* value (RE), the *view angle* value (AN) and two *ring* values (RI_A, RI_B). These values are discrete detector count numbers and also expressed in the polar coordinate system, with real numbers corresponding to the *radial distance* (d), the *angle* (α) and the two *axial distances* along the Z-axis (z_1, z_2).

Interleaving Implementation

Also part of the HR^+ functionality is, as explained in section 1.1.1, a general interleaving technique for increasing the spatial LOR sampling by linking detector combinations of two adjacent view angles [97]. While the interleaving technique is commonly used with PET scanners, its implementation varies.

In the actual interleaving implementation of an HR^+ , as shown in figure 2.4, a view angle always starts with an interleaved LOR, followed by a direct detector combination. According to this method, the number of contributing detector combinations in one single view angle is increased. This becomes especially clear on the right-hand side of the figure where near the centre of the FOV the interleaved LORs increase the spatial sampling due to the additional detector combinations.

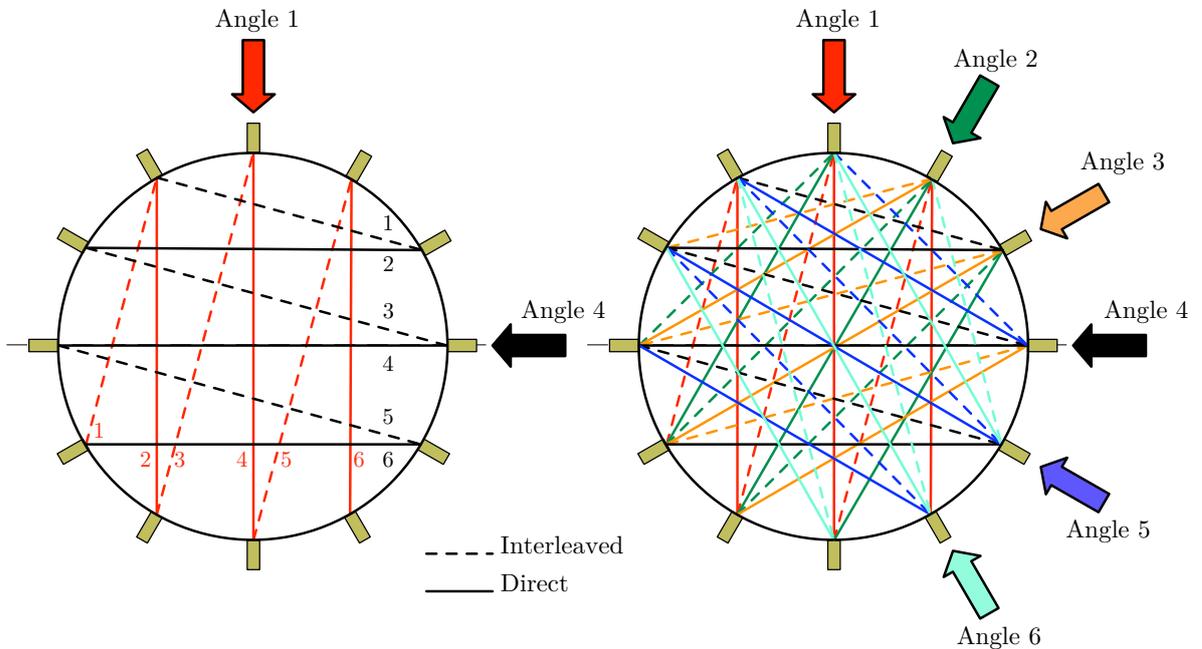


Figure 2.4:

Illustration of the interleaving technique implemented in the ECAT Exact HR^+ with a small sample of the 144 available view angles (on the left-hand side with 2 and on the right-hand side with 6 exemplary view angles): It demonstrates that by introducing interleaved detector combinations the spatial sampling, especially near the centre of the FOV, is increased.

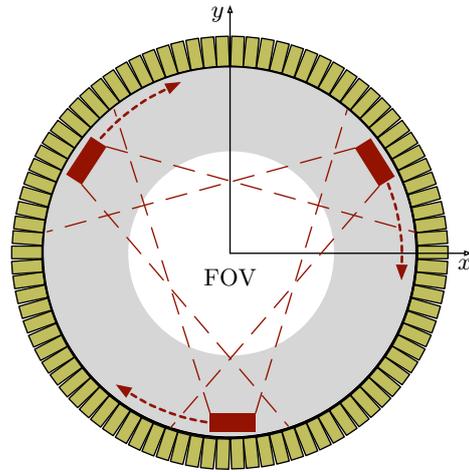
Attenuation Measurement

As shown in section 1.1.1, not all γ -quanta emitted by the positron-electron annihilations in the FOV are actually registered by the detectors of a PET scanner. Some of the γ -quanta are absorbed by the surrounding matter while they pass through. When performing analysis on PET data, a measurement and compensation of this attenuation is essential as otherwise the images would come up with an inhomogeneity.

For the ECAT Exact HR⁺ scanner, an attenuation measurement is performed via a *transmission* scan. During such a transmission three radioactive β^+ -emitting [⁶⁸Ge] line sources, which are normally located within a lead sheath, are automatically brought into the FOV of the γ -sensitive detectors (cf. figure 2.5).

Figure 2.5:

Three radioactive sources [⁶⁸Ge] automatically brought into the FOV of the HR⁺ during a transmission scan: By rotating these sources around the FOV of the scanner, a three-dimensional map with coefficients on the actual attenuation of the surrounding matter is calculated. During the image reconstruction this attenuation information is used to compensate for the inhomogeneity caused by the absorption of γ -quanta.



These positron emitting sources are rotating around the analysed subject during the whole transmission, thus illuminating the whole FOV of the scanner. According to the number of counts registered at the opposite of the current position of each radioactive source, the attenuation along the LOR caused by the surrounding matter (e.g. the bed, patient) is measured. This information is used and combined with attenuation data of an empty FOV (*blank scan*). As a result, attenuation factors for every sinogram bin, and thus a three-dimensional attenuation sinogram, are generated. This attenuation sinogram is then used during the image reconstruction and combined with the emission data such that the sinogram bins are scaled according to the attenuation factors.

For an HR⁺, an attenuation measurement is commonly performed in 2D mode. This is not only the case due to the reduced number of scattered events in 2D, but also due to the generally reduced amount of acceptable detector combinations. Otherwise, the high amount of activity resulting from the radioactive sources used would imply a loss of coincidences which is too high caused by the dead time of the acquisition system.

Acquisition Control System (ACS2)

The *acquisition control system* (ACS2) is part of the equipment of every ECAT-based PET scanner. It is the main system of the PET scanner and controls the gantry, samples the raw coincidence data and initiates the automatic movement of the bed in case of a whole-body scan [98]. Furthermore, it is responsible for saving the sampled coincidence data onto an internal

hard-disk drive. To control the PET scanner, the ACS2 has a direct fiber-optic connection to the electronics of the PET gantry through which all data communication is performed. The ACS2 itself is controlled by another workstation (ECAT console) which is used as the user front end. It comes with a graphical user interface, CAPP, and allows to perform evaluations on the sampled data as well as to generate standard medical report pages.

Data Formats

The data formats used by the HR⁺ to save the coincidence data are either the ECAT7 sinogram format or the list-mode data format. The ECAT7 sinogram format is a generic data container format for storing the binned coincidence data of a histogram-mode acquisition. Together with general acquisition related data, such as the patient meta data and the acquisition configuration, the coincidence data are stored in large three-dimensional arrays. The three dimensions of the stored data in those arrays correspond to the coordinate system description outlined in figure 2.3.

In contrast, the list-mode format, which is used to store the raw coincidence data of list-mode acquisitions, is a very basic data format. The coincidence data encoded according to the same LOR coordinate specifications as the compressed ECAT7 format are not binned but encoded into single 32-bit data words for every coincidence event. At a time interval of 1 ms, a *time word* is regularly inserted into the list-mode data stream, as shown in the general description of list-mode in section 1.1.1. In addition to the time information, these HR⁺ specific time words also contain information on the gating input of the PET scanner. The gating inputs are digital data inputs following the *Transistor-Transistor-Logic* (TTL) standard. They can be utilised to include status information into the list-mode data stream, e.g. the different heart phases derived from an ECG. Upon a positive TTL signal on one of the gating inputs the corresponding gating bits are set in the 32-bit long time words. They can thus be evaluated by the software, which is processing the list-mode data.

2.1.2 Motion Tracking: ARTtrack

For tracking the motion of a patient during an examination, a passive motion tracking system has been used throughout our studies. The one used here (ARTtrack, [87]) is a commercially available system equipped with flash lights in each of the four installed cameras (cf. figure 2.6). The integrated flashes emit light in the infrared spectrum and are synchronised with each other so that one flash does not blind the other. The cameras themselves are equipped with digital *Charge-Coupled Device* (CCD) chips sensitive to infrared light. The tracking targets have a reflecting surface which is optimised to reflect infrared light. Each camera captures the FOV of the PET scanner at a rate of 60 Hz and the tracking targets reflect the infrared light of the flashes directly to the tracking cameras. All cameras have a network connection to a control system operating a software (DTrack) for con-



Figure 2.6:
ARTtrack Motion Tracking
Cameras

trolling the measurement. In addition, the sampled data (images) of all cameras are processed by this software and due to the stereoscopic orientation of the tracking cameras, each coordinate of the visible markers are calculated (cf. section 1.2.2). Furthermore, by means of analysing the distances between each of the sampled marker coordinates, the tracking system is able to identify a set of markers as a unique *tracking body*. For every tracking body identified, the position within the coordinate system of the tracking device is calculated. In addition, rotation parameters are calculated relative to the orientation of the tracking body at calibration time.

Data Formats

The sampled data of the tracking devices are first processed by the DTrack software and then supplied to other workstations via a network connection.

For transmission of the motion data, it is packed as a single UDP datagram per sample and sent via a network connection to other workstations. Every $\frac{1}{60}$ th second, the DTrack software encodes the current motion information of each identified tracking marker and tracking body by using the translation vector (\vec{P}) and the three rotation angles (η, θ, ϕ) (cf. section 1.2.2). The workstation receiving the UDP data decodes this information and saves it into a proprietary list-type data format with timestamp information for each motion sample.

2.2 Optimised List-Mode Data Acquisition

As already outlined, the acquisition of list-mode data is desirable for several reasons. For the application of event-driven data processing methods the acquisition of list-mode data is a necessity. For some PET scanners, however, the acquisition of list-mode data is not originally supported by the scanner environment – sometimes it is not available at all. This is the result of times when the handling and processing of the very large amounts (generally multiple gigabytes) of data was difficult.

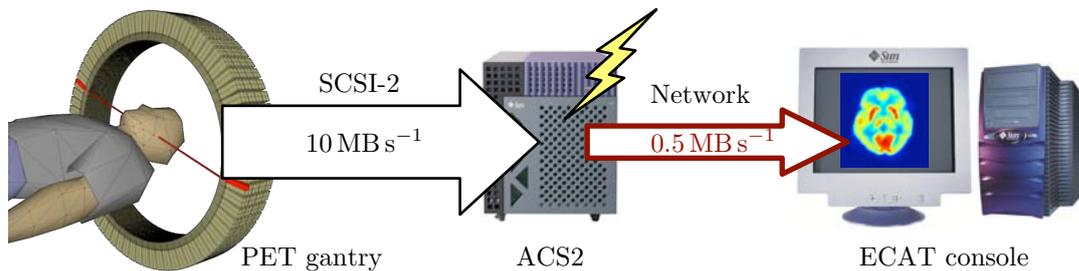


Figure 2.7:

Hardware limitations of the ACS2: Because of hardware limitations of the ACS2, the transfer of acquisition data to a workstation is practically limited to $\approx 0.5 \text{ MB s}^{-1}$, even if the data storage device itself supports higher bandwidths. For list-mode acquisitions, this leads to undesirable time delays which represent a severe limitation on the routine use of list-mode.

In our case, several aspects of the standard scanner environment prevent the routine use of list-mode with the ECAT Exact HR⁺ PET system. The missing capability to acquire list-mode data with the standard acquisition environment (ECAT console) directly is just one point

that has to be mentioned in this context. To acquire data in list-mode on an HR⁺ it is, for example, necessary to work with third party acquisition protocols which use proprietary software to transfer the list-mode data from the underlying acquisition control system (ACS2) to a processing workstation. More importantly, the limited network capabilities of the ACS2 system, as well as hardware and software stability issues, severely hinder the possibility of acquiring list-mode data with a standard HR⁺ PET scanner. To improve this situation, we have analysed the existing hardware and software constraints and have developed methods for routine acquisition of list-mode data on an HR⁺.

2.2.1 Shared Storage Device Acquisition

The ACS2 usually stores its acquisition data on an internal hard disk. Although this hard disk is connected to a standardised SCSI-2 bus which supports a maximum theoretical throughput of $\approx 10 \text{ MB s}^{-1}$, the main hardware of the ACS2 and its operating system components do not allow to transfer the data to other systems with rates higher than $\approx 0.5 \text{ MB s}^{-1}$. For the typical amount of approximately 4 GB to 6 GB of data for a list-mode acquisition, this results in a required transfer time of 2 h to 4 h.

To overcome this severe limitation, we connected a *shared storage device* to the ACS2 instead of the internal hard disk. This system complies to the official SCSI standard [99], which is well known for its backward compatibility, and therefore allows for transparent connection of newer SCSI compliant devices to the old SCSI-2 bus of the ACS2. Therefore, we chose a *Ultra160-SCSI RAID* system with a maximum theoretical throughput of up to $\approx 160 \text{ MB s}^{-1}$ and connected it to the ACS2. This device features a dual-channel functionality and thus makes it possible to connect two independent computer systems to the storage device while isolating the electrical circuits of the host systems from each other. This electrical isolation avoids interference between the different host systems while connected to the storage device.

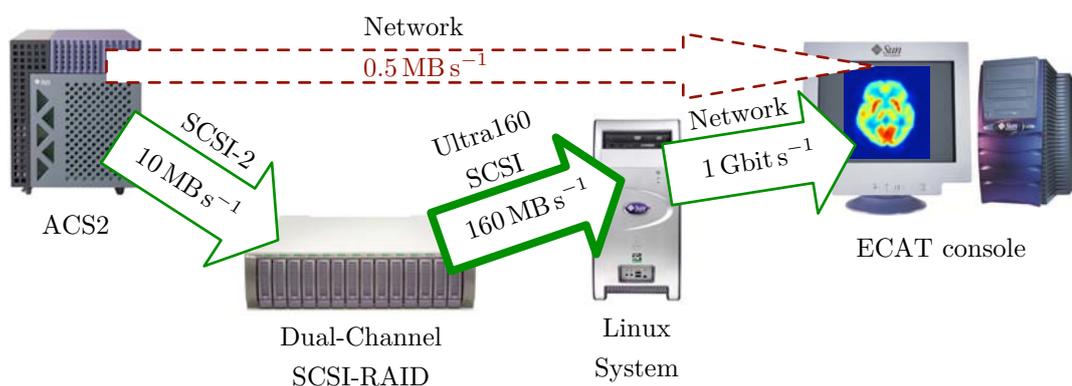


Figure 2.8:

Shared Storage Device Acquisition: A shared storage device (here an *Ultra160-SCSI RAID* system) is used to speed up the data transfer of acquisition data. By parallel connection of a *Linux* system, the acquisition data stored by the ACS2 can be retrieved as soon as the acquisition is finished.

By virtue of this galvanic decoupling, the two systems connected to the shared storage device can even utilise the storage with different transfer bandwidths. For the shared use on an ACS2

this means that, while on one device channel the ACS2 is connected via the slow SCSI-2 connection, another system can still address the full bandwidth of the Ultra160-SCSI device. In practice, this allows a workstation system to be directly connected to the shared storage device with an Ultra160 host adapter and thus with the highest possible bandwidth.

To implement such a shared storage device solution with the ECAT environment of an HR⁺ scanner, we connected the second channel of the RAID device to a server system with *Linux* as its operating system (cf. figure 2.8). While Linux itself comes with general support for SCSI and also with support for modern Ultra160-SCSI host adapters, it does not support the file system data structures that are used by the proprietary ACS2 system. Since the ACS2 is based on the *VxWorks 5.2* operating system, it uses the *VxWorks extended DOS file system (VXEXT)* to store its data. This file system is a derivative of the standardised FAT16 file system commonly available for many different operating systems. Due to certain requirements of the VxWorks operating system, however, VXEXT comes with a slightly different data structure. In contrast to FAT16, the VXEXT file system supports, for example, the storage of files with a filename length of up to 40 characters. In addition, it permits to maintain more than 2 gigabyte of data, which is necessary due to the large list-mode data. These differences render the VXEXT file system incompatible with the official FAT16 specification [100]. Therefore a standard FAT16 file system driver, as supplied by Linux, cannot be used. In order to still be able to access the data stored by the ACS2, we have reverse-engineered the data structures of the VXEXT file system and have developed an implementation for the Linux operating system using the C programming language [27].

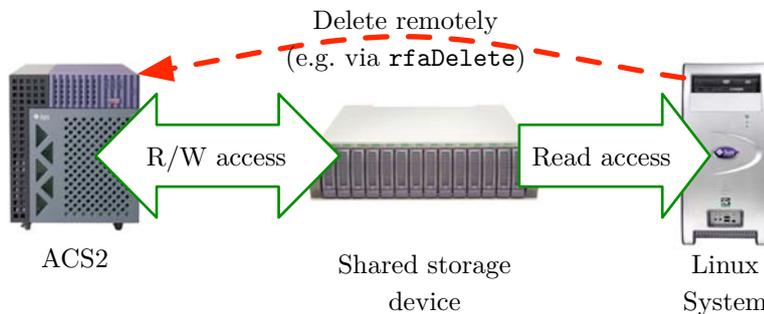


Figure 2.9:

Data access to Shared Storage Device: Because of file system cache considerations, a simultaneous write access to the *shared storage device* is not possible from multiple systems. However, only read access to the acquisition data from Linux is required, namely for transferring data to a processing workstation. Instead of direct write access, a remote execution of existing commands (e.g. `rfaDelete`) can be issued, if necessary.

Due to the data access management and data caching policies of an operating system's storage process in general, a simultaneous write access of both systems (ACS2 and Linux) to the VXEXT data partition is not possible. This is especially true since in case of data changes from the ACS2, any other system connected to the same file system is not automatically notified of the changes. While this type of data change notification is a common functionality of *distributed file systems* (e.g. DFS, QFS, etc.), VXEXT, as an ordinary file system is not capable of such distributed notifications and thus requires other means for a possible notification of data changes. One approach is to access the storage from the Linux system in read only mode while only the ACS2

is allowed to change information on the device. This is possible, because the only write access that is required from Linux is to delete a set of files after they have been copied successfully to an external workstation. Such a delete operation can be achieved via a remote execution of shell commands (e.g. `rfaDelete`) which are available on every ECAT console. Therefore, we have developed command scripts to execute these commands as soon as the copy operation of the acquisition data has been finished (cf. figure 2.9).

The Linux system also has to refresh the file system meta-information regularly in order to identify if data have been changed on the storage device. Otherwise it cannot keep track of potentially modified data since the last access of the ACS2 system and may return file system information which is already obsolete at the time of reading. This is particularly critical if the ACS2 has modified the data between different acquisition sessions of the PET system while the Linux system is still transferring data from the storage device. Therefore, to arrange for a synchronisation of the write operations of the ACS2 and the read operations performed by the Linux system, we have developed an automatic *remounting mechanism*. The VXEXT file system on the shared storage is automatically unmounted via the standard auto-mounter service of the Linux operating system. This happens after a fixed time period in case no further access is required. Furthermore, it is automatically remounted in case of a new read access from the Linux system. This has the effect that on every new remount of the VXEXT file system, the meta information is refreshed and thus the Linux system is able to list the data content which was present at the time of the mount operation.

Evaluation Protocol

To evaluate the correct operation and performance of the *shared storage device* method, an overall amount of 40 consecutive list-mode acquisitions have been performed and evaluated in four different test phases:

1. In the first test phase the hardware compliance, i.e. the stability and compatibility with the ACS2, has been evaluated via ten test acquisitions. As criterion for stability and compliance, the number of found dropouts or instabilities of the ACS2 system has been quantified.
2. In a second test phase, the data consistency of retrieved list-mode data has been analysed. The sampled data of ten test acquisitions was in one case transferred via an access through the implemented VXEXT file system and in another case via the original network-based *file transfer protocol* (FTP) of the ACS2. Both data sets have then been compared via a binary compare mechanism to identify potential divergences and thus inconsistencies.
3. In a third test phase the time required for the transfer of the acquisition data has been characterised. In this comparison we measured the time needed for the transfer of the data via the traditional transfer protocol of the ACS2 and for the transfer via the optimised access to the storage device.
4. In a fourth and final test phase, the stability of the auto-mounter functionality has been evaluated. The data on the storage device have been accessed immediately after each of

the ten test acquisitions. Simultaneously, the operation of automatically unmounting and remounting the storage device was reviewed.

2.2.2 DAQ-based Data Acquisition

One unresolved issue with the shared storage device method is, that a transfer of list-mode data is only possible after a PET examination has finished. This is uncritical for a regular histogram-mode examination due to the low amount of data. However, it is desirable to access the coincidence data in real-time while they are being sampled by the PET scanner. An immediate access to the raw data of the PET scanner is useful for time consuming list-mode processing methods, such as motion compensation or flexible re-framing via a list-mode sorter. To arrange for such a real-time data access, we have carried out a detailed technical analysis of the specific hardware components of the ACS2.

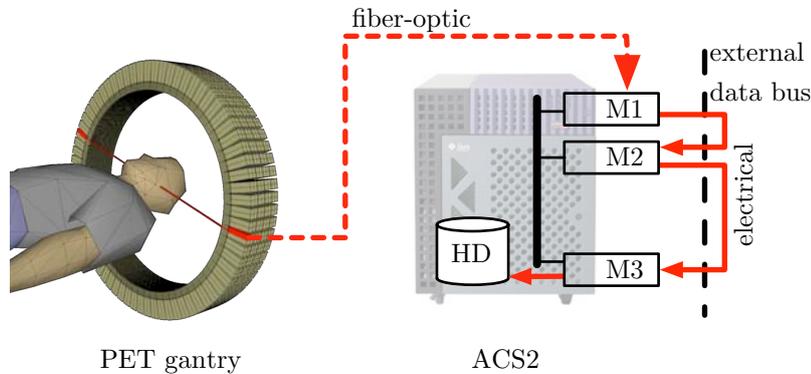


Figure 2.10:

Schematic outline of ACS2 hardware: Based on the Versa Modular Eurocard (VME) standard, the ACS2 consists of several modules (M1-M3) where the raw coincidence data are routed and processed until they are stored on an internal storage device (HD).

The hardware of the ACS2 is based on the *Versa Modular Eurocard* (VME) standard [101] and comes with several custom VME modules connected to an internal system bus. While some of these modules communicate through the internal system bus only, the main data processing modules communicate through an additional, proprietary data bus. This external data bus transfers digital data through consecutive flat ribbon cables accessible at the backplane of the ACS2 (cf. figure 2.10). These cables have 32 data channels and the digital data signals are based on the *Transistor-Transistor-Logic* (TTL) standard which are synchronised via an embedded clock signal. In principle, starting at the *fiber-optic module* (M1), the raw coincidence data from the PET gantry are converted into electrical TTL signals. In a consecutive step, the data are routed through different sub-modules until they pass the *rod converter module* (M2) where information on the rod transmission source position is calculated. From there, the data are either routed through the *hardware coincidence sorter* (VSB), or through the *read/write module* (M3) until they are saved on the internal data storage device. If operated in histogram-mode, the data routed through the external bus are directly sorted via the VSB module into a sinogram. Otherwise, during list-mode processing, the coincidence data are forwarded to the *read/write module* where they are directly stored on the hard disk in the list-mode data format.

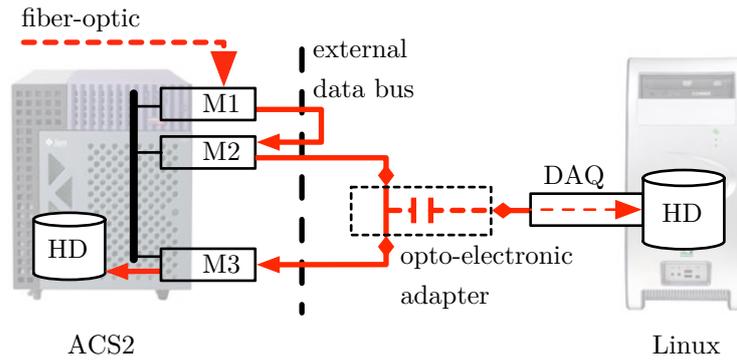


Figure 2.11:

Parallel data acquisition using DAQ cards: Digital acquisition cards (DAQ) are used to directly acquire coincidence data from the electronic components of the ACS2. The list-mode data can be obtained directly before the data is written to the internal storage device of the ACS2. This is done by sampling the data from the flat ribbon cables of the external data bus.

Based on this hardware analysis, a direct sampling of the coincidence data from the flat ribbon cables was considered (cf. figure 2.11). To directly retrieve the coincidence data from the external bus, we installed a set of *digital acquisition cards* (DAQ) into a computer system operating Linux. These data acquisition cards are commercially available PCI-based solutions [102] which allow for a sampling of the data of 16 digital signals each. By combining two cards, 32 digital TTL signals are sampled in parallel, permitting to cover the full 32bit bandwidth of an HR⁺ list-mode format. A TTL-based clock signal is provided together with the 32 data channels on the cables for the synchronisation of the digital signals. This clock signal is required because the ACS2 itself transfers the data on the external bus asynchronously without a fixed sample rate. This means that for every consecutive coincidence event the 32 data bits, in addition to the clock bit, are set by the fiber-optic module – all other modules are reading the clock signal status only. On each positive clock edge, the state of the TTL data signals are sampled and thus correspond to an actual valid bit combination. By forwarding this clock source to the available input on the DAQ cards, a synchronous sampling of the coincidence data has been achieved.

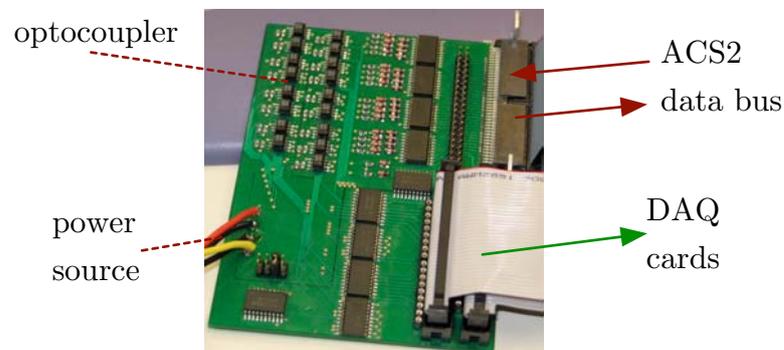


Figure 2.12:

DAQ hardware adapter card: During our work on the DAQ-based list-mode acquisition method, we have developed an adapter card. It uses opto-electronical components to decouple the electrical circuits of the ACS2 and the Linux system.

In order to retrieve the data from the cables without imposing any risk on the ACS2 circuits, an additional adapter card has been developed using opto-electronical components for decoupling the electrical circuits of both systems. This adapter card is shown in figure 2.12 and has been developed in co-operation with *Seiler IT Services* [103]. As illustrated, the adapter card is connected to the ACS2 data bus and the data bus of the DAQ cards in parallel. In addition, an external power source with 5 V (DC) is connected and provides the required power for the optocouplers. If connected, the data routed through the Y-cable on the ACS2 data bus side of the adapter are converted into light signals within the optocouplers and immediately converted back into electrical signals which are forwarded to the DAQ cards.

Evaluation Protocol

To verify the correct operation of the adapter card and the DAQ-based data sampling in general, an overall amount of 42 list-mode-based test acquisitions have been performed. During these test acquisitions, the stability, the data consistency as well as the performance of the data acquisition method has been analysed in four test phases:

1. In the first test phase, twelve list-mode acquisitions were initiated to verify that the DAQ cards or the use of the adapter card do not result in interference with the original ACS2 operations. For three of the twelve test acquisitions the adapter card has been removed, while for another three it has been connected to the ACS2 bus, but kept inactive. Then, the adapter card has been kept physically connected and activated for another three test acquisitions. Furthermore, the adapter card has been connected to the ACS2 bus for another three test acquisitions and the coincidence data has been sampled via the DAQ cards.
2. In the second test phase, the data sampling consistency has been analysed. In ten test acquisitions list-mode data have been sampled with the DAQ-based acquisition method and the traditional ACS2 acquisition method in parallel. If differences in the data exist due to discrepancies in the new DAQ method they ought to be identified by comparing the parallel acquired list-mode data.
3. Furthermore, the capability of the DAQ-based solution to sample list-mode data in real-time has been evaluated in a third test phase. In ten test acquisitions, the delay until coincidences could be actually sampled by the DAQ cards has been quantified. For this analysis the sampling software has been modified to calculate the time required for the internal sampling of the coincidences.
4. Finally, in a fourth test phase the stability of the DAQ-based sampling at high count rates has been evaluated. During ten test acquisitions, the overall amount of local radioactivity has been increased stepwise. This increase in radioactivity has been carried out to test the maximum count rate and thus to identify the sampling rate at which the DAQ hardware is not able to sample coincidences anymore.

2.2.3 Event-based Attenuation Measurement

With an HR⁺ scanner, a transmission scan to acquire the attenuation coefficients of the examined subject is usually performed in histogram-mode. This approach does not allow for the application of any time resolved motion compensation method such as, e.g., gated acquisition. This fact is unfavourable since a standard transmission scan commonly lasts about 10 minutes, i.e. a duration with a likelihood for patient movement already. If not corrected, the patient motion leads directly to an incorrect determination of the attenuation coefficients. To overcome this issue the technical requirements have been analysed and appropriate methods have been developed to perform the acquisition of a transmission in list-mode.

During a standard acquisition with an ECAT Exact HR⁺ a graphical interface, CAPP, is used to prepare and start the acquisition. This graphical environment executes various command-line programs in the background of the ECAT workstation during an acquisition. However, the graphical interface does not provide any functionality for an execution of the commands necessary to set up the PET scanner for list-mode acquisition.

Table 2.1:

List-mode Processing Command Overview: Overview of commands required for the set up and processing of a list-mode acquisition with an ECAT Exact HR⁺ PET scanner.

Command	Function
acqasync	Enable/Disable asynchronous statistics collection
acslock	Lock/Unlock the scanner for exclusive operations
move	Extract/Retract the septa for 2D/3D processing, Extract/Retract the transmission sources
rtsconf	Reinitialise the scanner environment
rtslm	Enable/Disable the list-mode processing
rtsacqm	Configure the acquisition mode (windowed/normal)
rtsdpln	Configure the data storage mode (2D/3D)
rtsdfrm	Configure the acquisition duration
rtslist	Set up the filenames for storing the list-mode data
rtsstrt	Signal to start processing
rtswait	Wait until acquisition is finished
setenergy	Set the permissible upper and lower energy levels
singles	Enable/Disable the <i>Singles</i> rate polling

Out of the several dozen command-line programmes available to specify certain aspects of the acquisition, only a small subset is necessary to set up and initiate a list-mode acquisition. The relevant commands have been identified by means of an analysis of the standard protocols as well as by comparison of the available scripts for configuring the scanner for list-mode processing. Table 2.1 summarises these commands together with a short explanation of each command.

Similar to the technical problem of emission-based list-mode acquisition, the standard graphical interface of an HR⁺ scanner allows for a transmission to be processed in histogram-mode only. Moreover, the acquisition of a transmission in list-mode is not available through already existing command-line scripts or programmes provided by the scanner manufacturer. During the execution of the standard graphical protocols, however, command-line programmes are executed in the background of the ECAT workstation, which set up the PET scanner for the

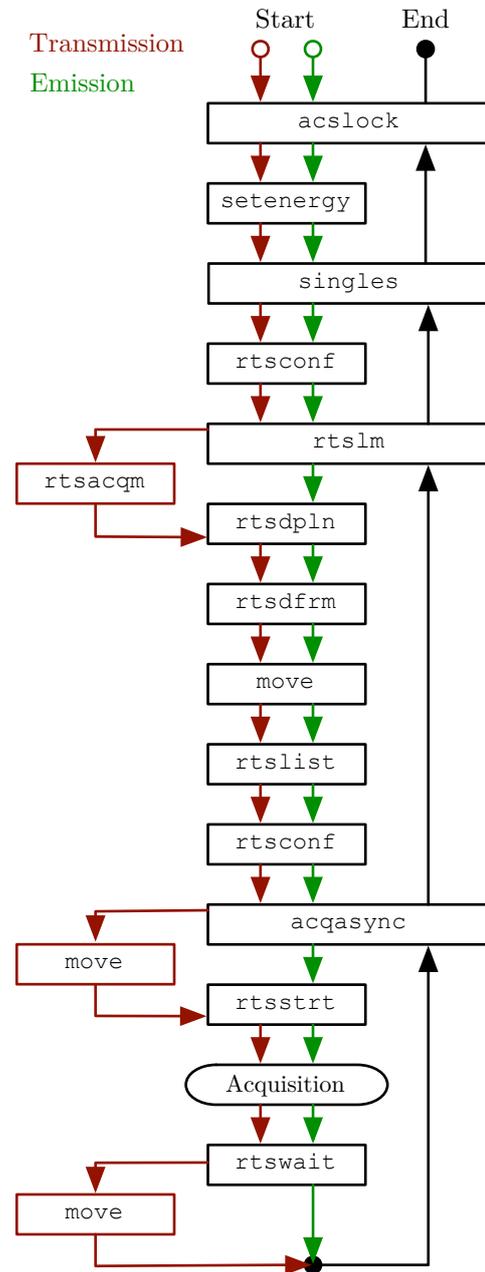


Figure 2.13:

ECAT list-mode command flow-chart: During a list-mode acquisition different command-line programmes are executed to set up an ECAT Exact HR⁺ scanner (cf. table 2.1). Here, the identified command execution flow-chart for the two main acquisition modes *transmission* (red) and *emission* (green) is shown. A difference between the two execution paths is the additional extraction and retraction of the radioactive transmission sources (*move*) as well as the set up of the acquisition mode to *windowed* via the *rtsacqm* call.

respective processing (emission, transmission). The relevant commands that set up an HR⁺ scanner for the processing of a transmission scan, including the extraction and retraction of the radioactive transmission sources, have been identified by an analysis of the histogram-mode protocols. The main differences have been found via a comparison of the commands to execute a list-mode emission on an HR⁺ scanner (cf. table 2.1) with the standard command execution of a histogram-mode-based transmission. This allowed for generating a command execution plan for a potential list-mode-based transmission measurement. A flow-chart of this execution is shown in figure 2.13.

After having implemented the new protocol for an event-based transmission, the acquired list-mode data was found to be slightly different from the standard list-mode format. While during an emission scan coincidences are always encoded as the same kind of *event words*, the data of a list-mode file acquired during a transmission scan has been found to differentiate between *windowed* and *non-windowed* coincidences. ‘Windowed’ in this context refers to coincidences that are actually known, due to the current transmission source position, to be directly originating from the rotating transmission sources while non-windowed coincidences are known to be of an arbitrary source and thus neglected. However, as the standard list-mode data format of an HR⁺ does not allow to mark coincidences according to their windowed state, a further analysis has revealed that windowed coincidences are marked as *scattered* events in a transmission-based list-mode data stream. According to these results an already existing list-mode sorter software (`lmSorter`) has been adapted to only sort events marked as *scattered events*.

Evaluation Protocol

To evaluate the new protocol for a transmission in list-mode 30 test acquisitions have been performed. In addition, the correct operation of the modification of the list-mode sorter software has been evaluated. For this evaluation the test acquisitions have been split into two different test phases:

1. In a first test phase, ten test acquisitions have been performed with the goal to evaluate the general stability of the new transmission protocol and its involved components. After each test acquisition the number of drop-outs and error messages found at the ECAT system has been quantified.
2. In the second test phase, the data consistency of the sampled list-mode data of 20 test acquisitions has been evaluated. During the test acquisitions a cylinder phantom has been placed on the bed of the PET scanner. For each of the ten acquisitions the phantom has on one occasion been sampled with the traditional histogram-mode-based transmission, and on another occasion with the new list-mode-based method. As a criterion for the concordance between the two images an intensity correlation analysis has been performed. In this analysis correlation plots on the data of two volume images are compared. The intensity of a voxel in one image (x) is set in the plot according to the intensity of the same voxel found in the other image (y). Therefore, two images which correlate perfectly can be identified by showing a voxel intensity correlation in the 2D plot with a slope of $m = 1$ through the origin of ordinates.

2.3 Motion Quantification

For a quantification of patient motion, methods are required to obtain and prepare the motion data of a motion tracking system before it can be used for a motion compensation. This includes, for example, methods for a conversion of the motion data into the respective coordinate system of the PET scanner as well as the development of an appropriate patient motion target.

2.3.1 Cross-Calibration

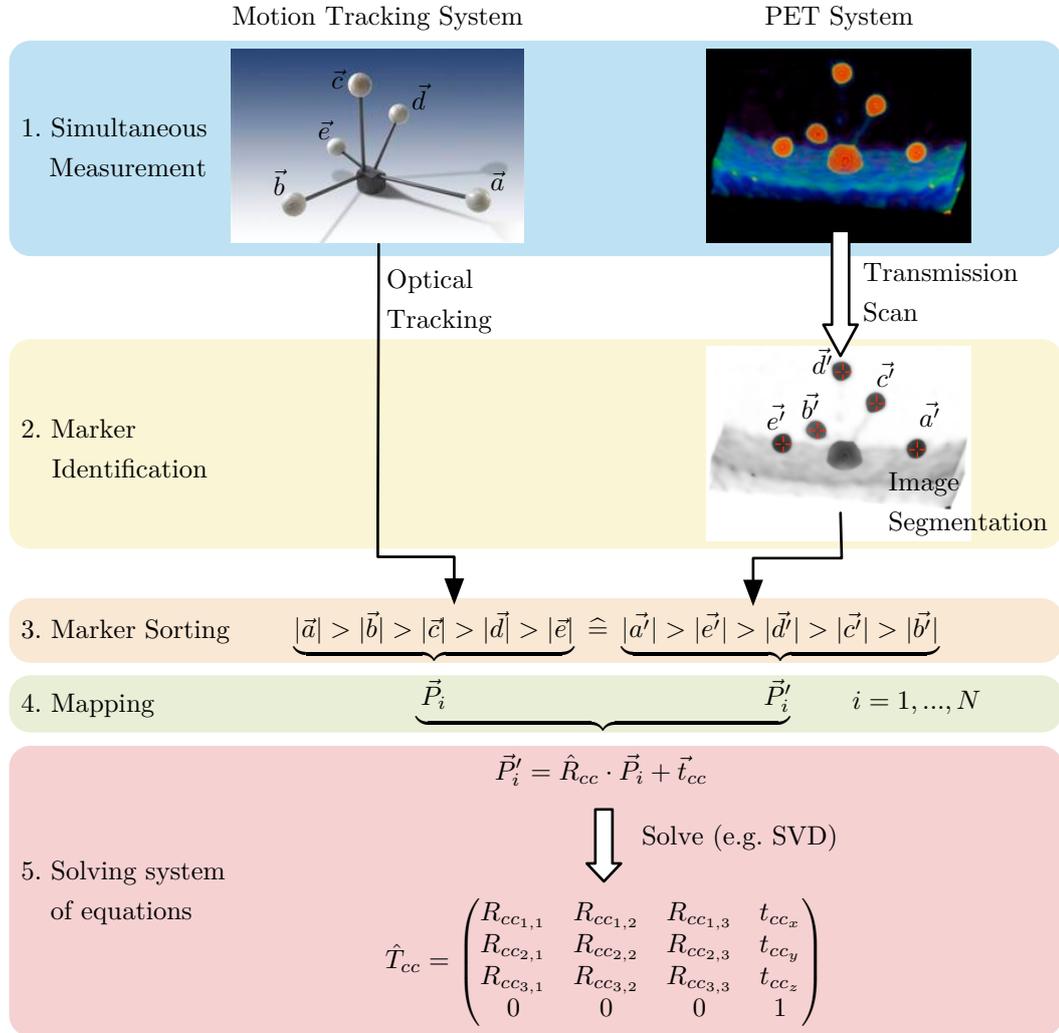
The motion tracking system, as a self-contained system, uses its own coordinate system. This obviously implies that if motion data need to be interpreted in the PET scanner's coordinate system, a transformation between those two coordinate systems will be required. This *cross-calibration* consists of the determination of a transformation matrix which converts the motion data from the tracking system into the coordinate system of the PET scanner. To formalise this conversion and directly combine it with the motion information supplied by the tracking system, a point \vec{P} within the coordinate system of the tracking system is transformed into the moved point \vec{Q} within the coordinate system of the PET scanner via

$$\begin{bmatrix} \vec{Q} \\ 1 \end{bmatrix} = \begin{bmatrix} \hat{R}_{cc} \cdot \hat{R}_{mc} \cdot \hat{R}_{cc}^{-1} & -(\hat{R}_{cc} \cdot \hat{R}_{mc} \cdot \hat{R}_{cc}^{-1}) \cdot \vec{t}_{cc} + \hat{R}_{cc} \cdot \vec{t}_{mc} + \vec{t}_{cc} \\ 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} \vec{P} \\ 1 \end{bmatrix}, \quad (2.1)$$

where \hat{R}_{cc} is the three-dimensional rotation matrix and \vec{t}_{cc} the three-dimensional translation vector of the homogeneous cross-calibration matrix \hat{T}_{cc} (cf. figure 2.14) and $\hat{R}_{mc}, \vec{t}_{mc}$ the respective rotation and translation parameters of the motion data obtained from the external motion tracking system [31].

Although the motion defined by $\hat{R}_{mc}, \vec{t}_{mc}$ is known due to the sampling by the tracking system, a method to calculate the parameters ($\hat{R}_{cc}, \vec{t}_{cc}$) of the cross-calibration matrix (\hat{T}_{cc}) is required. To determine these parameters, we have developed a method to match the position of a set of markers obtained by simultaneous measurement in both, the tracking system and PET system. For this process, a minimum of four spherical motion markers are placed within the FOV of the PET scanner and the tracking system. While the position of each marker within the tracking coordinate system is directly obtained from the recorded motion data, the localisation within the PET coordinate system requires an additional method to make the markers visible within a PET image. To achieve this visualisation, the markers have to be normally composed of a γ -ray-emitting material (e.g. ^{68}Ge). For this material, an emission scan can measure the radioactive distribution and allows to derive the position of each marker. Another possibility is to use markers which are composed of a material which has an attenuation high enough to make them visible by a transmission scan. As the first method requires additional considerations in terms of radiation protection regulations, we concentrated on the transmission-based solution and used markers composed of aluminum.

Segmentation algorithms, as described in [105], have been used to automatically identify and calculate the centre of each of the spherical markers within the reconstructed PET transmission image. During this segmentation three-dimensional ROIs are delineated by means of applying an empirically determined threshold on the intensities in the volume data. Furthermore, the

**Figure 2.14:**

Cross-calibration method illustrated: Due to the different coordinate systems, a *cross-calibration* transformation matrix \hat{T}_{cc} has to be calculated. In several consecutive steps, the coordinates of a minimum of four markers are measured simultaneously with the motion tracking system and by extracting the marker locations from a transmission image. The resulting coordinates are sorted and mapped according to their Euclidean length and then fitted into a system of equations. Afterwards this system of equations is solved (e.g. via a *Singular Value Decomposition* (SVD) algorithm [104]) so that the rotation \hat{R}_{cc} and translation parameters \vec{t}_{cc} are extracted.

activity weighted average of all voxels in each ROI is used for determination of the centre of gravity in such a segmented ROI. This transformation yields a representation of all voxels in the ROI whose radial coordinates (and corresponding intensities) define points on the average radial profile. The known radius of the markers, and therewith the known sphere volume, are provided as an *a priori* information to filter ROIs which are not related to the tracking markers. This helps our algorithms to identify the ROIs of the markers in the volume data automatically by masking out found volumes not matching the reference volume.

After successful identification of the markers in each of the coordinate systems, each marker position in one coordinate system has to be associated with the respective position of the same marker within the other coordinate system. The placement of each marker at an individual distance to each of the other markers allows for deducing a criterion for this marker assignment: by sorting all calculated distances of the marker locations according to their length, the identified marker positions in each coordinate system are mapped to the sorted distances in the other coordinate system. This allows for performing a one-to-one mapping of the found markers in both coordinate systems. To calculate final the cross-calibration parameters $(\hat{R}_{cc}, \vec{t}_{cc})$, the location mapping of each marker is fitted into a system of equations like

$$\begin{aligned}\vec{P}'_1 &= \hat{R}_{cc} \cdot \vec{P}_1 + \vec{t}_{cc} \\ \vec{P}'_2 &= \hat{R}_{cc} \cdot \vec{P}_2 + \vec{t}_{cc} \\ \vec{P}'_3 &= \hat{R}_{cc} \cdot \vec{P}_3 + \vec{t}_{cc} \\ \vec{P}'_4 &= \hat{R}_{cc} \cdot \vec{P}_4 + \vec{t}_{cc}\end{aligned}\tag{2.2}$$

where \vec{P}_i are the coordinates (x, y, z) of marker i directly obtained from the motion tracking data and \vec{P}'_i the extracted location parameters of a transmission scan performed in the PET coordinate system. The rotation matrix \hat{R}_{cc} and the translation parameters \vec{t}_{cc} are derived by solving this system of equations via algebraic methods such as *Singular Value Decomposition* (SVD) [104]. Finally, the cross-calibration parameters $\hat{R}_{cc}, \vec{t}_{cc}$ are put into a homogeneous transformation matrix \hat{T}_{cc} which is used in any future mapping of coordinates between the two systems.

For obtaining the position information of the sampled tracking markers, a software package (**trkCalib**) has been developed. This software implements the algorithms to calculate the cross-calibration matrix out of position information retrieved from both measurement systems (PET system and motion tracking system) in parallel. The software permits to load a motion tracking data file (***.trk**) in addition to a standard ECAT transmission image. For this process, a reconstructed image of the transmission scan of a cross-calibration measurement is used. When this image is loaded, **trkCalib** analyses the image data and automatically identifies the centre of the spherical markers via the aforementioned image segmentation algorithms. After the calculation of the cross-calibration matrix, the parameters of this matrix are saved into a transformation matrix file (***.tfm**), which is afterwards used during a motion analysis. In addition, to allow for an intuitive use of the calibration software, a graphical user interface has been developed as well. This interface allows for loading the motion tracking data and the attenuation image data via graphical elements. After processing the image segmentation, it shows the identified

centres of the markers in the loaded PET image. A screenshot of this graphical user interface with example data is shown in figure 3.3.

Evaluation Protocol

To evaluate the workings of our cross-calibration software and to verify if the cross-calibration method is capable of calibrating the two coordinate systems adequately, we performed ten test acquisitions with the PET scanner and the motion tracking system. Various analyses were done in two separate test phases to identify problems that might exist:

1. In the first test phase, the stability of the cross-calibration method in general as well as the correct operation of the `trkCalib` software has been evaluated. For this purpose, the recorded motion tracking data of five cross-calibration measurements, as well as the image data of the performed transmission scans have been loaded in `trkCalib`. The number of drop-outs and errors during the transmission scan and the motion tracking was counted and taken as criterion for the stability of the method in general. In addition, the correctness of the matrix calculation has been evaluated by verifying that the calculated cross-calibration matrix (T_{cc}) is a valid transformation matrix. Therefore, the determinant of the cross-calibration matrix $\det(T_{cc})$ has been calculated for every test calibration. If the calculated determinant differs from one, the cross-calibration calculation is classified as a failure – a sign that the calculated matrix does not conserve the orientation (right-handed vs. left-handed) when used.
2. In the second phase, another five PET acquisitions have been performed with a calculated cross-calibration matrix. In these list-mode-based emission scans lasting 10 minutes, a radioactive [^{68}Ge] point source with 1 mm diameter and a radioactivity of 6.1 MBq has been brought into the FOV of the PET scanner. While keeping the point source at a fixed position in the first half of each acquisition, it has been moved manually to arbitrary locations during the last half. In addition, it has been motion tracked by attaching it to a standard tracking target via adhesive tapes. As a result of the constant motion the emission data showed the usual blurring after image reconstruction. To verify that the cross-calibration method can calibrate the two coordinate systems sufficiently, two separate event-driven motion compensations have been performed for the last half of each list-mode data. In one of the motion compensations the emission data have been corrected with the previously calculated cross-calibration matrix. For the other compensation, however, the data have been corrected for motion with an identity matrix as cross-calibration, thus with a ‘disabled’ calibration between the two coordinate systems. By comparing the resulting motion compensated images to the image of the first half of each acquisition, the effect of a valid cross-calibration becomes visible. Therefore, we have compared the images and rated them according to their concordance to the point source images at rest. This can be seen as criterion for the accuracy of the cross-calibration matrices.

2.3.2 Motion Target

For an adequate correction of patient motion, it is essential that the motion tracking mechanism is able to report the movement of a patient not only with a high resolution, but also with a high certainty. The problem of an accurate external measurement of motion is mostly subject to the way a tracking system is registering the motion. While the method of motion tracking varies widely (cf. section 1.2.2), most of the motion tracking solutions have in common that a specific tracking target is required. This individual target is monitored by the tracking system and motion information is calculated according to a constant observation of that target. This, of course, implies that the target is able to convey the motion of the monitored anatomic area on the patient sufficiently. A central problem is therefore to ensure a fixed relation between the actual patient motion and the ability of the motion target to ‘transfer’ this motion to the tracking system – something we refer to as *motion transferability*.

The fixation of the motion target on the patient is therefore a highly critical point. As a ‘foreign’ body, the motion target has to be on the one hand properly attached to the monitored area of the patient so that it transfers the occurring motion sufficiently. On the other hand, the motion target has to be suitable for daily use, especially within a clinical environment. For the latter, it is also important that the target fits comfortably and is not experienced as an obstacle during the PET examination, which might otherwise result in an additional unrest motion of the patient.

For this study, where only the tracking of motion of a human’s head is examined, certain head motion target solutions have been partly developed and evaluated. One solution utilises common ski goggles by attaching the motion target on the top of the frame of the goggles. Another tracking body is, in contrast, directly attachable to the head surface of a patient via adhesive tape. As a main goal of the evaluation, a qualitative and quantitative statement on the suitability of the individual motion targets and their fixations on the patient’s head should be given.

Maxilla-based Reference Target

To evaluate the motion transferability and general stability of the motion targets objectively, a reference target has been developed. This target is directly attached to the upper jaw (*Maxilla*) of a test person and has not been developed for clinical use (cf. figure 2.15). Thanks to a previously performed teeth moulding procedure, the reference target is attachable to the Maxilla in a reproducible way. The teeth moulding target is worn during the test acquisitions so that the actual head motion is directly reported to the motion tracking system.

The way of fixation of this motion target represents, in principle, a *gold standard* in terms of motion transferability. This is true because the upper jaw is directly connected to the skull of a human head and thus reports any motion of the entire head immediately. Therefore, the motion targets that are to be evaluated are directly compared to motion registered simultaneously by this reference target.



Figure 2.15:

Teeth moulding-based reference motion target: As a reference method for measuring the patient's head motion, a separate reference motion target has been developed. This target consists of a teeth moulding form which is attached to a motion tracking target. Due to the direct measurement of the motion from the Maxilla of the test patient, this target is used as a *gold standard* to compare the motion transferability of existing targets, but also to support the development of an optimised head motion tracking target.

Ski Goggles-based Head Target

As one of the earlier approaches for an adequate patient target, a target model had been developed in previous studies [31]. This target is based on ski goggles to which a motion tracking body is attached at the upper frame of the goggles. The used motion tracking body consists of a standard *tree target* with five reflecting markers located at positions where each marker has a unique distance to the other markers (cf. figure 2.16).



Figure 2.16:

Ski Goggles-based Head Target: In the course of a previous study [31], a ski goggles-based head motion target had been developed. It consists of a ski goggles frame to which a standard *tree target* for optical motion tracking is attached. Using attached elastic straps with velcro fasteners at each end, the target can easily be fastened and released.

The goggles frame consists of soft plastic and has foamed material around the inner edges of the frame which are attached to the patient's face during an examination. At the front of the frame, a large transparent screen exists, which itself is composed of a solid plastics material, providing the whole goggles its stability. The pair of goggles themselves are attached to the patient's head via elastic straps on both sides of the frame. To provide for an easy fastening, the elastic straps have been modified to be closed via velcro fasteners at each end. In order to avoid a high attenuation factor due to the used materials, the tree target at the top of the goggles, including the spherical markers, is completely composed of *polyvinyl chloride* (PVC).

Os frontale-based Head Target

As part of a co-operation with the Forschungszentrum Jülich [106], a custom head motion target has been developed and evaluated. It is composed of a transparent plastics material with a possibility to place a maximum of four optical markers on a planar platform at the top of the target. The target itself is fixed on the skin of a patient within the area of the *Os frontale*. Suction cups are used to attach it to the skin. In addition, adhesive tape is used at the fringes of the caps to secure them on the head surface (cf. figure 2.17).

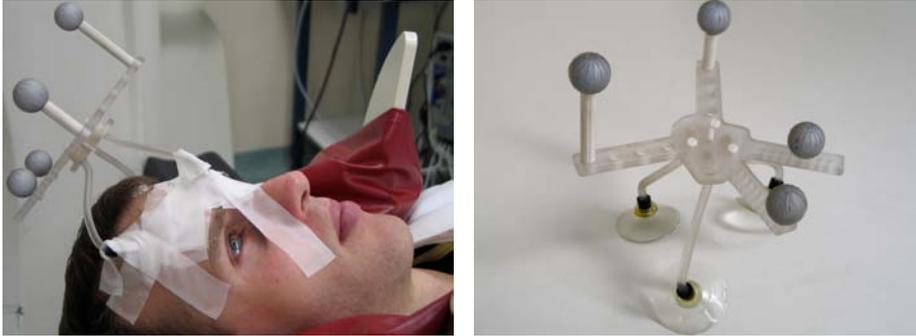


Figure 2.17:

Os frontale-based head target: Image of a custom head motion target developed in co-operation with the Forschungszentrum Jülich [107]. Through a fixation of the target on the head surface within the area of the *Os frontale*, the motion transferability ought to be improved compared to the ski goggles-based motion target approach.

While the target was already evaluated at the Forschungszentrum Jülich in previous experimental studies [107], its motion transferability within a clinical environment ought to be evaluated in this thesis. In addition, its compatibility with the ARTtrack motion tracking system has been analysed and evaluated.

Safety Goggles-based Head Target

To provide for another alternate head motion target, a third possibility for an adequate head motion target has been developed. This target uses a safety goggles frame with a standard tree target attached at the upper side. One of the main differences between this safety goggles-based target and the ski goggles-based one is, that the target comes with a larger and completely transparent screen at the front (cf. figure 2.18), which is composed of solid plastics.

The inner part of the frame, which is in contact with the patient skin, is composed of flexible plastics material. Another difference, compared to the ski goggles target, is the large skin overlap of the frame. Because of that, the general seating of the goggles frame on the patient's skin is improved. Similar to the head targets already discussed, the same tree target with five optical markers is attached to the upper side of the target. When worn, the orientation of the tree target on the safety goggles however is slightly different than that on the ski goggles. Furthermore, the safety goggles-based target is fixed on the patient's head via a single elastic strap that is to be put around the head of the patient.



Figure 2.18:

Safety Goggles-based Head Target: As an optimised motion target, a new target which is based on a safety goggles frame has been developed. The target is composed of fully transparent plastics material where the elastic frame overlaps on each side of the target. This assures a good attachment onto the patient's skull and thus a good motion transferability in general.

Evaluation Protocol

For a proper evaluation of the aforementioned head motion targets, a separate evaluation protocol has been developed (cf. figure 2.19). This protocol is used during several motion tracking acquisitions. During those consecutive acquisitions, a volunteer is asked to perform different types of predefined movements while wearing one of the head motion targets together with the Maxilla fixed reference target. These motions are executed in a specific direction so that potential differences in the fixation of each motion target are easily identified. This will indicate which target is best suited for transferring the head motion to the target of the motion tracking device. The motion parameters of both targets – the forehead fixed motion target and the Maxilla fixed motion target – are compared and the distance between the two targets is calculated over time. In an ideal situation, the distance between those two targets is always constant.

In addition to the evaluation of the motion transferability of the targets, the general patient's compliance is a criterion which has been taken into consideration. After each acquisition, the volunteer has been interviewed if the head motion target felt comfortable during the acquisition or if there were any areas in the fixation which caused a feeling of discomfort during the acquisition. Compliance scores between 1 (Excellent) and 5 (Fail) have been assigned for this purpose. The effort needed to attach and remove the motion targets has been evaluated as well. Similar to the patient's compliance, this effort has been rated between one and five. Both evaluation criteria are important for the usability in a clinical environment.

2.3.3 Motion Analysis

Drawing conclusions from the raw motion data on the actual magnitude of motion can be problematical. Motion data are, for example, usually provided in the form of three translations along, and three rotations around, the coordinate axes. These raw parameters (the *six degrees of freedom*), however, are not suited for a direct and intuitive assessment of the magnitude of three-dimensional motion and its effect on PET data. However, the magnitude of the translations can be deduced directly from the translation parameters of the motion data. But the three rotation parameters (the rotation angles η , θ , ϕ) and the non-commutative relationship between

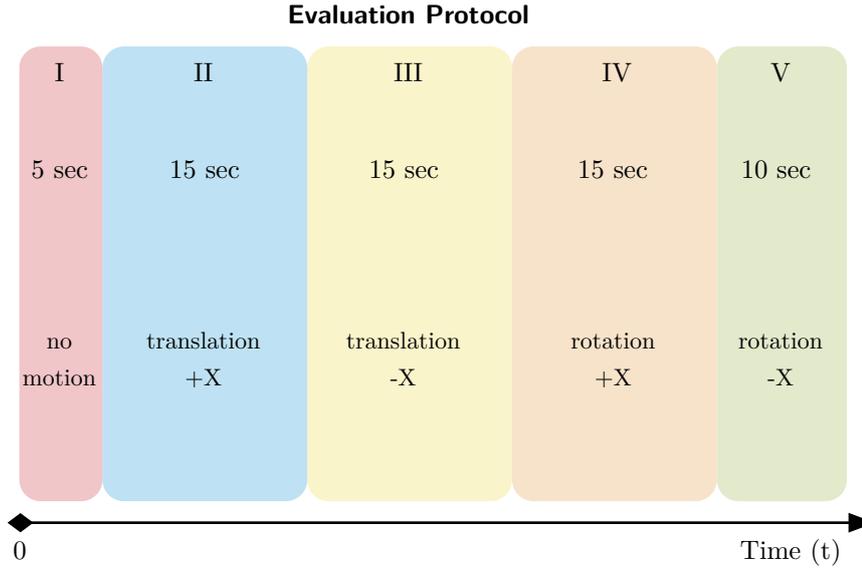


Figure 2.19:

Illustration of the motion target evaluation protocol: In five consecutive phases each motion target is evaluated according to its motion transferability. In phase I a five seconds interval with no motion is used to verify the correct stability of the motion tracking in general. In the beginning of phase II to V the volunteer is asked to perform a single motion in a specific direction.

the application of a translation and a rotation make it hard to deduce an overall magnitude statement from just those raw motion parameters. The actual magnitude of rotation angles is, for example, dependent on the given centre of rotation. Therefore, it is necessary to execute the spatial transformations defined by the six motion parameters for each point within a specific ROI. This permits conclusions about the actual influence of the motion at that very coordinate. In addition to the problem concerning the interpretation of motion tracking data in general, suitable motion threshold values have to be specified – not only to identify ‘significant’ motion, but also to reduce the amount of relevant motion data. In doing so, the required computation time for motion compensation methods can be minimised.

Therefore, procedures for an automated analysis of motion data addressing the aforementioned problems have been developed. These procedures make use of the statistics toolkit ‘R’ [108] and directly process the raw motion data. This toolkit facilitates statistical analyses on large data sets by providing a functional programming language environment.

Data Processing and Data Filtering

To smooth background noise instabilities which are generally present in the sampling of motion data, we filter the data with a *simple moving average* (SMA) algorithm [109] which is defined by

$$SMA_t(x, n) = \frac{1}{n} \sum_{i=0}^{n-1} x_{t-i} = \frac{x_t + x_{t-1} + x_{t-2} + \cdots + x_{t-n+1}}{n}, \quad (2.3)$$

where $SMA_t(x, n)$ is the moving average of n consecutive values of the motion parameter data set x up to and including a given time t . The six independent parameters of the motion data (rotations, translations) are used independently for the input (x) of the moving average. After

the parameter smoothing, the whole data set (x_t) is searched for equal consecutive motion parameters so that the data set is reduced by elimination of the smoothed duplicates. Furthermore, the rotation matrices are recalculated from the smoothed rotation parameters so that a valid transformation matrix in the form of a 3×3 matrix is available (i.e. with a determinant of one).

Position-relative Motion Analysis

A method has been developed to easily quantify the magnitude of patient motion for a specific coordinate within a target ROI. It allows to specify any three-dimensional coordinate within the FOV of the PET scanner. For this particular coordinate, statistical parameters on the motion magnitude over time are calculated and graphically prepared. By spatially transforming the specified coordinate vector (x,y,z) with each motion transformation matrix supplied by the tracking system, an array with distances of the coordinate point to its original position is calculated over time. Because of this position-relative contemplation of motion data, a direct assessment of the magnitude of motion for a specified target region is possible. Due to the utilisation of the rotation parameters, for example, and thus by calculation of the distance the dependency of the rotation parameters on the actual centre of rotation is clarified. By assessment of the calculated distances rather than by assessing the raw parameters, a direct statement on the motion magnitude is derived.

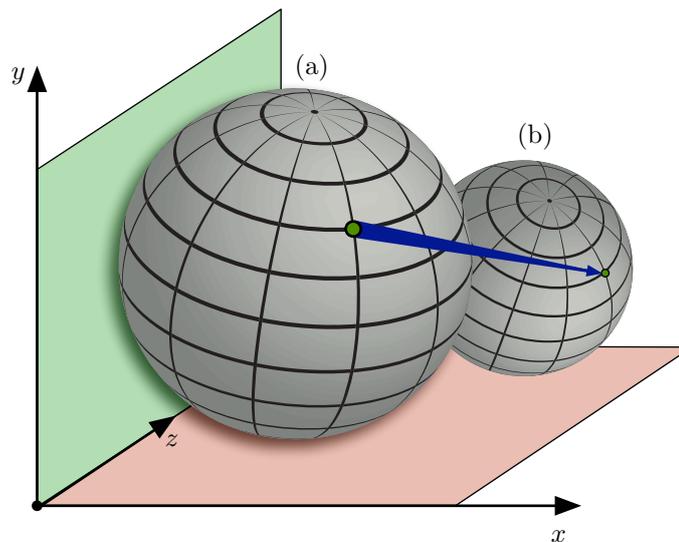


Figure 2.20:

Sphere surface with virtual grid: Illustration of the position-relative motion analysis method. A grid is placed on a virtual sphere within the coordinate system of the PET scanner (a). All intersecting points on the grid (here shown on one example intersection point in green) are transformed with the motion information of the tracking device (b). To assess the magnitude of the motion in relation to the position in the FOV, the distances of each grid point relative to their origin before the motion transformation are calculated and statistically analysed.

For the goals of this thesis, we have enhanced this technique to be applicable to the analysis of head motion. In the literature, the mean perimeter of a human head is documented to vary between ≈ 54 cm for female adults and ≈ 56 cm for male adults, respectively [94, 95]. To calculate the ‘real’ magnitude of the motion from the registered patient motion, we transfer the motion parameters onto the surface of a virtual sphere with a perimeter comparable to that of a ‘mean’

human head. For our application, and to incorporate a safety boundary, we are in fact using a virtual sphere with a diameter of 20 cm, which corresponds to perimeter of ≈ 62 cm. By spatially transforming each intersection point of a virtual grid placed on the imaginary patient head with the motion tracking data, the distance of each grid intersection point from its original position is calculated (cf. figure 2.20).

Therefore, by analysing all calculated distances of the intersection points of the grid, the coordinates on the sphere – and thus on the imaginary patient head – with the largest distance and smallest distance are identified. This helps to identify motion intensive and non-motion intensive areas on the sphere and thus on the imaginary patient head. Similar to that, the magnitude of motion for any coordinate within the whole FOV of the PET scanner can be calculated. Together with information on the coordinate with the highest and lowest fluctuation during the acquisition, a distance frequency analysis for the ‘worst’ and ‘best’ point on the sphere is performed. According to this analysis, not only the type of motion but also the influence of motion is identified, thus providing a better prediction of the expected image degradation. In addition, this analysis should help to identify and describe the motion according to its nature, thus classifying it as a *one-time motion* and/or *continuous drifting motion*, as well as pointing out the areas where the motion had most influence on the PET data.

When applying our position-relative motion analysis, the motion data is also automatically transformed into the PET coordinate system. This is performed by applying the aforementioned cross-calibration (cf. section 2.3.1) to every motion data received from the tracking device. Furthermore, to filter out ‘minor’ motion data during the motion analysis a fixed threshold of 1 mm to 3 mm is applied. This empirically determined threshold has been chosen to be slightly lower than the known resolution of the PET scanner – which is between 5 mm to 6 mm for the ECAT Exact HR⁺. Based on this value only motion which results in a maximum distance on the imaginary patient head higher than the threshold (within a 10 % to 90 % confidence interval) is considered to be ‘significant’. To provide this type of motion filtering to a motion compensation method the transformations which have been found to imply a motion to be ‘significant’ are saved to a data file only. This file then contains only the significant motion data with already cross-calibrated transformations. As a consequence of this method the motion compensation is applied only for relevant motion data, which in fact potentially reduces the amount of time required for the compensation.

Report Document

For preparing and using the aforementioned motion analysis in a clinical environment, the statistical data have been graphically prepared. A software package (`analyze_trk`) has been developed which is using the ‘R’ statistical programming language. In addition to applying the position-relative motion analysis and outputting the ‘significant’ motion data for a correction, this software package contain routines to graphically prepare the motion analysis in a clinical report document. Through the generation of this report document, the physician should be provided with a tool for better assessment of the motion of the patient and its impact on PET data. In addition, it should help to decide if a motion compensation is really required and if the additional time for this correction is justified.

The report document contains multiple pages with graphical plots on the motion analysis. An example of this report document can be found in section 3.2.3. For a clinical application of this report document it has been split into three different pages:

Raw Motion Page: To illustrate each single motion parameter, one page of the report document contains graphical plots on each of the six degrees of freedom (3 translations, 3 rotation angles). Apart from the motion data of the patient target, this page also contains motion plots on a special reference target. This target is located at a fixed position within the FOV of the motion tracking cameras. It is included as a quality control, so that with an incorrectly calibrated motion tracking system, motion would appear in the data of the reference target. Therefore, in case of a correctly calibrated and configured tracking system, the translation and rotation parameters of the reference target ought to be zero throughout the whole motion acquisition. An example of this report page can be found in the result section of this thesis (cf. figure 3.9).

Sphere Analysis Page: On a second page the data on the imaginary patient head analysis are shown (cf. figure 3.10). In three main graphical plots the time dependent distances of each grid point on the sphere are illustrated. For the first plot the distance of each grid point from its original position at acquisition start is calculated and plotted over time. As a result a grey area with the calculated distances of all grid points is generated. One line along the abscissa in the grey area corresponds to a single intersection point on the sphere. In addition, red and blue lines mark the ‘worst’ and ‘best’ point on the sphere, corresponding to the sphere coordinates with the highest and lowest motion identified. Note, however, that these points do not necessarily correspond to the ‘best’ and ‘worst’ point on the sphere at any time. Moreover, these points are the coordinates on the sphere which were found to carry the minima and maxima of the calculated distances throughout the whole motion analysis. Furthermore, as the position-relative motion analysis allows for defining additional coordinates within the FOV of the PET scanner we output data in the range plot for two additional coordinates. For these coordinates we have chosen the ‘mean’ coordinates in the FOV where usually the left and right *striatum* is to be found. For the determination of these ‘mean’ coordinates, we have analysed a set of 20 PET images and manually identified the coordinates at which the left and right striatum could be found with the patient fixation used at our facility. To allow for a better interpretation of the plotted motion, in relation to the acquisition time, the splitting of the acquisition into dynamic frames is also illustrated via grey vertical dotted lines in the range plot. Furthermore, the times at which a motion was found to be ‘significant’ according to the threshold are displayed via green vertical lines and small green triangle signs on the abscissa. For illustrating the information on the ‘best’ and ‘worst’ point of the sphere movement analysis, histograms of the frequencies of the distances of each of the points has been included as well. These frequency plots correspond to the projection of the respective ‘best’ and ‘worst’ lines from the abscissa onto the ordinate of the range plot.

Furthermore, to allow for an intuitive inspection of the motion intensive areas on the imaginary patient head, an additional illustration with a three-dimensional output has

been developed. This three-dimensional view of the virtual sphere is, however, an interactive graphic which can be used during an interactive analysis with our software package. It contains routines which allow for an interactive inspection of the sphere. By moving the mouse of the computer system a sphere view can be directly rotated and zoomed in real-time. Even if the three-dimensional plots are not part of the report document, the information shown in these interactive 3D graphics is, however, also included in the report page via a tiled image map. Both, the interactive 3D sphere graphics and the 2D image map in the report page illustrate the fluctuation of the distances of each grid point over time. In fact, the standard deviation of the distances of each grid point on the sphere is calculated and illustrated. In a second image map, the mean distance of each grid point from its origin at acquisition start is illustrated. Both image maps carry anatomic labels to allow for a better identification of the motion intensive areas based on the human head anatomy (cf. section 1.3). An example of these sphere plots and image maps is shown in figure 3.11.

Motion Summary Page: On the final report page, the information of all previous pages are summarised in a plain text document. Together with the values on the raw motion of the reference and head motion target, the distance values on the ‘best’ and ‘worst’ point on the sphere are presented. In addition, the motion values on the user selectable coordinates within the FOV (e.g. the ‘mean’ striatum coordinates), are shown as well.

The complete motion analysis implies a large amount of statistical values and plots. Therefore, a final motion score is part of the summary page. This score divides the motion magnitude into five distinct values, as shown in table 2.2.

Table 2.2:

Motion Analysis Scores: For an intuitive assessment of the motion analysis empirical motion scores are assigned in the motion summary page of the report document. Based on the maximum within the 10 % to 90 % percentile of the ‘worst’ point distances found on the imaginary patient head, the score is assigned to different classes.

Score	Description	Maximum Distance (in 10 %–90 % percentile)
1	no motion	0.0 mm – 0.5 mm
2	low motion	0.5 mm – 3.0 mm
3	moderate motion	3.0 mm – 5.0 mm
4	significant motion	5.0 mm – 7.0 mm
5	high motion	> 7.0 mm

These score values have been empirically assigned to be dependant on the maximum value found in the 10 % to 90 % percentile of the distances for the ‘worst’ point on the sphere. An example of this page is shown in figure 3.12.

2.4 Event-Driven Motion Compensation

Certain aforementioned practical implications of an event-driven motion compensation restrict its use in clinical PET. The difficulty of lost counts in the sampled data due to a transformation

of LORs outside the FOV of the scanner, in combination with the resulting image artefacts, is only one reason for this limitation. The following section outlines the methodology of all the required post-correction methods as parts of an event-based compensation and discusses known shortcomings. Finally, improvements to those methods are presented.

2.4.1 Normalisation Correction

One of the necessary post-correction methods of an event-driven motion compensation takes care of the differences in the detector sensitivities (normalisation) when reassigning detector tuples (*normalisation correction*). As discussed in section 1.1.1, each detector has a different γ -detection sensitivity due to physical differences in the scintillator material and due to differences in the signal processing electronics of the PMTs. During a normalisation scan, correction factors are usually calculated to compensate for those differences.

The reorientation of LORs during an event-driven motion compensation represents an invalidation of this standard normalisation procedure. Due to the spatial transformation of an LOR, resulting in a reassignment of a detector tuple $d_{(i,j)}$ to a different detector pair $d_{(i',j')}$, the normalisation coefficient $\eta_{(i,j)}$, as shown in equation 1.1, needs to be adapted. For this operation, a normalisation correction factor $f_{(i,j,i',j')}$ is calculated for detector tuple (i, j) by

$$f_{(i,j,i',j')} = \frac{\eta_{(i',j')}}{\eta_{(i,j)}}, \quad (2.4)$$

where $\eta_{(i,j)}$ is the normalisation value for the initial detector tuple $d_{(i,j)}$, and $\eta_{(i',j')}$ the normalisation value for the new, motion compensated tuple $d_{(i',j')}$. The resulting normalisation correction factor is calculated and applied to every transformed LOR so that the standard normalisation map can be used during image reconstruction.

To achieve this normalisation correction within the ECAT environment of an HR⁺ scanner, the data format of the normalisation file has been implemented and routines to calculate these normalisation correction factors have been developed. In addition, the normalisation maps used by the HR⁺ also contain factors such as the correction for the dead time of the electronics of the PET scanner. Therefore, procedures for extracting and combining these factors with the normalisation correction factors have been implemented.

Evaluation Protocol

For an evaluation of the implemented normalisation correction methods, a number of ten list-mode-based test acquisitions have been performed. To confirm that the extracted and calculated normalisation coefficients are able to normalise the coincidence data of the test acquisitions sufficiently, the list-mode data has been sorted into standard ECAT sinograms and reconstructed afterwards. In one data set of the tests, the sinograms have been generated by applying the normalisation factors during the image reconstruction, thus performing a standard image reconstruction. For the other data set, the list-mode sorter software has been altered to normalise each LOR directly before sorting it into a sinogram. This generates already normalised sinograms during the list-mode sorting process so that the images do not have to be normalised during an image reconstruction.

By comparison of the two different normalised data sets, the validity of our normalisation routines has been evaluated. On the basis of a calculation and analysis via voxel-based intensity correlation plots the data of the different normalised PET images have then been compared to each other. Following the correlation comparison described in section 2.2.3, the differences in the normalisation correlation have also been quantified by calculating the deviation of the image intensities from a slope of one (identity line).

2.4.2 LOR Discretisation Correction

The previous assumption that an LOR is the connecting line between the centres of two γ -detectors, leads to a problem when spatially reorienting an LOR with the event-based correction algorithm shown in section 1.2.3. The spatial placement of the γ -detectors in a cylindrical arrangement requires certain geometrical specifics to be considered for reorienting LORs.

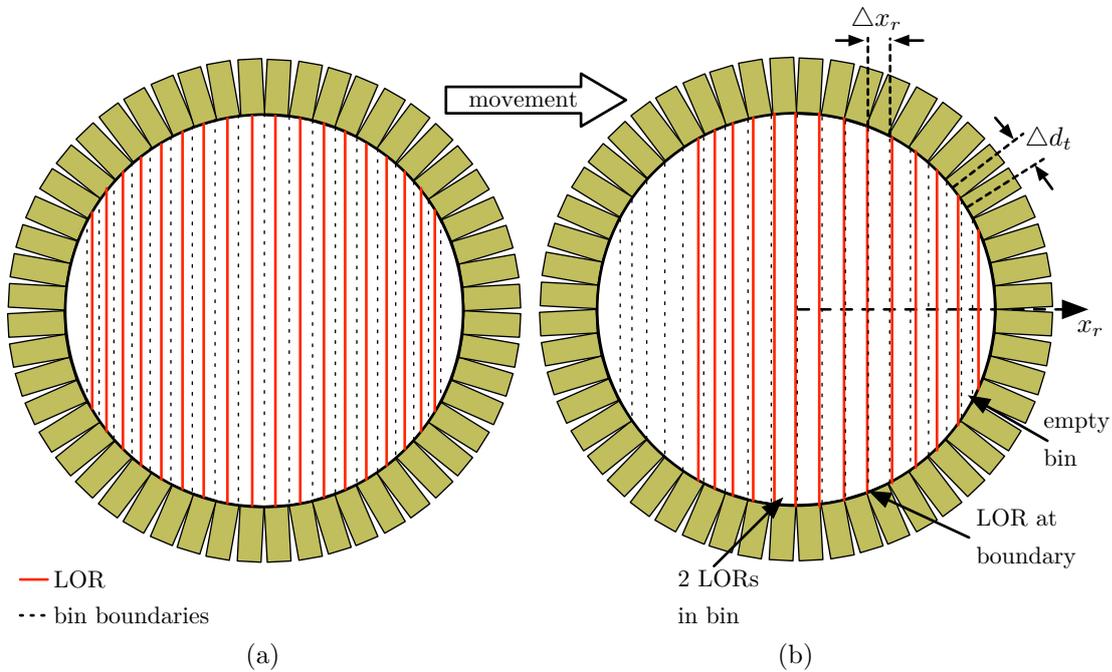


Figure 2.21:

Illustration of the LOR discretisation problem: Part (a) shows the LORs as the connecting lines between the centres of opposite detector tuples. By considering a simple translation in x_r direction (b) it becomes clear, that due to the different transaxial spacing Δx_r , certain LORs are moved exactly between two adjacent bins. But cases where LORs are transformed so that empty bins or 'overfilled' bins occur are also shown.

An LOR may be, for example, spatially transformed so that it is relocated exactly between two adjacent detectors and thus cannot be assigned uniquely to a new detector pair (*uncovered bins*). In addition, the radial placement (x_r) of the detectors results in a different transverse spacing Δx_r between adjacent detectors, as shown in figure 2.21. This spacing is given by

$$\begin{aligned} \Delta x_r &= \Delta d_t \cdot \sqrt{1 - \left(\frac{x_r}{R_d}\right)^2} \\ \Delta d_t &= 2\pi \cdot \frac{R_d}{N_d} \end{aligned} \quad (2.5)$$

where Δd_t is the constant centre-to-centre spacing between detectors in the transverse direction, N_d the number of detectors in one transaxial ring and R_d the radius of the PET ring [59]. Therefore, with increasing x_r , the effective transverse spacing Δx_r of an LOR decreases constantly. Hence, a spatial LOR transformation may result in several source LORs being assigned to the same new detector tuple and thus to the same sinogram bin. Such bins are called *overfilled bins* because the whole portion of a single LOR with a larger Δx_r is assigned to a new LOR with a smaller Δx_r , even if it only covers a portion of the new LOR boundaries after the spatial transformation. In practice, this leads to a non-uniform distribution of counts within the motion compensated PET data and thus to image artefacts.

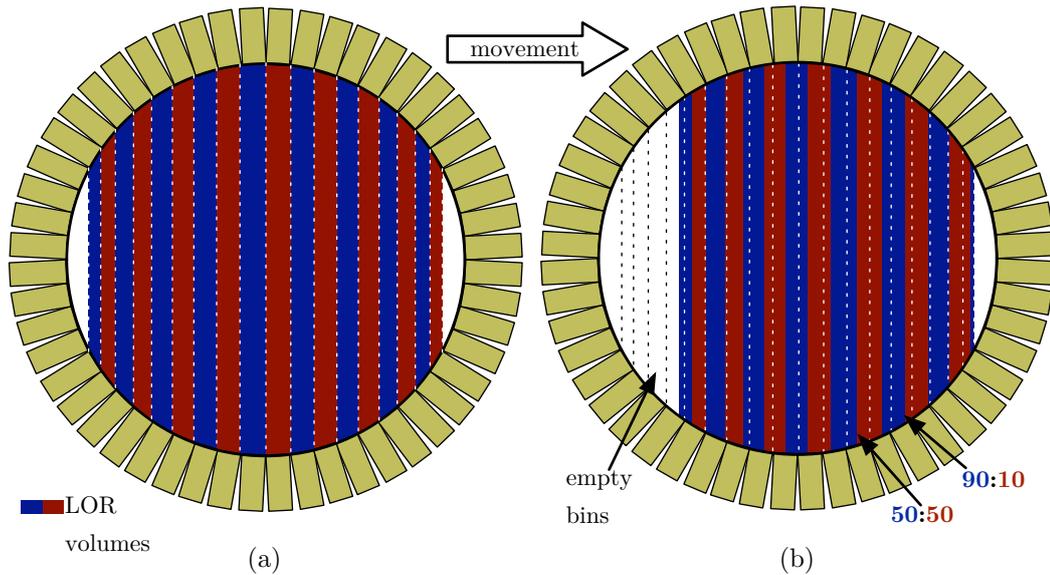


Figure 2.22:

LOR discretisation correction: The LOR discretisation correction considers LORs as three-dimensional volumes confined by the eight corners of the two respective detectors. Here, the method is illustrated for the 2D case (a). After transformation, the discretisation correction identifies the amount of overlap and calculates weighting factors for every LOR (b). This factor is then applied to the target bins which overlaps the transformed LOR.

To compensate for these effects, we consider an LOR to be a volume, which is confined by the planes connecting all four edges of the two opposite detectors – a cuboid with a rectangular cross-section [31]. Due to this discretisation, an LOR transformation cannot lead to uncovered or overfilled bins, except for situations where LORs are transformed outside the FOV. If several of those transformed LOR volumes cover the same sinogram bin, a weighting factor is calculated. It is proportional to the amount of overlap between the transformed volume and the particular bin constraints (cf. figure 2.22).

The calculation of these weighting factors is generally a time intensive task due to the arbitrary orientation of an LOR after spatial transformation, thus resulting in an arbitrary three-dimensional intersection area where the transformed LOR volume hits the detector ring. Therefore, a simplified LOR discretisation correction calculation has been used in the past. This simplification reduces the amount of computations for the calculation of the overlapping area by assuming a constant bin width in axial direction. Thus, only the overlapping 2D area in transverse direction has to be calculated and not the intersections in all three dimensions [15].

In practice, however, this geometric simplification still causes LORs to be mismatched or incorrectly assigned to a new detector pair, especially for motion with rotations along the X-axis (which are quite common). Due to the arbitrary orientation of LORs after their spatial transformation, the intersection area of an LOR with the fixed detector dimensions may cause the axial dimensions of the intersection to be larger than the constant axial bin width used during the simplified algorithm.

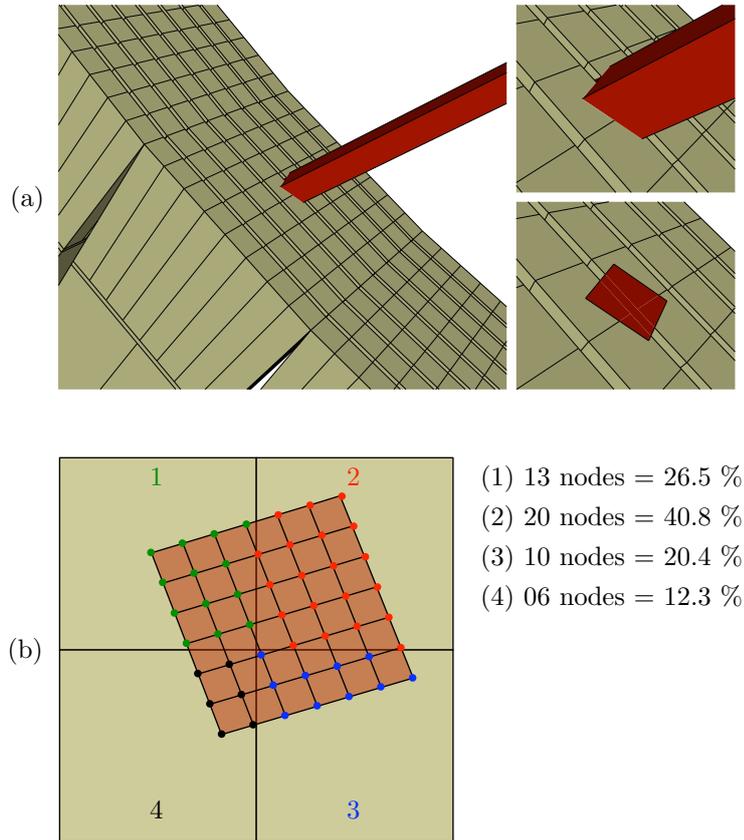


Figure 2.23:

Optimised LOR intersection calculation: Illustration of the LOR volume intersection problem and the optimised method to calculate the amount of overlapping with individual detectors. In the top row (a) the problem with arbitrary overlapping of transformed LOR volumes with the detector surfaces is demonstrated (here with 4 involved detectors). In the bottom row (b), the workings of the optimised algorithm to identify the weighting of the overlaps is shown. An imaginary grid is placed on the overlapping surface and the amount of overlapping is approximated by counting the nodes of a grid that fall into the area of the respective detectors.

To compensate for this problem, a different intersection calculation algorithm has been developed. It calculates the weighting factors by using a full three-dimensional intersection calculation. Instead of transforming the centre coordinate of the LOR endpoints only, all of the eight endpoint coordinates of the confined LOR volume are transformed. After their transformation, the coordinates of the points where the corner marks of the transformed LOR intersect with the detector ring are calculated – similar to the general correction algorithm procedure outlined in figure 1.21. To efficiently identify which of the detectors on a ring are hit by the transformed LOR dimensions, a virtual grid is placed between each of the four transformed corner marks.

Furthermore, an algorithm is used to process on each intersection on the grid: The virtual grid points which are located within the area of a detector are counted and then the weighting factors for a specific LOR are calculated from this information. Thus the proportion of a single LOR volume is distributed over numerous detectors ‘hit’ by the intersections of both end areas of the cuboid LOR volume. This procedure is illustrated in figure 2.23.

Evaluation protocol

An artificial list-mode data stream has been generated to evaluate the correct operation and the accuracy of the optimised LOR discretisation correction. The list-mode data contains LORs of a specific segment of a sinogram only, which, when sorted, end up in a predefined pattern allowing an easy identification of each LOR. This pattern is shown in figure 2.24 on a single projection $p(x_\phi, y_\phi, \phi, \theta)$ with fixed ϕ and θ .

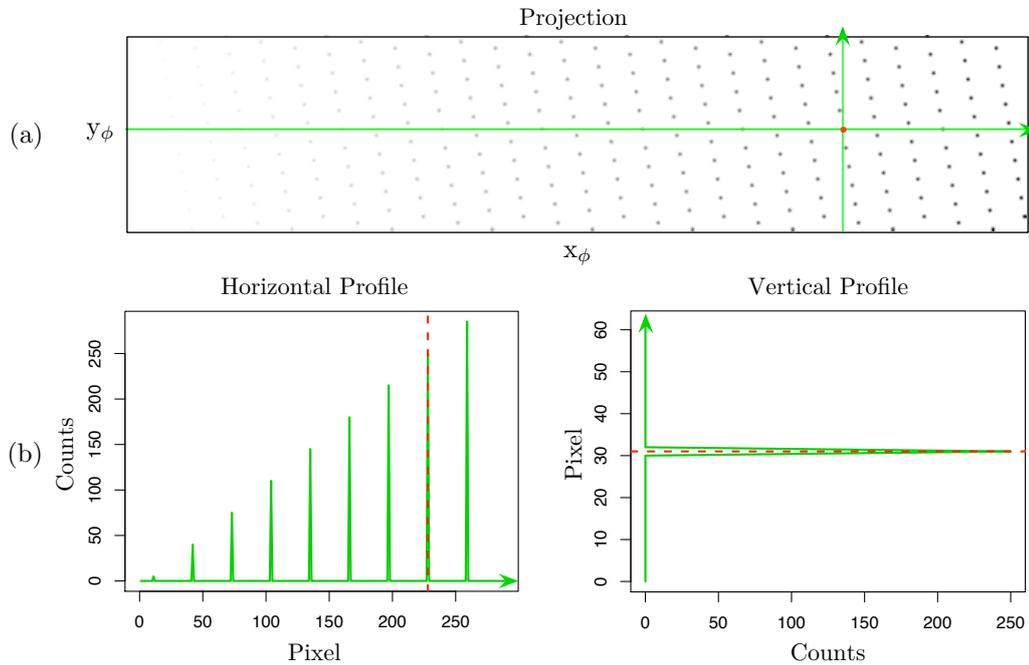


Figure 2.24:

Artificial List Mode Pattern: For the evaluation of the optimised LOR discretisation correction an artificial list-mode data stream has been generated. This data stream contains an LOR pattern in a single sinogram view only. Part (a) shows the projection data $p(x_\phi, y_\phi, \phi, \theta)$ of this view together with the horizontal and vertical profile (b) drawn along the green lines. This illustrates that for every inserted detector combination (the single dots) the number of counts is subsequently increased.

For each artificial detector combination inserted into the list-mode data stream, the number of inserted counts for this LOR is increased by one, which results in an increased amount of counts per LOR. This can also be seen in part (a) of the figure by comparing the left-hand side of the projection image with the right-hand side. At the right-hand side, the LORs with higher count rates end up in the projection with a darker colour, thus more visible.

After having generated the artificial list-mode data stream, it has been spatially moved in the y_ϕ direction of the sinogram space (i.e. the Z-axis of the PET scanner). For this process, the aforementioned event-driven motion compensation has been used with the optimised LOR

discretisation correction enabled. To compare the optimised correction to the previous non-optimised solution, the non-optimised correction has been applied to one additional data set. In both tests (optimised vs. non-optimised), the patterned list-mode data has been moved along the Z-axis so that the spatially transformed LOR volumes should overlap with neighbouring detectors by $1.5\times$ of the voxel width.

To verify that the new discretisation correction method is able to optimise the motion compensation sufficiently, the sinograms of the reoriented list-mode data have been analysed. By drawing profiles along y_ϕ right over the sinogram bins with the numbered LORs, the weighting of the optimised correction has been visualised. The distribution of the counts has been visualised in a separate plot with the counts found on a drawn profile line. By comparison of the count distribution of the optimised and non-optimised data the better assignment of counts due to the percentage of known LOR volume overlap has been verified.

2.4.3 Out-of-FOV Correction

If LORs are moved outside the FOV during spatial transformations or if transformed LORs cannot be assigned to valid detector tuples even after an LOR discretisation correction, the event of that particular LOR is discarded (cf. figure 2.25). Hence, such discarded Out-of-FOV (OFOV) events add to the data consistency problem, discussed in section 1.1.1. While slightly lower statistics might not be necessarily critical, the existence of OFOV events, however, lead to a non-uniformity in the intensities of the image data and thus to image artefacts after the image reconstruction (cf. figure 2.26).

The spatial transformation of coincidences causes LORs to be discarded and results in the existence of *empty bins* at the opposite direction of their occurrence within the FOV (cf. figure 2.22). Such *empty bins* lead to underestimated areas in the sinogram, which contribute to the non-uniformity problem already mentioned. Furthermore, even though OFOV events mostly occur at the outer limits of the FOV they also contribute to an underestimation of tracer quantities near the centre. This is caused by the non-local property character of image reconstruction algorithms in general (cf. section 1.1.2). Thus, empty bins are the reason for a slightly reduced resolution in the centre of images if an event-driven motion compensation is applied (cf. figure 2.26).

For a correction of such underestimated areas in motion compensated PET data, an *Out-of-FOV correction* can be applied [15]. Based on the assumption that an LOR may be spatially transformed outside the FOV for a limited time only, this post-correction method tries to compensate for the underestimation by introducing an additional scaling factor for each sinogram bin. Despite the working of the motion compensation method, the introduced scaling factor $f_{OFOV(i,j)}$ have to be calculated for *every* detector combination by

$$f_{OFOV(i,j)} = \frac{t}{t - t_{OFOV(i,j)}} \text{ with } t_{OFOV(i,j)} < t, \quad (2.6)$$

where t is the duration of the whole acquisition and $t_{OFOV(i,j)}$ is the accumulated time the corresponding LOR, confined by detector tuple (i, j) , was found outside the FOV. Thus, the underestimated areas in the sinogram are theoretically restored by multiplication of the calculated

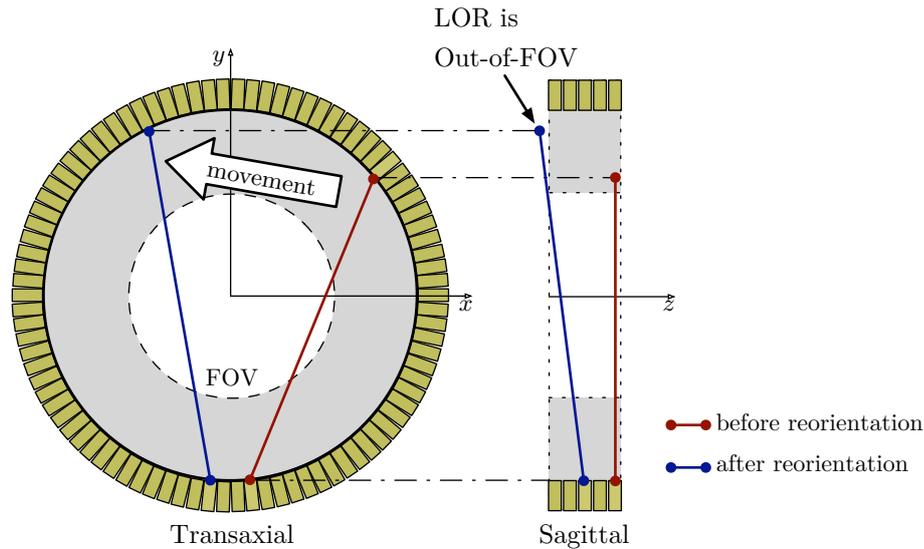


Figure 2.25:

Out-of-FOV problem illustrated: During an event-driven motion compensation an LOR may be transformed outside the available FOV of the scanner geometry. This results in *lost counts* as a direct consequence of this Out-of-FOV problem.

(LOR-wise) OFOV correction factor with the registered counts in the corresponding sinogram bins.

This is, of course, only true for sinogram bins which at the time those factors are calculated/applied contain any counts at all. In fact, this is a clear limitation of the Out-of-FOV post-correction method. In cases where a patient, for example, moves immediately after the acquisition start such that a certain fraction of LORs are outside the FOV for the rest of the acquisition, an OFOV correction is not able to restore the underestimated sinogram bins. In addition, looking at the ratio of equation 2.6, it becomes clear that if $t_{OFOV_{i,j}}$ tends towards t , the resulting correction factor will become very large.

Like all statistically based processes, the data registration process in PET is prone to a certain amount of uncertainty (statistical error) in the measured data. In fact, the data sampling process of PET is subject to the effects of the *Poisson distribution*. While the image reconstruction process of PET considers this fact and contains mathematical methods which account for this statistical error, the OFOV scaling becomes problematic. As the OFOV correction scales the number of registered counts, it invalidates the relation between the amount of counts and their ‘expectable’ statistical error. This causes the image reconstruction to account for an amount of uncertainty in the data which in fact is not the case after the OFOV scaling. Therefore, the OFOV correction also scales the statistical error and thus leads to an ‘over-scaling’ of sinogram bins for high OFOV factors. As a result of this over-scaling, image artefacts are generated (cf. figure 2.26 and 2.27), which represent an obstacle for clinical evaluation of such data sets.

In addition to these statistical limitations, an OFOV correction also has to be applied to every spatial transformation that is used during a motion compensation process. Thus, for every ‘significant’ movement identified during a motion analysis (cf. section 2.3.3), the OFOV correction factors have to be calculated and then applied to every reoriented LOR. This is a time consuming task as there are generally numerous relevant motions. Besides there is the

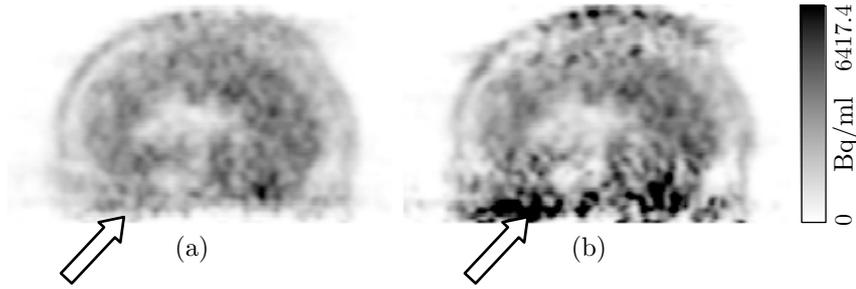


Figure 2.26:

Out-of-FOV problem shown on PET images: A comparison of two sagittal images of the same patient scan. On the left-hand side (a) a motion compensated image without any correction of the Out-of-FOV events is displayed. Aside from a generally reduced contrast in that image, the underestimation at the top and bottom can be identified as a result of the lost events. The same detail, but with an activated Out-of-FOV correction is shown on the right (b). Because of the Out-of-FOV correction, image artefacts are visible, resulting from the over-scaling of counts due to high OFOV factors. The same issue is shown in figure 2.27.

requirement to calculate the OFOV factors for *every* possible detector combination, rather than for every individually registered LOR. Thus, for a PET geometry with millions of possible detector combinations, this easily results in several million spatial transformations that have to be performed and analysed. In case of an ECAT Exact HR⁺ this results in $\approx 8.5 \times 10^7$ required transformations for every applied motion transformation. In practice, the OFOV post-correction is therefore the most time consuming part during an event-driven motion compensation.

One of the main drawbacks of the original OFOV scaling factor calculation is, that the spatial transformations for the realignment of LORs are always calculated relative to a single reference position only – often the position where the patient rested at acquisition start. This is generally a good choice for a motion compensation because of the fact that immediately before an emission scan, a transmission is performed – thus, it is desirable to align the emission data to the transmission data. For an event-driven motion compensation and its involved post-correction algorithms, however, this leads to the OFOV scaling issue because a common patient head motion mainly consists of a steady drifting motion: At the end of a long acquisition a patient’s head is normally located at a position farther away from the start position than this would be the case for a short acquisition. For an OFOV post-correction this results in more OFOV events if compared to motion that may consist of a fluctuation around the same spatial position only. In addition, the steady drifting implies generally higher OFOV scaling factors as some bins are constantly transformed outside the FOV. Furthermore, even for motion that consists of several distinct movements only, the scaling factors increase significantly if these motions occur near the frame boundaries of a dynamic study or if the motion results in some LORs never entering the FOV again (cf. figure 2.27).

Optimisation

To address the problem of image artefacts caused by the over-scaling of sinogram bins we developed an optimised OFOV scaling factor calculation method. This new method is based on the assumption that a minimisation of the OFOV factors automatically leads to a reduction of

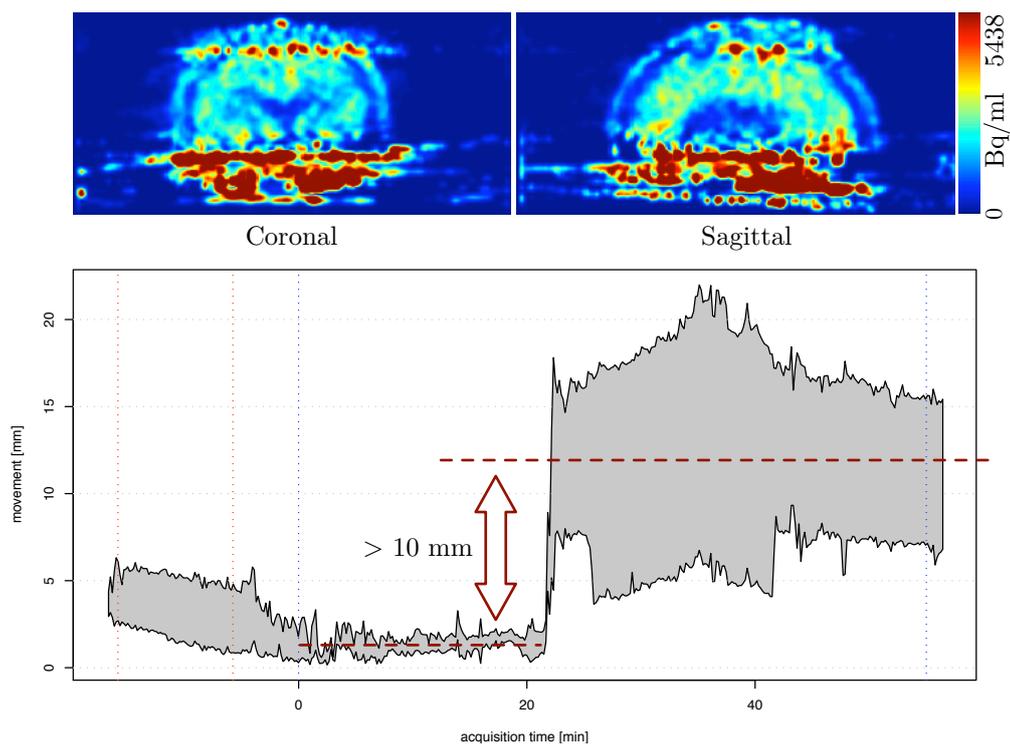


Figure 2.27:

Out-of-FOV problem illustrated on motion data: Due to a large distinct motion, typical Out-of-FOV image artefacts are generated when applying an OFOV correction during an event-driven motion compensation. In the above plot, a typical motion with a range plot of the lowest and highest motion during the acquisition is shown. Due to the large motion at around the 20th minute post acquisition start, a large portion of LORs are transformed outside the FOV without reentering the FOV at a later time. This causes the OFOV correction factors to increase and thus image artefacts to appear, as shown in the above coronal and sagittal images of a brain scan.

the oversampling and thus to a reduction of image artefacts in motion compensated data. In contrast to the original OFOV calculation method, the new method performs analyses on the motion data first. In a pre-calculation step, one or more reference positions are identified and used for the calculation of the spatial transformations. A pre-condition for identifying those reference positions is, that these are average positions in which the patient is estimated to have rested for most of the time. This reduces the OFOV factors due to a reduced magnitude of the applied motion compensation.

For identification and application of those reference positions, the motion is analysed and applied in different steps (cf. figure 2.28):

1. In a first step, the points in time where a motion is considered to be ‘significant’ are identified. Subsequently, all those points are analysed in order to select the largest time interval between two points as the time when the patient is assumed to have remained still for most of the acquisition duration (i.e. the *reference position*).
2. In a consecutive step, transformation matrices are calculated for all other selected positions in time in order to transform LORs to these reference positions. In a dynamic study, this step is repeated for every frame so that a number of reference positions (t_{opt}) are calculated and selected throughout an acquisition. If the OFOV factors are calculated during such an optimised motion compensation, they are not only minimised due to the smaller spatial magnitude of the motion transformations. Also, thanks to this optimisation, an OFOV correction is not required for the selected reference position as the subject is assumed to have remained still during that time interval.
3. As this optimised motion compensation calculates one or more reference positions, an additional alignment between those optimised reference positions is required. Unlike the original OFOV processing, numerous images of the motion compensated data need to be reconstructed for each selected and spatially corrected reference position (t_{opt}). Afterwards, these ‘reference position wise’ images are spatially aligned to the same orientation of the transmission data via rigid body transformations (cf. figure 2.29).

As a result of this process, a single motion compensated data set is generated out of several distinct images which have been previously aligned to separate reference positions.

Implementation

For the implementation of the aforementioned OFOV correction method, we adapted different software packages. For the calculation of the reference time intervals (t_{opt}) and the respective transformation matrices we have modified the motion analysis software `analyze_trk` (cf. section 2.3.3). As a result of this modification, the motion analysis outputs two different motion transformation data files. One file, which is used during the event-driven motion compensation step, carries the parameters to spatially reorient each LOR to the calculated reference position within a frame (inner frame transformation file – `*.ifm`). In contrast, the second file (external frame transformation file – `*.efm`) carries motion parameters for the rigid body transformation of all the frame images to an acquisition-wide reference position – usually the position at acquisition start.

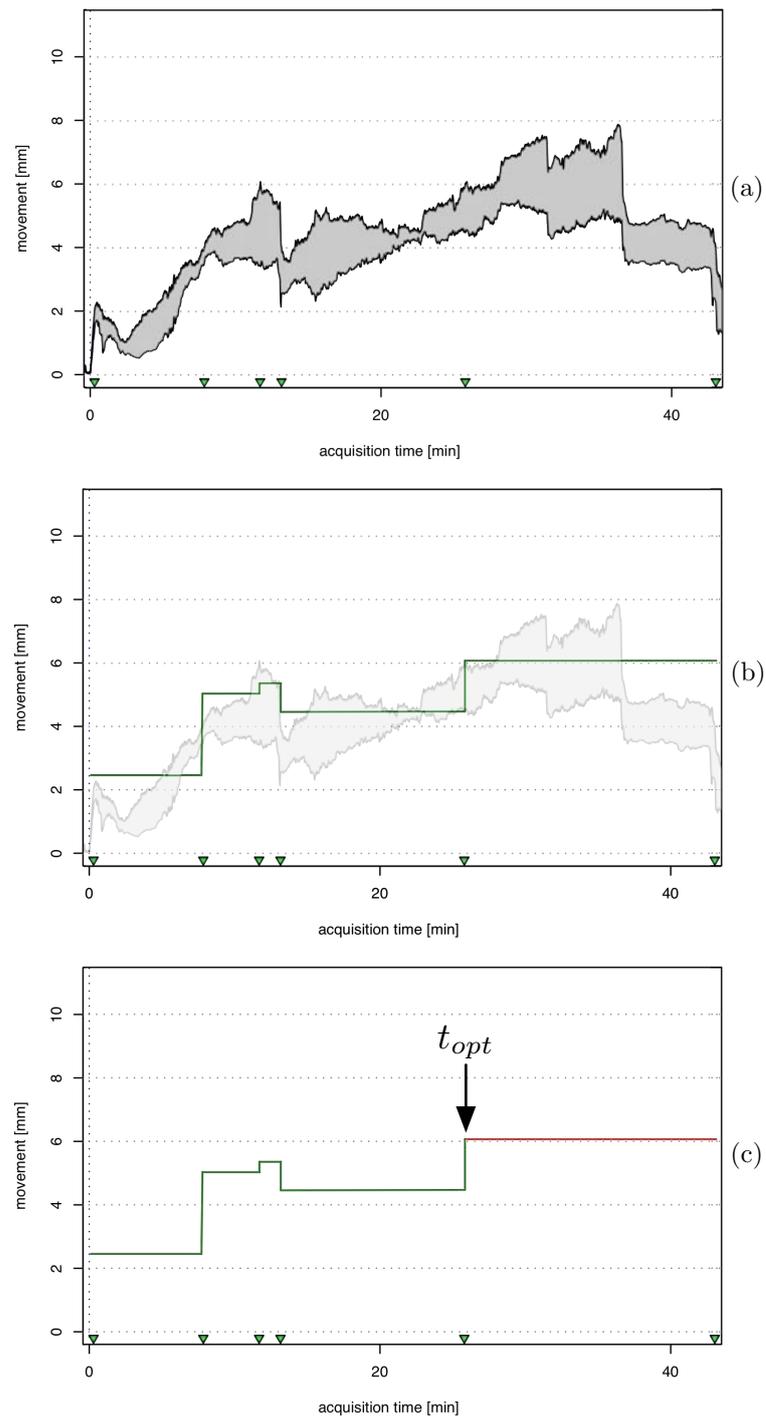


Figure 2.28:

Out-of-FOV optimised Motion Analysis: To identify the point in time where a patient is estimated to have rested for most of the acquisition, the patient motion is analysed via a three step process, illustrated above. In a first step (a), ‘significant’ motion is identified according to the optimised motion analysis discussed in section 2.3.3. In a second step (b) the average spatial position for each of the distinct motion intervals is calculated. A third step (c) finally chooses the motion interval which serves as the *reference position* (t_{opt}). The largest interval is identified and spatial transformations are calculated to transform LORs to the selected reference interval.

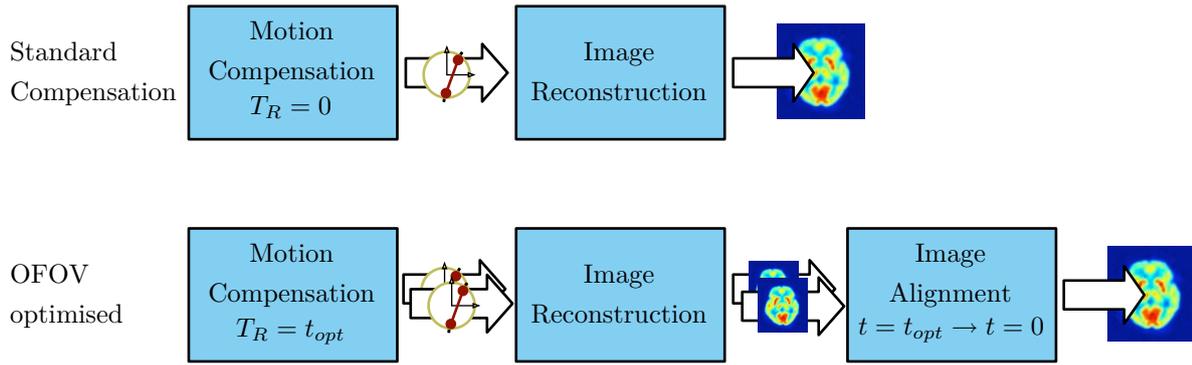


Figure 2.29:

Out-of-FOV optimised Motion Compensation: Compared to a standard event-driven motion compensation, an Out-of-FOV optimised compensation generates multiple data sets in accordance to the numerous reference positions identified during a motion analysis. These motion compensated data sets are then reconstructed individually and spatially aligned to each other. Due to this additional step, the calculated OFOV factors are reduced, which is a pre-condition for an Out-of-FOV aware motion compensation without image artefacts.

As a result of this two-fold motion compensation process, the two transformation data files have to be applied at different stages of the motion compensation method. This has required to modify existing software components (`1mMover`) and to develop additional software for performing the rigid body transformations of PET images. In a first step, the event-driven motion compensation has been modified to reorient the LORs upon the motion information supplied in the `*.ifm` data file, thus reorienting LORs to the frame-wise reference position (t_{opt}). In a next step, a software has been developed to duplicate the attenuation data of the performed transmission scan for every frame of a dynamic PET acquisition. This command script-based software automatically executes an image reconstruction of the attenuation data and then duplicates it for every frame of the emission data. Afterwards, the frame-wise attenuation images are spatially reoriented by a software implementing rigid body transformations on PET image data (`tfmMover`). The transformation parameters used for this reorientation are read from the parameters stored in the `*.efm` data files, and therefore the attenuation images are reoriented to the same reference position as their emission counterparts. Afterwards, the attenuation coefficients of the reoriented attenuation images are calculated via a forward projection of the images. As a result, frame-wise attenuation data sinograms are available, which are aligned to the same frame-wise reference positions as the motion compensated list-mode data.

During the next step, the reoriented list-mode data is sorted into a sinogram and then used together with the spatially aligned attenuation sinograms during an image reconstruction. The result of this process is an emission data set which carries separately reoriented image data for every frame of a dynamic PET study. To align each of the frame images to the acquisition-wide reference position, the `tfmMover` software is used with the inverse of the transformations stored in the `*.efm` data file. After this final step, a complete motion compensated data set is available with all frames motion compensated to an individual *inner-frame* reference position, but aligned to each other via a rigid body transformation to a unified *external-frame* reference position (the position at acquisition start). The complete workflow of this processing is also illustrated in figure 2.30.

Evaluation protocol

For the evaluation of the new optimised OFOV correction method, a number of six example list-mode data sets have been evaluated. These data sets have been selected from a pool of routine PET examinations which have been performed in list-mode with the motion tracking system monitoring the patients' head motion. The acquisition parameters of these dynamic PET examinations are summarised in table 2.3.

Table 2.3:

Out-of-FOV Evaluation Study Parameters: For the evaluation of the new OFOV method a set of six list-mode studies have been used. This table summarises the main parameters of these studies including the mean number of counts per second (cps) as well as the maximum motion found during the 'worst' point motion analysis (cf. section 2.3.3).

	Radio-pharmaceutical	Dosage	Duration	\varnothing cps	# frames	max. motion
(1)	[¹⁸ F]DOPA	316 MBq	55 min	293991	27	9.4 mm
(2)	[¹⁸ F]DOPA	171 MBq	55 min	116324	27	19.4 mm
(3)	[¹⁸ F]FDG	289 MBq	20 min	431941	4	4.9 mm
(4)	[¹⁸ F]OMFD	333 MBq	60 min	241383	28	5.1 mm
(5)	[¹⁸ F]FDG	244 MBq	20 min	396362	4	2.4 mm
(6)	[¹⁸ F]DOPA	252 MBq	55 min	232274	27	9.3 mm

For the selected acquisitions, the list-mode data have been motion compensated for one test with the *non-optimised* and for a second test with the *optimised* OFOV post-correction method activated. For the final evaluation those different motion compensated data sets have been analysed in two phases:

1. In the first phase of the evaluation the qualitative improvement of the images caused by the new OFOV method has been evaluated. Via a subjective comparison of the non-optimised and optimised PET image the reduction of the image artefacts has been analysed. If the reduction of artefacts was clearly visible the application of the optimised OFOV correction method was rated as successful.
2. During the second test phase the *lmMover* software has been modified to output all calculated OFOV factors – $f_{OFOV(i,j)}$. The OFOV factors of the optimised and non-optimised OFOV processing have been saved to a data file and afterwards statistically analysed. As the reduction of the OFOV factors is a direct indicator of the performance of the optimised OFOV correction method, the distribution of the OFOV factors have been analysed via a quantitative analysis. In a standard box plot output not only the distribution of the factors, but also the outliers in the non-optimised OFOV correction have been illustrated. Judging on this graphical output and the statistical parameters the improvement in the scaling factor calculation due to the general reduction of the OFOV factors has been illustrated.

2.5 Clinical Integration

Before implementing the aforementioned methods in a clinical environment, certain practical aspects have to be considered.

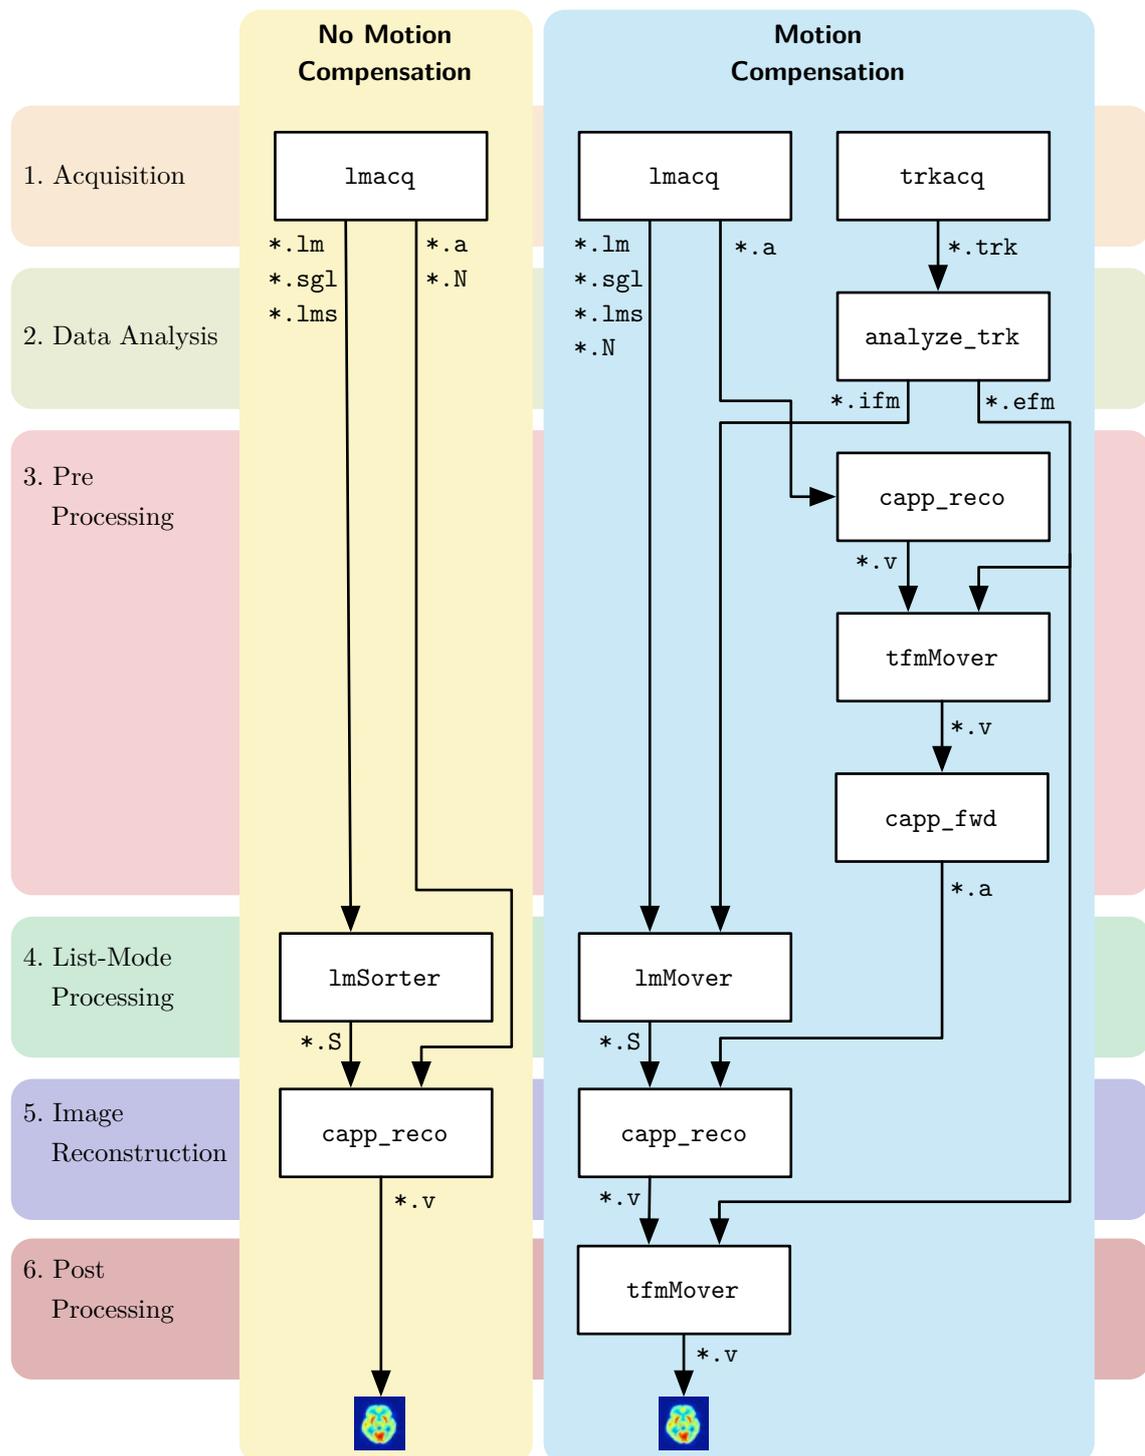
A clinically applicable acquisition of list-mode requires, for example, the sampling of the coincidence data either via the presented DAQ methods or via accessing the list-mode data through the shared storage device solution, as shown in section 2.2. However, before using these non-standard acquisition methods in a clinical environment they have to be embedded in the standard protocols which are regularly used by the medical staff. In addition, the methods should not require any extra time during a medical examination. The same applies to the shell commands that have to be executed in order to set up the PET scanner for list-mode processing. Furthermore, after a list-mode acquisition the required post-processing of the list-mode data (sorting, motion compensation, etc.) is another time-critical task. This means, that either a simple sorting into a sinogram or the application of an event-driven motion compensation have to be processed fast enough so that they do not represent an obstacle for clinical applications. Finally, as the standard acquisition environment usually consists of a graphical user interface, a graphical implementation of the aforementioned acquisition methods is desirable as well. To account for these things, six command-line-based tools as well as two graphical interface-based applications have been developed.

Our tools and the workflow of their application are illustrated in figure 2.30. A detailed description on their implementation is shown in the follow paragraphs.

lmacq – List-Mode Acquisition Tool

One of the developed tools with a graphical user interface (`lmacq`) implements the optimised list-mode acquisition methods of section 2.2. This application has been designed to be invoked by a protocol block in the standard ECAT user interface. To facilitate its use by the medical staff, the graphical elements of this user interface have been developed to be compatible with the user interface of the PET scanner. For this reason, the medical technicians have assisted in defining the requirements and layout of the user interface. A screenshot is shown in figure 3.17 of the results chapter.

When executed, the application automatically retrieves patient information via a database connection to the ECAT system. Furthermore, it allows for a manual definition of acquisition related and patient related information directly before an acquisition. After confirmation, the application performs all necessary steps to set up the PET scanner for list-mode. This also includes the process of executing a transmission scan either before the emission phase or afterwards. In addition, the application contains graphical elements allowing for monitoring the progress of an acquisition as well as for aborting it immediately on user request. After an acquisition is finished, it automatically issues a transfer of the list-mode data (`*.lm` and `*.sgl` files) to a given processing workstation. On that workstation the list-mode data is either sorted into an ECAT sinogram (via `lmSorter`) or an event-driven motion compensation is automatically issued (via `lmMover`). Afterwards, the sinogram data file (`*.S`) is forwarded to the standard image reconstruction (`capp_reco`) of the PET scanner. As soon as this process is finished, the data is stored into the internal patient database of the PET scanner so that it is available as any other acquisition data on the HR⁺. Wherever possible, this stepwise acquisition and processing of list-mode data has been developed to be performed via the available standard ECAT commands, as shown in section 2.2.3.

**Figure 2.30:**

Clinical Integration Workflow: For the integration of our methods in a clinical environment, a number of seven software applications have been developed. As shown, these tools (e.g. `lmMover`, etc.) are used either for the processing of an event-driven motion compensation or for using list-mode for a normal examination only. The figure illustrates the developed tools and their application. In addition, the output and input data files are listed via their file extensions.

trkacq – Motion Tracking Acquisition Tool

For the acquisition of the motion data, an additional graphical user interface (`trkacq`) has been developed. This user interface uses the same graphical toolkit (Qt4) [110] and elements as the `lmacq` application. It has been designed to be automatically executed by an ECAT protocol block as well, as shown in figure 3.18. If this generic protocol is used, both the `lmacq` and `trkacq` applications are executed and allow for configuration of all necessary parameters for the acquisition of the motion tracking and list-mode data. After the motion tracking acquisition is finished, a motion data file (`*.trk`) is generated and forwarded to the motion analysis application (`analyze_trk`), as discussed in section 2.3.3. As soon as ‘significant’ motions have been identified and extracted from the motion data, the calculated inner-frame transformation matrices are saved (`*.ifm`) and forwarded to the event-driven motion compensation application (`lmMover`). In parallel a generated transformation data file on the external-frame motion (`*.efm`) is supplied to the routines transforming the attenuation data.

To provide a live view of the motion tracking, the `trkacq` application contains a graphical plot element. During the motion acquisition, an estimate of the magnitude of the motion is calculated every second and then displayed via a red horizontal polygon along the abscissa of the plot. Similar to the user interface of the list-mode acquisition, this interface also allows for aborting and restarting a motion tracking acquisition at any time. A screenshot of this interface is shown in figure 3.17 of the respective results chapter.

lmSorter – List-Mode Sorting Tool

In case no motion compensation is required (e.g. no significant motions were found), the list-mode data is sorted into a standard ECAT sinogram (`*.S`) only. For this process a command-line-based software has been developed (`lmSorter`) and enhanced to utilise multiprocessor systems. During the sorting process, the list-mode data stream is split into independent units which are processed by individual processors of a computer system. This multiprocessor optimisation ought to speed up the sorting of list-mode so that no additional delay for the software-based sorting, in contrast to the PET scanner’s own hardware sorting, is required. Furthermore, as this application has a command-line interface only, the `lmacq` application has been modified to execute `lmSorter` as soon as the list-mode acquisition is finished. In addition, the resulting sinogram is automatically forwarded to the image reconstruction process (`capp_reco`) which calculates the image data from the sinogram.

lmMover – Event-based Motion Compensation Tool

The event-driven motion compensation algorithms, as shown in section 2.4, have been integrated in a command-line-based software (`lmMover`). Similar to the `lmSorter` application, this software has been optimised to process the list-mode data with multiple processors. Despite the working of the sorter application, the multiprocessor optimisations of `lmMover` are, however, limited to the parallel processing of acquisition frames. This is required because the OFOV correction is dependent on the actual motion data. So for every new motion transformation data an

OFOV correction have to be issued. A detailed description of this multiprocessor optimisation is discussed in [23].

The `lmMover` application is executed as soon as the `lmacq` and `trkacq` applications are finished with their acquisition of the list-mode and motion tracking data. It uses the list-mode data files (`*.lm` and `*.sgl`) which are supplied by the aforementioned list-mode acquisition methods. In addition, it uses a meta-data file (`*.lms`) which contains all acquisition relevant and patient relevant information. This file is created by the `lmacq` application and meant to carry all relevant information on a patient examination (patient name, acquisition duration, etc.). After the successful processing of the motion compensation, a standard ECAT sinogram (`*.S`) is generated. Similar to the processing of the `lmSorter` application, the resulting sinogram file is forwarded to the image reconstruction tool (`capp_reco`). However, in contrast to the plain sorting of list-mode data, a motion-corrected PET image is afterwards post-processed by the `tfmMover` application which generates the final motion-corrected image data.

tfmMover – Rigid-body Transformation Tool

As the optimised OFOV correction method (cf. section 2.4.3) requires the alignment of the attenuation data, a command-line tool for performing rigid-body transformation on ECAT volume images have been developed (`tfmMover`). This application is used twice during an event-driven motion compensation. At the first use it is executed with the external-frame motion data file (`*.efm`) and the reconstructed attenuation data (`*.v`) which is generated by an image reconstruction of the measured attenuation (`*.a`). As a result of this process, an image data file is generated with the attenuation data aligned to the same orientation of each acquisition frame (t_{opt}). During the second use of `tfmMover`, the inverse of the data contained in the external-frame motion data file (`*.efm`) is used for the spatial transformation. Right after the successful image reconstruction (`capp_reco`) of the motion compensated sinogram, the frame data in the image file (`*.v`) is spatially reoriented to a unified reference position (e.g. the position at acquisition start). This is illustrated in figure 2.29.

capp_reco/capp_fwd – Wrapper Scripts to CAPP reconstruction

To be able to perform an image reconstruction and forward projection independently from the user interface of the HR⁺, two command-line-based tools (`capp_reco` and `capp_fwd`) have been developed. These two tools are command scripts which take a standard ECAT file (`*.S`, `*.a` for `capp_reco` and `*.v` for `capp_fwd`) as an input data file and either performs an image reconstruction or forward projection of ECAT image data. When these tools are executed on a system with an installed CAPP software package (Clinical Applications Programming Package), the same reconstruction and forward projection routines are used like on the graphical user interface of the ECAT console.

2.6 Clinical Evaluation

To evaluate the aspects of the application of our event-driven motion compensation in a clinical environment, two main tests have been performed. In a first evaluation a radioactive point

source has been attached to a motion tracking target and constantly moved during an emission scan. For the second evaluation the motion tracked PET data of routine patient acquisitions have been analysed and motion compensated respectively. Via a comparison of the motion compensated data sets to the uncorrected data the qualitative and quantitative influence of the motion compensation has been evaluated.

I. Point Source Evaluation

For the evaluation of the ability and accuracy of our optimised motion compensation a radioactive [^{68}Ge] point source with a diameter of ≈ 1 mm and an activity of 3.83 MBq has been attached via adhesive tape to a standard motion tracking target. During a 3D list-mode emission of 10 minutes, the point source has been manually moved during the last half of the acquisition to arbitrary locations. The effective motion magnitude is illustrated in figure 2.31. According to the motion analysis (cf. section 2.3.3) the performed motion shows a maximum distance from the initial position of approximately 11 cm. In fact, the distances of the ‘worst’ point on the sphere surface (as the figure demonstrates) are between 0.1 mm and 40.3 mm within the 10 % to 90 % percentiles of the distribution of the distance for the whole motion acquisition.

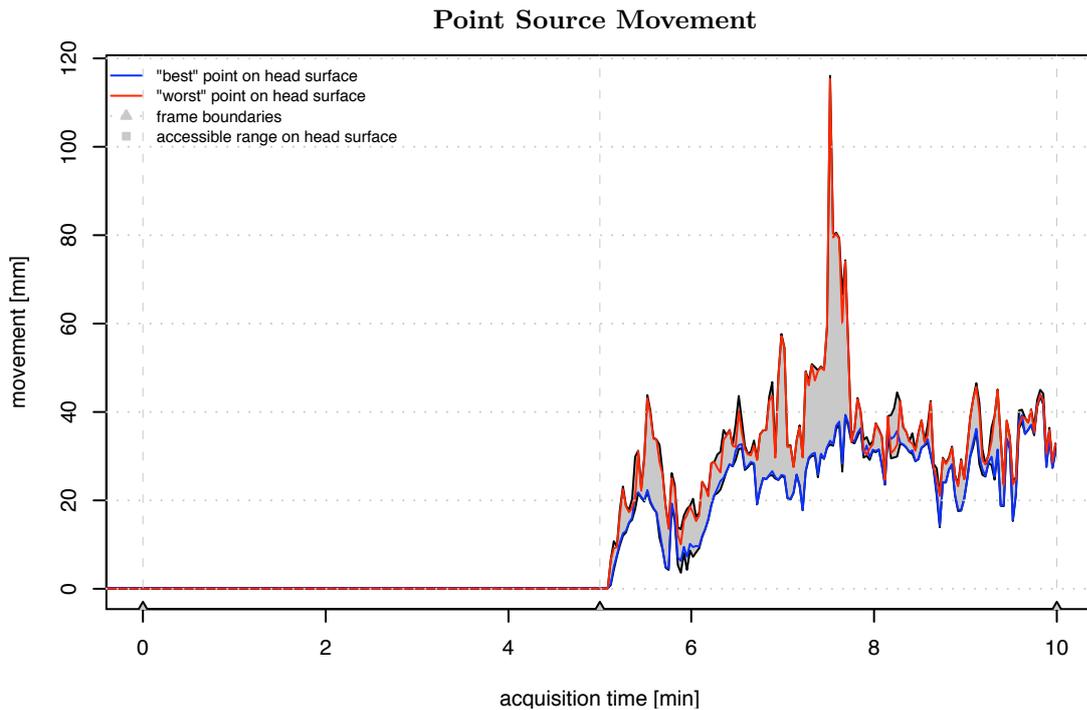


Figure 2.31:

Point Source Movement: For the evaluation of the accuracy of the event-driven motion compensation a radioactive point source has been moved in arbitrary locations within the FOV of the PET scanner. The given figure illustrates the distribution of the motion according to the motion analysis method shown in section 2.3.3.

By applying the optimised event-driven motion compensation method, the ability to restore the data of the moved point source have been evaluated. For this purpose three images have been reconstructed into $128 \times 128 \times 63$ matrices with a voxel size of $2.6 \times 2.6 \times 2.4$ mm (weighted OSEM reconstruction, 6 iterations, 16 subsets, no filter). One image has been reconstructed from the

uncorrected data set of the last half of the test acquisition and shows the typical blurring due to the performed motion. In contrast, a second image has been reconstructed from the motion compensated data set instead. Furthermore, to be able to compare the motion compensated image to an image generated from data without any motion a third image (reference image) has been reconstructed from the first half of the test acquisition. Based on these three images a qualitative rating of the performed motion compensation has been applied.

For a quantitative evaluation of the motion compensation accuracy, an analysis of the influence of the motion compensation on the effective image resolution has been performed. For a simple object geometry and an isotropic *Gaussian Point Spread Function* (PSF) the image resolution (FWHM) can be derived by calculating the convolution of the object function and the PSF. For spheres, to which a ‘point’ source with a diameter of ≈ 1 mm belongs to, such an equation has already been given by [111]. To calculate the image resolution of the reference and the corrected image we have used a tool (`psfFit`) which has been developed in our facility [105, 112]. This tool has already been evaluated in our facility and is used as a standard tool to determine the effective resolution of a PET scanner out of a phantom measurements with a number of standard sphere inserts.

II. Patient Data Evaluation

To evaluate and demonstrate the new motion compensation method under clinical conditions the patient motion of a number of fifteen clinical acquisitions have been analysed. These 15 list-mode-based acquisitions have been selected from a pool of routinely performed [^{18}F]DOPA brain acquisitions. They have been carried out with the motion tracking target attached to the patients’ head as a quality control method. As a criterion for the selection, the estimated motion score as well as the number of significant motions found in the respective motion report document (cf. section 2.3.3) have been used. The acquisition and motion parameters of these patient measurements are shown in table 2.4.

Furthermore, one study of the fifteen patient measurements have been selected for a more detailed analysis of the influence of the motion compensation. Based on the motion magnitude parameters shown, study 3 has been selected. Because of the occurred motion during this examination the patient movement, with a maximum of ≈ 19 mm and a mean of ≈ 11.3 mm, suggests to have a severe effect not only on the quality of the image but also on the quantitate results of a *tracer kinetics analysis*. To emphasise this decision, the motion plot of this examination is shown in figure 2.32. It shows a generally large movement at the acquisition time of approximately 20 min post injection. For a qualitative evaluation of the motion compensation an uncorrected and corrected volume image with a dimension of $128 \times 128 \times 63$ have been reconstructed from the list-mode data of study 3 and compared to each other (CAPP reconstruction, weighted OSEM, 6 iterations, 16 subsets, 4 mm Hann filter).

In addition to the qualitative inspection of visible differences and potential motion artefact eliminations, a tracer kinetics analysis has been performed on both images of study 3. For the quantification of the binding of dopamine within the area of the striatum a number of 9 three-dimensional elliptical ROIs have been positioned in each respective PET image. This ROI delineation has been carried out with the ROI visualisation and evaluation tool *ROVER* [105,

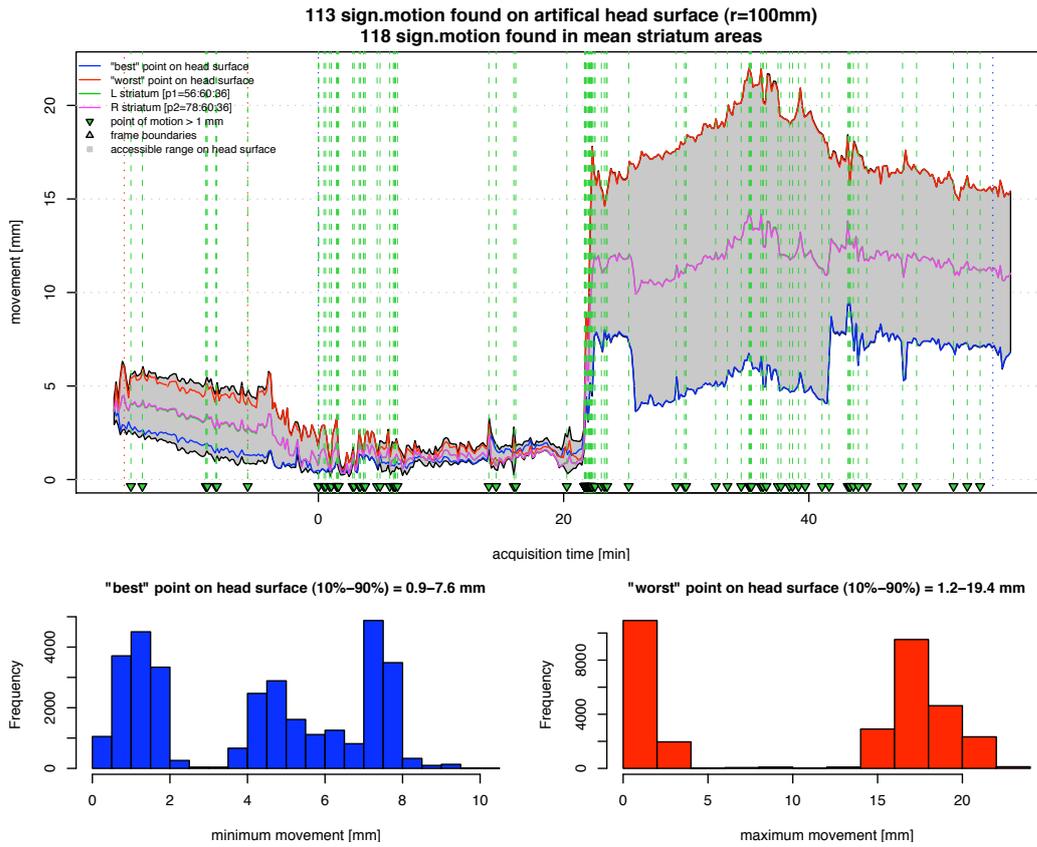


Figure 2.32:

Patient Examination Movement: During the clinical evaluation of motion affected patient data, one patient data set has been selected for a detailed analysis. The shown plot shows the motion data of this patient examination with the significant motions (> 1 mm) marked with green vertical lines in the range plot. In addition, the frequency plots on the 'best' and 'worst' point motion on the imaginary patient head are shown.

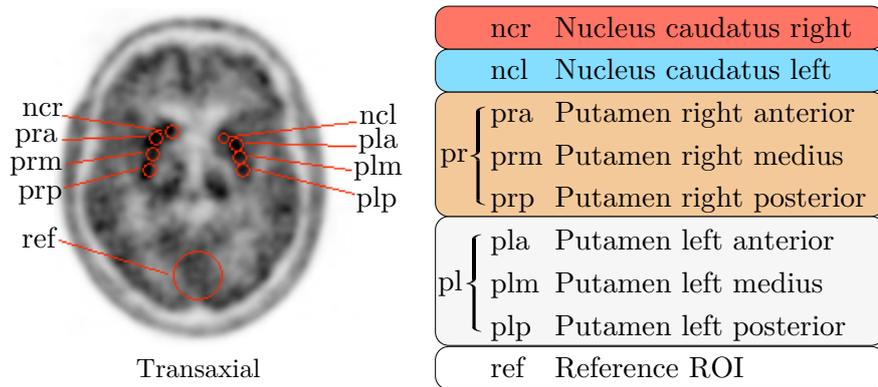


Figure 2.33:

ROI delineation procedure: During the quantitative evaluation of the accuracy of the motion compensation a number of nine three-dimensional elliptical ROIs have been delineated within the selected patient data. This delineation procedure has been carried out by an experienced physician in accordance to the established protocol for the analysis of [¹⁸F]DOPA studies [113–115].

Table 2.4:

Patient Examinations Parameters: For the evaluation of the clinical compliance of the new optimised motion compensation method a number of 15 routinely performed [^{18}F]DOPA acquisitions brain acquisition have been analysed. Based on the motion score and distribution of the motion one acquisition has been selected for application of the motion compensation and for a further detailed quantitative evaluation.

	Dosage	Duration	# frames	# 'significant' ($> 1\text{ mm}$) motions	max. motion	mean motion	motion score
(1)	291 MBq	55 min	27	17	3.7 mm	3.1 mm	3
(2)	252 MBq	55 min	27	26	5.0 mm	3.6 mm	3
\Rightarrow (3)	171 MBq	55 min	27	113	19.4 mm	11.3 mm	5
(4)	269 MBq	55 min	27	70	7.9 mm	5.7 mm	5
(5)	299 MBq	55 min	27	45	5.7 mm	5.0 mm	4
(6)	316 MBq	55 min	27	801	9.4 mm	6.3 mm	5
(7)	254 MBq	55 min	27	41	8.6 mm	6.6 mm	5
(8)	318 MBq	55 min	27	19	8.5 mm	6.0 mm	5
(9)	250 MBq	55 min	27	14	8.4 mm	6.0 mm	5
(10)	312 MBq	55 min	27	22	5.4 mm	3.2 mm	4
(11)	292 MBq	55 min	27	12	6.3 mm	3.7 mm	4
(12)	300 MBq	55 min	27	16	3.4 mm	3.4 mm	3
(13)	183 MBq	90 min	27	125	8.4 mm	5.2 mm	5
(14)	275 MBq	55 min	27	92	6.5 mm	3.6 mm	4
(15)	264 MBq	55 min	27	17	7.1 mm	5.3 mm	5

116]. An experienced physician delineated each of the 9 ROIs in both data sets according to the standard protocol used at our facility. This method is described in [113–115] and used for the quantification and evaluation of the dopamine binding for [^{18}F]DOPA related questions (e.g. *Parkinson's disease*). Figure 2.33 illustrates this delineation on an example data set. Eight of the delineated ROIs have a volume of 5.24 ml ($1.03 \times 1.03 \times 1.70\text{ cm}$) where two of these ROIs (ncl and ncr) have been placed where the *Nucleus caudatus* (left/right) has been most clearly seen. The same procedure has been repeated for the other six ROIs (pla,plm,plp and pra,prm,prp) which have been placed on the *Putamen* (left/right). In addition, the data in all three ROIs delineating each Putamen have been combined in a ROI (pl and pr) summarising the activity of the other three ROIs. Furthermore, as part of the ROI delineation procedure a ninth elliptical ROI has been placed in the area of the *Occipital lobe*. This ROI has a volume of 79.28 ml ($3.5 \times 3.5 \times 1.70\text{ cm}$) and is used as a reference/background ROI for the further quantitative tracer analyses.

To visualise the tracer accumulation over time *time-activity-curves* (TAC) have been calculated. These TACs illustrate the accumulation of the injected tracer within a specific ROI throughout the acquisition (in steps of acquisition frames). For the uncorrected vs. corrected comparison we calculated TACs for every delineated ROI. By comparison of the TACs of the corrected and uncorrected PET image a difference in the tracer accumulation of both volume image and thus also a restoration of the TAC due to the motion compensation ought to be identified.

The same TACs serve as an input for a quantitative tracer kinetics analysis on the base of an *irreversible reference tissue two compartment model* [4, 117]. In this analysis the tracer kinetics

parameter R_0k_3 is commonly used for the determination of the dopamin metabolism. For this purpose an in-house software (`roitool`, FZ Dresden-Rossendorf, ABX Radeberg) has been used which is based on the graphical method for the determination of that value described by Patlak and colleagues [118, 119]. This tool has been established as the standard method in our facility for the calculation of a *Patlak plot*, from which the R_0k_3 parameter of each ROI is derived from. In clinical practice the value of this parameter is usually inspected in relation to the values found in a healthy patient collective and thus a statement on the healthiness of the dopamin metabolism in a human brain is given. As the ROI delineation and thus the calculation of this parameter is usually highly influenced by the existence of motion artefacts, a direct comparison of the R_0k_3 parameters found in a corrected and uncorrected image ought to demonstrate the accuracy of our event-driven motion compensation. For this process we have calculated the R_0k_3 parameter for both sets of the uncorrected and corrected ROIs on the PET data for minute 4 to 40 (frame 16–24).

In addition to the calculation of the R_0k_3 parameter for the specific ROIs, parametric images have been calculated for the uncorrected and corrected data set. These images have been generated with `roitool` and have the same dimensions as the original PET images ($128 \times 128 \times 63$). To generate these parametric images, for each voxel of the image a TAC is calculated and used for the determination of a voxel-wise R_0k_3 parameter, as already explained. As a result, the value of a voxel in a parametric image corresponds to the R_0k_3 value calculated for the same voxel in the original PET volume image. One peculiarity of such a parametric image is, that artefacts due to motion are easily visible due to high R_0k_3 values appearing within untypical regions, e.g. along the skull of a human head. Therefore, in addition to the aforementioned qualitative comparison of the uncorrected and corrected PET image a comparison of the parametric images of the uncorrected and corrected data set have been carried out. In a qualitative analysis of the differences in the R_0k_3 images, the influence of the motion compensation and thus the ability to correct motion artefacts in the R_0k_3 images are immediately visible.

3 Results

You are today where your thoughts have brought you; you will be tomorrow where your thoughts take you. You cannot escape the results of your thoughts.

(Author Unknown)

3.1 Optimised List-Mode Data Acquisition

As shown in section 2.2, the development of methods for a direct acquisition of list-mode data with an ECAT Exact HR⁺ PET scanner is a necessity. In the following section, the results of these methods to improve the list-mode capabilities of an HR⁺ are shown.

3.1.1 Shared Storage Device Acquisition

After the installation and configuration of the dual-channel Ultra160-SCSI RAID system, the new list-mode acquisition method has been evaluated. In 40 consecutive list-mode acquisitions the correct operation and performance of the device has been analysed and evaluated according to the test protocol discussed in section 2.2.1. The results on these evaluations are summarised in table 3.1.

Table 3.1:

Shared Storage Device Evaluation Results: For the evaluation of the *shared storage device* method 40 list-mode test acquisitions have been performed (cf. section 2.2.1). During the test acquisitions four evaluations have been performed to verify the correct operation of the acquisition method.

	Evaluation Method	# Acq.	Criterion	Value
(1)	Hardware Compliance (stability/compatibility)	10	# of dropouts	0
(2)	Data Consistency (test data binary compare)	10	# differences	0
(3)	Speed Comparison (data access speed)	10	speed-up rate	≈ 140×
(4)	Auto-Mounter Stability (remounting/remote-deletion automatic)	10	# of dropouts	0

After having performed the test acquisitions with the connected RAID system, no problems with the general stability of the hardware were observed. Rather, the RAID system was transparently accepted as a usual hard disk device by the ACS2 system. Due to the backward compatibility of the SCSI standard, the slower SCSI-2 connection between the ACS2 and the Ultra-160 RAID system did, in fact, not cause any failures during the evaluation.

To verify the simultaneous operation of the RAID system with both systems (ACS2 and Linux workstation), parallel data transfers were initiated during the test acquisitions. In addition to a general evaluation for potentially existing hardware interferences between the two systems, our file system implementation (VXEXT) has been evaluated. A binary comparison of the sampled list-mode data has been performed and did not show any differences. Moreover, a high speed-up of the transfer of acquisition data to the local storage device of the Linux system has been observed. The list-mode data of the test acquisitions has been transferred with a maximum throughput of $\approx 70 \text{ MB s}^{-1}$. This represents a $\approx 140 \times$ speed-up in comparison to the ACS's own capabilities to transfer acquisition data to another workstation via a network connection (cf. figure 3.1).

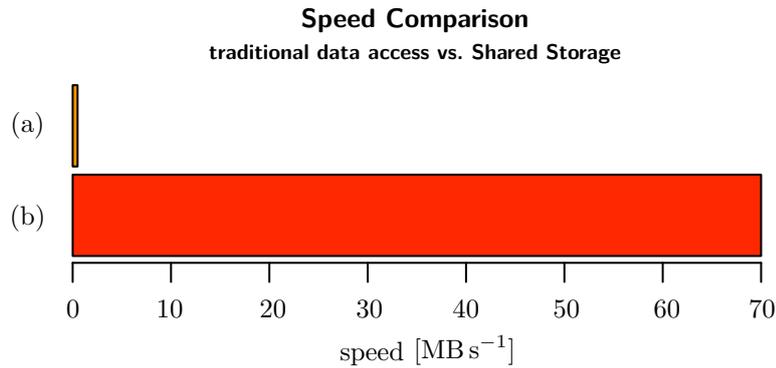


Figure 3.1:

Data Transfer Speed Comparison: In ten test acquisitions the data access speed of the new *shared storage device* method (b) and the traditional data access method of the ACS2 (a) has been evaluated. While the speed of the network-based data access of the ACS2 has been found to be $\approx 0.5 \text{ MB s}^{-1}$, the new access method allows for an access of data with $\approx 70 \text{ MB s}^{-1}$. This represents a $\approx 140 \times$ speed-up.

As discussed in section 2.2.1, a simultaneous write access from both systems (ACS2 and Linux) to the shared storage device is not possible – but not strictly required either. As suggested, the auto-mounter of the Linux system was modified to mount the file system of the RAID system on demand only. In the test acquisitions the auto-mounter functionality has been evaluated as well. No problems with this on-demand method of accessing and transferring the coincidence data were found. In addition, the remote deletion via the `rfaDelete` command proved to work reliably, thus no dropouts during the test acquisitions have been observed.

3.1.2 DAQ-based Data Acquisition

As outlined in section 2.2.2, a hardware adapter card has been developed to isolate the electric circuits of the ACS2 from the circuits of the digital acquisition cards installed in a Linux system. To verify the correct operation of our adapter card and the DAQ method in general, we performed 42 test acquisitions. In four test phases we performed different types of evaluations. Tables 3.2 summarises the results of the evaluation tests.

In the first phase twelve test acquisitions have been performed for the evaluation of the hardware compliance with the standard ACS2 hardware. As a result of these test acquisitions the number of dropouts have been found to be zero. Moreover, no dropout or system instabilities

Table 3.2:

DAQ-based Data Acquisition Evaluation Results: In ten test acquisitions the performance and stability of the DAQ-based list-mode data acquisition method has been evaluated (cf. section 2.2.2).

	Evaluation Method	# Acq.	Criterion	Result
(1)	Hardware Compliance (stability/compatibility)	4×3	# of dropouts	0
(2)	Data Consistency (test data binary compare)	10	# differences	$< 10^{-6} \%$
(3)	Real-Time Performance (delay until data retrieval)	10	data retrieval delay	< 1 sec
(4)	Load/Stability Test (increasing count rate)	10	# of dropouts	0

of the ACS2 were found as a result of the connected and/or enabled adapter card. In a second test phase the data sampling consistency has been evaluated via ten list-mode acquisitions with the data acquired in parallel via the new DAQ method and the traditional ACS2 data acquisition method.

After a detailed binary comparison of the sampled coincidence data of the test acquisitions, no relevant differences in the sampled data stream were found. In fact, due to the hardware caching of the DAQ cards, a typical amount of approximately $\approx 5 \pm 1$ coincidences could not be sampled at the beginning and at the end of each acquisition. Apart from these few lost counts, the rest of the typically 2 GB to 4 GB large data stream of each test acquisition (i.e. which sums up to $\approx 10^8$ to 10^9 coincidences) were found to be consistent with the data set obtained via the original data acquisition methods of the ACS2. In fact, the inconsistency of the data has been found to be lower $6 \times 10^{-6} \%$.

For the third test phase the real-time performance has been evaluated. According to ten test acquisitions the delay until a coincidence is available to be processed by the Linux system has been quantified. As a result of the tests the data retrieval delay was found to be always less than one second and thus within an acceptable real-time acquisition range. In the final test phase the general stability of the DAQ-based sampling at high count rates has been evaluated via further test acquisitions. As a result of ten test acquisitions, the DAQ method was found to be able to sample up to the count rate at which the electronics of the ACS2 itself had to drop coincidences due to the high load. The variation of the applied radioactivity during the test acquisitions did not result in any shortcomings or dropouts, either.

3.1.3 Event-based Attenuation Measurement

After implementation of the command scripts to execute a list-mode-based transmission scan (cf. section 2.2.3), the new method has been evaluated via 30 test acquisitions. During the evaluation the tests have been grouped in two phases. In the first test phase the general stability of the list-mode-based transmission scan has been evaluated. During the second phase another ten acquisitions have been performed for a comparison of image data generated at one occasion of a standard histogram-mode acquisition and at another occasion of a list-mode-based transmission. The results of these test phases are summarised in table 3.3.

Table 3.3:

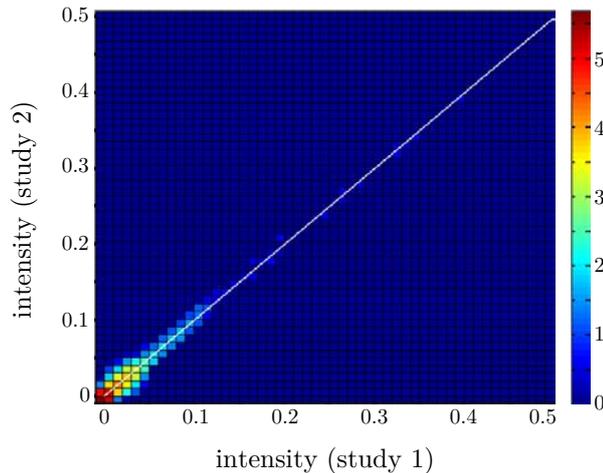
Event-based Attenuation Measurement Evaluation Results: In two test phases a number of 30 transmission acquisition have been performed with the new event-based method. According to an evaluation protocol (cf. section 2.2.3) the stability and the data consistency have been evaluated.

	Evaluation Method	# Acq.	Criterion	Value
(1)	Stability (stability/compatibility)	10	# of dropouts	0
(2)	Data Consistency (image correlation test)	2×10	# differences	< 3.5 %

After the first test phase no dropouts or errors during the tests have been identified. Even after having varied the acquisition duration or by manually aborting acquisitions no dropouts during the test acquisitions have been found. Rather, the new transmission protocols have been working fully transparently in the ECAT environment of the HR⁺. During the second test phase another 20 test acquisitions have been performed with a phantom located in the FOV of the PET scanner. For the first half of the test acquisitions a transmission in traditional histogram-mode has been performed and compared to the data of the other half of the acquisitions performed with the new event-based transmission method. After having sorted the list-mode data into sinograms via the modified sorter software, the images of both (histogram-mode-based and list-mode-based) sinograms were reconstructed for the comparison. As part of the second test phase the images have been compared to each other via an intensity correlation analysis. An example correlation plot of this analysis is shown in figure 3.2.

Figure 3.2:

Logarithmic intensity correlation plot (2D histogram): The differences between an image generated from a histogram-mode transmission (study 1) and an image generated from an event-based attenuation measurement (study 2) are shown. The plot shows that differences between the two images exist mainly in lower intensities. In fact, the difference between the images can be read from the deviation of the intensities from the white identity line. After quantification, it was found to be < 3.5 %.



This voxel-wise analysis suggests that there is a minor difference between a histogram-mode-based and list-mode-based transmission scan. In fact, the difference in correlation of the voxel intensities in the images of both methods has been found to be lower than 3.5 % according to the correlation analysis.

3.2 Motion Quantification

As shown in section 1.2.2, additional methods are required for the quantification of patient motion within the FOV of a PET scanner. One method is the required cross-calibration between the two independent coordinate systems. Another method is the support of an adequate motion target which is able to convey the patient motion accurately. Furthermore, a detailed analysis of the motion is required to distinguish between ‘significant’ and ‘non-significant’ motion.

3.2.1 Cross-Calibration

As a result of the development on the cross-calibration method outlined in section 2.3.1, a software package (`trkCalib`) has been developed. Besides the implementation of the cross-calibration algorithms this software also provides a graphical user interface. A screenshot of this interface is shown in figure 3.3.

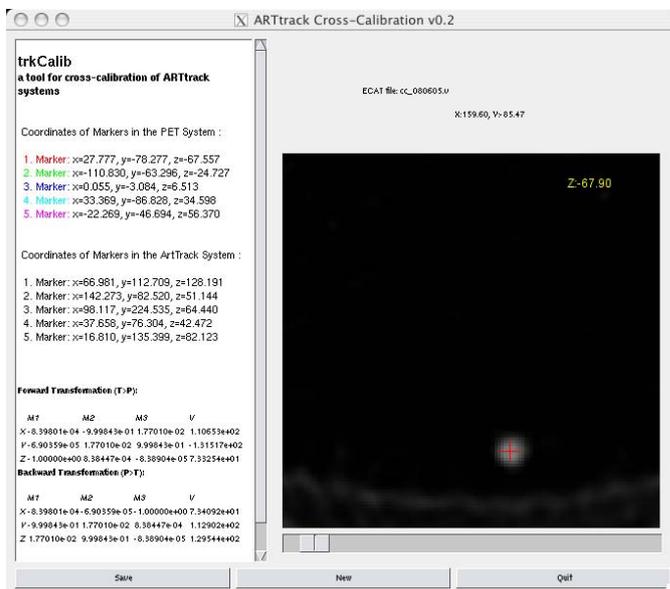


Figure 3.3:

Screenshot of the graphical `trkCalib` user interface: On the left-hand side of the interface, the retrieved and calculated locations of each marker of the cross-calibration measurement are displayed. On the right-hand side, a graphical view of the transmission image (here with one marker) is shown. Each marker is displayed with a cross at its identified centre of mass. According to the identified centres of a minimum of four markers, `trkCalib` is able to calculate the coordinate transformation matrix.

The user interface is split into two main areas. One area (right-hand side) covers an interactive display of the loaded transmission image of the cross-calibration. In this display the found centres of the markers are highlighted with colours. The other area (left-hand side) displays text output of the processing of the matrix calculation and displays the final parameters of the cross-calibration matrix. In addition, the `trkCalib` software contains export functions which allow for saving the cross-calibration matrix into a separate data file (`*.tfm`). This file is then used during a motion analysis and compensation.

To evaluate the operation of our software and to verify that the cross-calibration algorithms are able to calibrate the two coordinate systems adequately, we performed ten test acquisitions. In a first test phase five cross-calibrations have been performed and the cross-calibration matrices have been calculated with the `trkCalib` software. The stability of the cross-calibration transmission measurements has been inspected. No dropouts or errors during the execution of this new transmission protocol have been observed. In addition, a validity check of the cross-calibration matrices have been performed by calculating the determinant of the transformation matrix. In all cases the determinant has been calculated to be one and thus valid.

Table 3.4:

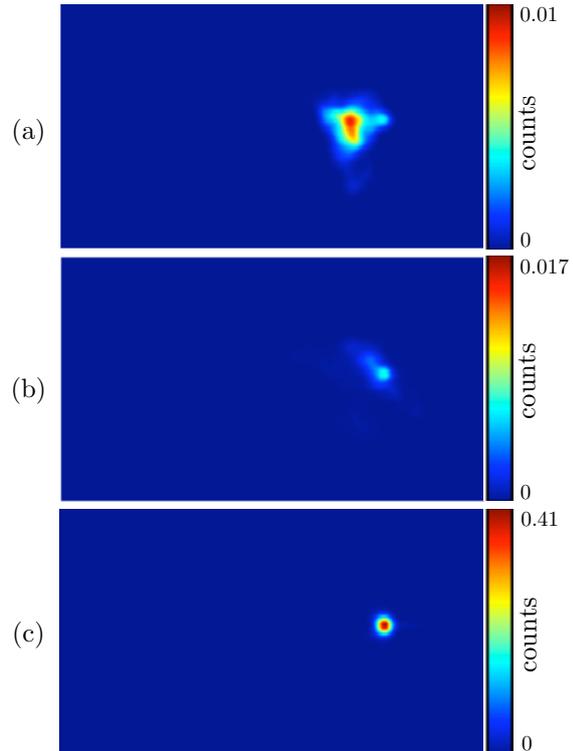
Cross-Calibration Evaluation Results: In two test phases a number of ten transmission acquisition have been performed. According to an evaluation protocol (cf. section 2.3.1) the stability and the correctness of the cross-calibration method have been evaluated.

	Evaluation Method	# Acq.	Criterion	Value
(1)	Stability/Validity (stability/matrix validity)	5	# of dropouts or invalid matrices – $\det(T_{cc}) \neq 1$	0
(2)	Calibration Plausibility (comparison calibrated/uncalibrated)	5	# incorrect calibrations	0

In the second test phase the cross-calibration method has been evaluated in general. In five test acquisition a radioactive point source has been placed within the FOV of the PET scanner. By constantly moving the point source an image blurring has been achieved and an applied motion compensation ought to correct the resulting motion artefacts. On the comparison of a fully motion compensated image with an image that has been motion compensated without any cross-calibration, the plausibility of the cross-calibration method have been demonstrated. The result of such a comparison is shown in figure 3.4.

Figure 3.4:

Coronal images of the results of the cross-calibration evaluation: All images are based on the same acquisition where during the list-mode acquisition the point source has been constantly moved. At the top row (a), the typical image blurring due to the moved point source is visible. In the centre image (b), the list-mode data have been motion compensated but with a disabled cross-calibration (unity matrix used). At the bottom row (c), the motion data has been fully calibrated to the PET coordinate system and then used during the event-driven motion compensation. The reduced blurring and the higher contrast in image (c) shows, that the calculated cross-calibration matrix calibrates the two coordinate systems accurately.



The results of this qualitative comparison on the calibrated and uncalibrated images show, that due to the application of the cross-calibration matrix the motion compensated data is an improvement. Compared to the uncalibrated image, the structure of the point source becomes visible again. For the uncalibrated image motion artefacts are, however, still visible.

3.2.2 Motion Target

After evaluation of the individual motion targets introduced in section 2.3.2, the results of the motion transferability tests has been statistically analysed and graphically visualised. In addition, the patient compliance has been rated according to an interview of the volunteer person performing the test acquisitions. In a third step, the clinical application of the motion target was rated corresponding to the amount of required time to attach the target on the patient head.

In the following pages the results of the test acquisitions are shown with respect to the individual motion phases (I-V) that were defined in section 2.3.2. In figure 3.5, 3.6 and 3.7, the results of a comparison between the reference target (Maxilla fixed target) and the individual head motion target are presented. Throughout these figures each of the six distinct motion parameters (translation/rotation) are displayed. The motion of the reference target is shown via a black line, whereas the red line illustrates the motion parameters of the used head motion target of the same acquisition.

Evaluation Summary

The raw motion parameters of the evaluation figures (3.5-3.7) provide information on the actual correlation of the sampled motion of the Maxilla fixed reference target with the head target. While the difference in one of the motion parameters of these comparison plots becomes visible quite easily, a general classification of the magnitude of the correlation between the two targets is, however, not easily possible.

To allow for a better assessment of the correlation, the distance between the two targets during the test acquisitions have been calculated and plotted over time. Figure 3.8 summarises the results of this comparison, which would show a flat horizontal line if the correlation between the two targets were ideal. This case would imply a Euclidean distance variance Δd of zero. Therefore, a minimal Δd suggests a good motion transferability of a target, bearing in mind that the Maxilla fixed reference target is considered a *gold standard* in terms of motion transferability.

Table 3.5:

Motion Target Evaluation Results: The three different motion targets have been evaluated according to three different criteria. During test acquisitions the distance between the motion target and a reference target has been monitored over time (Δd). In addition, the patient compliance and the effort to attach and remove the target have been rated.

	Motion Target Type	Δd	Patient Compliance	Clinical Requirements
(1)	Ski Goggles	4.2 mm	3 (Satisfactory)	2 (Good)
(2)	Os frontale	2.3 mm	5 (Fail)	4 (Barely passing)
(3)	Safety Goggles	0.6 mm	2 (Good)	2 (Good)

In addition to the quantitative comparison of the motion transferability, the targets have been rated according to different criteria. After the test acquisition the volunteer has been asked to rate each target between scores 1 (Excellent) and 5 (Fail). In addition, the effort for the attachment and removal of the target has been rated. Table 3.5 summarises the results of these evaluations.

Ski Goggles Target vs. Reference Target

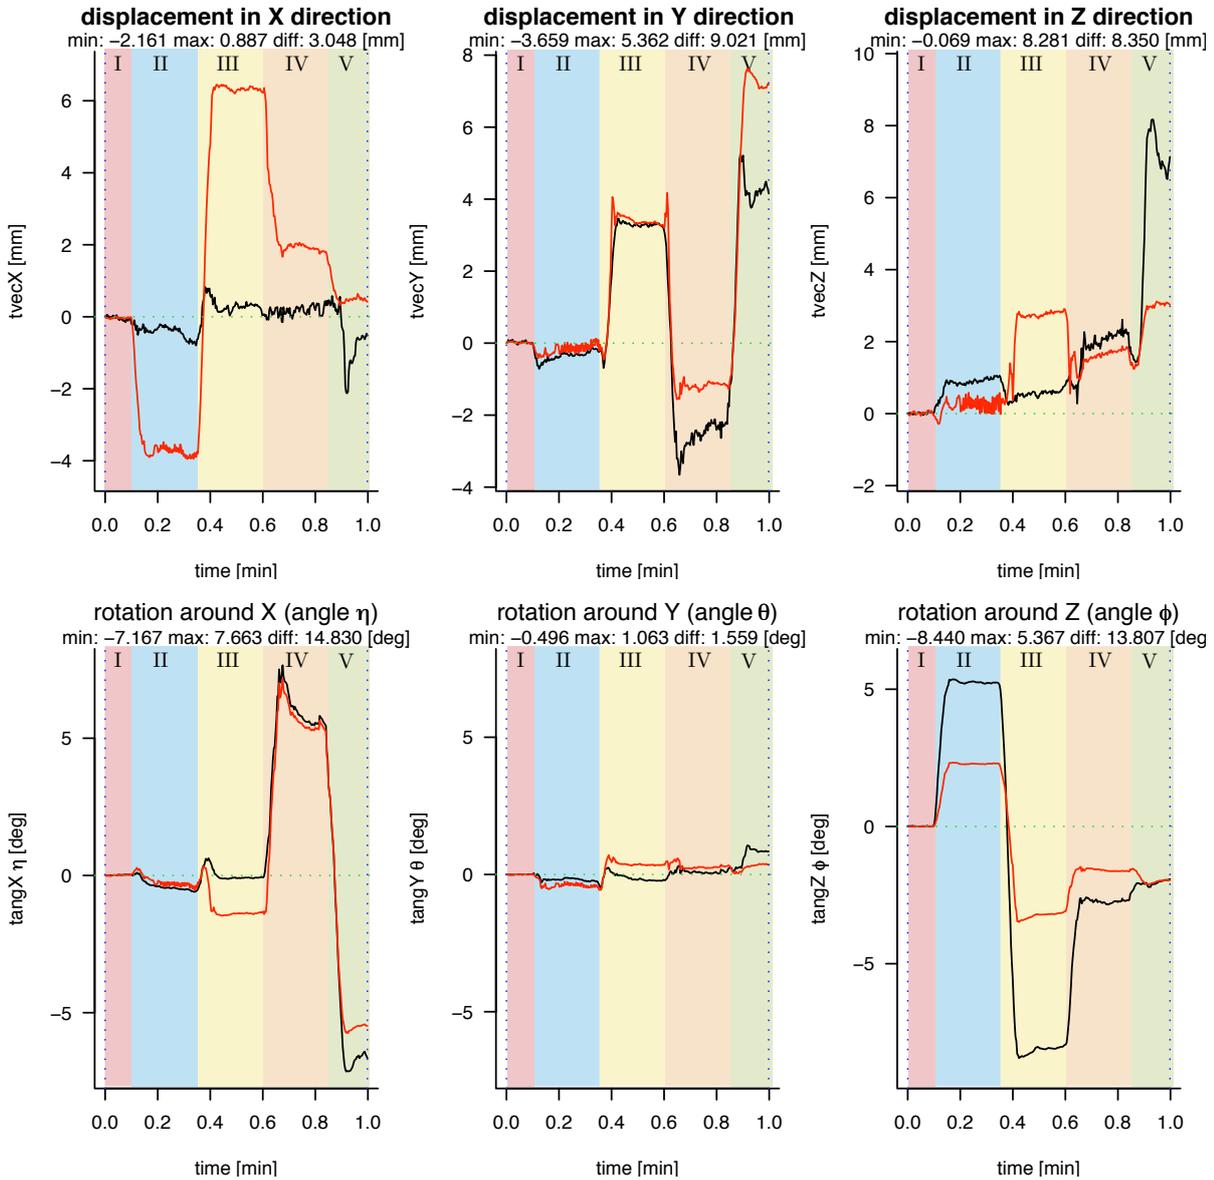
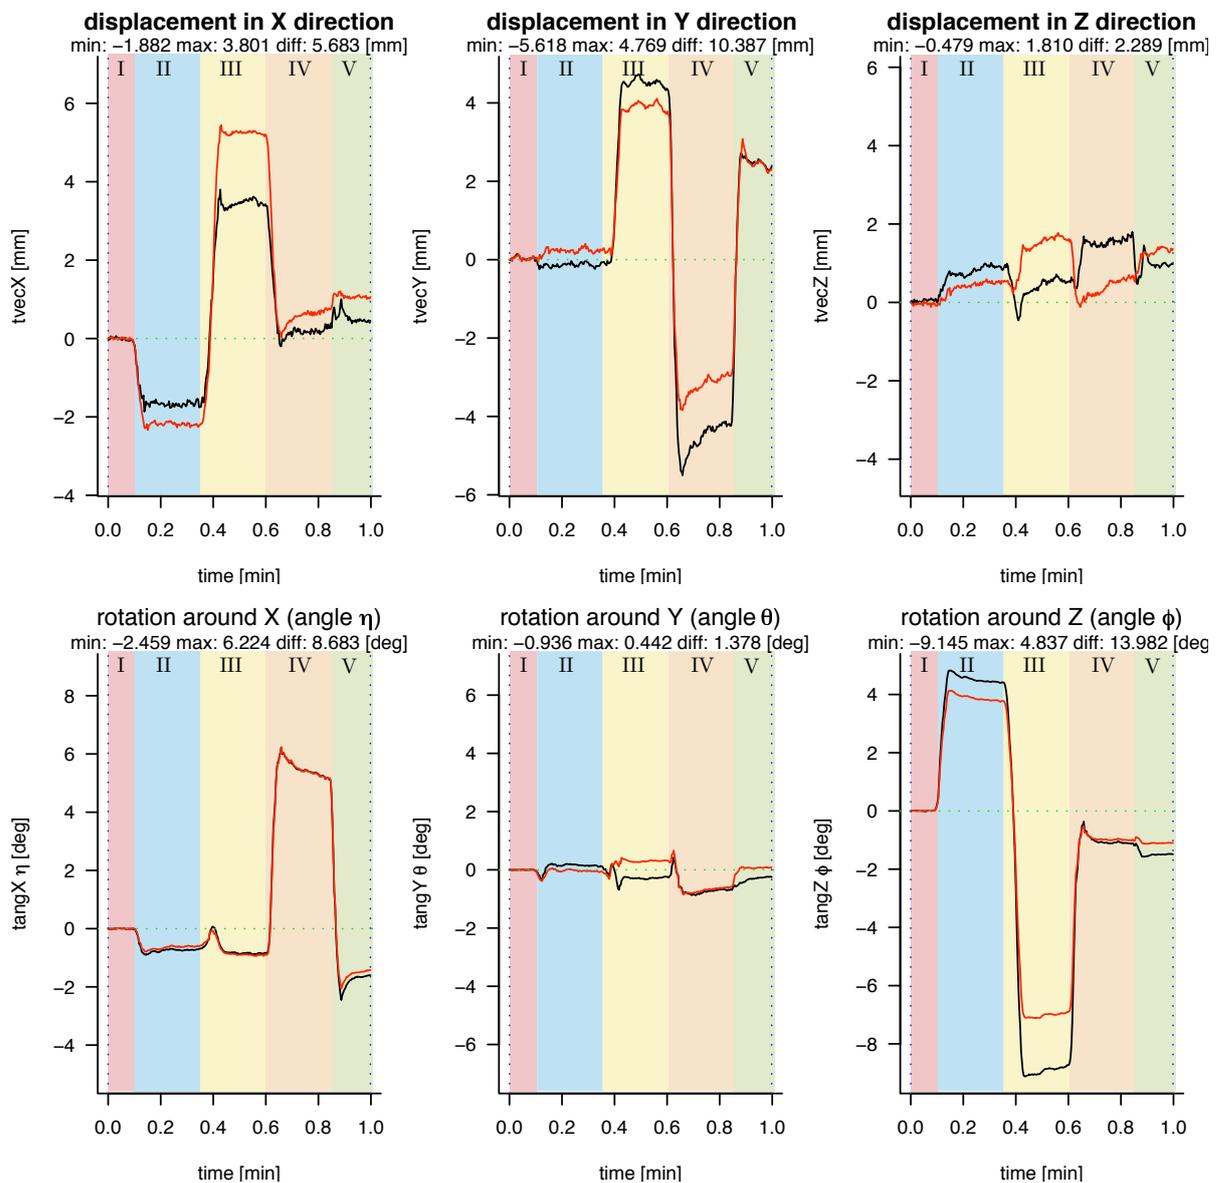


Figure 3.5:

Results of the ski goggles head target evaluation: All six motion parameters of the sampled motion are shown in individual plots. Here the motion of the *ski goggles head target* (red line) and the *Maxilla fixed reference target* (black line) are shown. In addition, the plots are shown in respect to the different movement phases (I-V) during the acquisition.

Os frontale Target vs. Reference Target

**Figure 3.6:**

Results of the *Os frontale* head target evaluation: All six motion parameters of the sampled motion are shown in individual plots. Here the motion of the *Os frontale* head target (red line) and the *Maxilla fixed reference target* (black line) are shown. In addition, the plots are shown in respect to the different movement phases (I-V) during the acquisition.

Safety Goggles Target vs. Reference Target

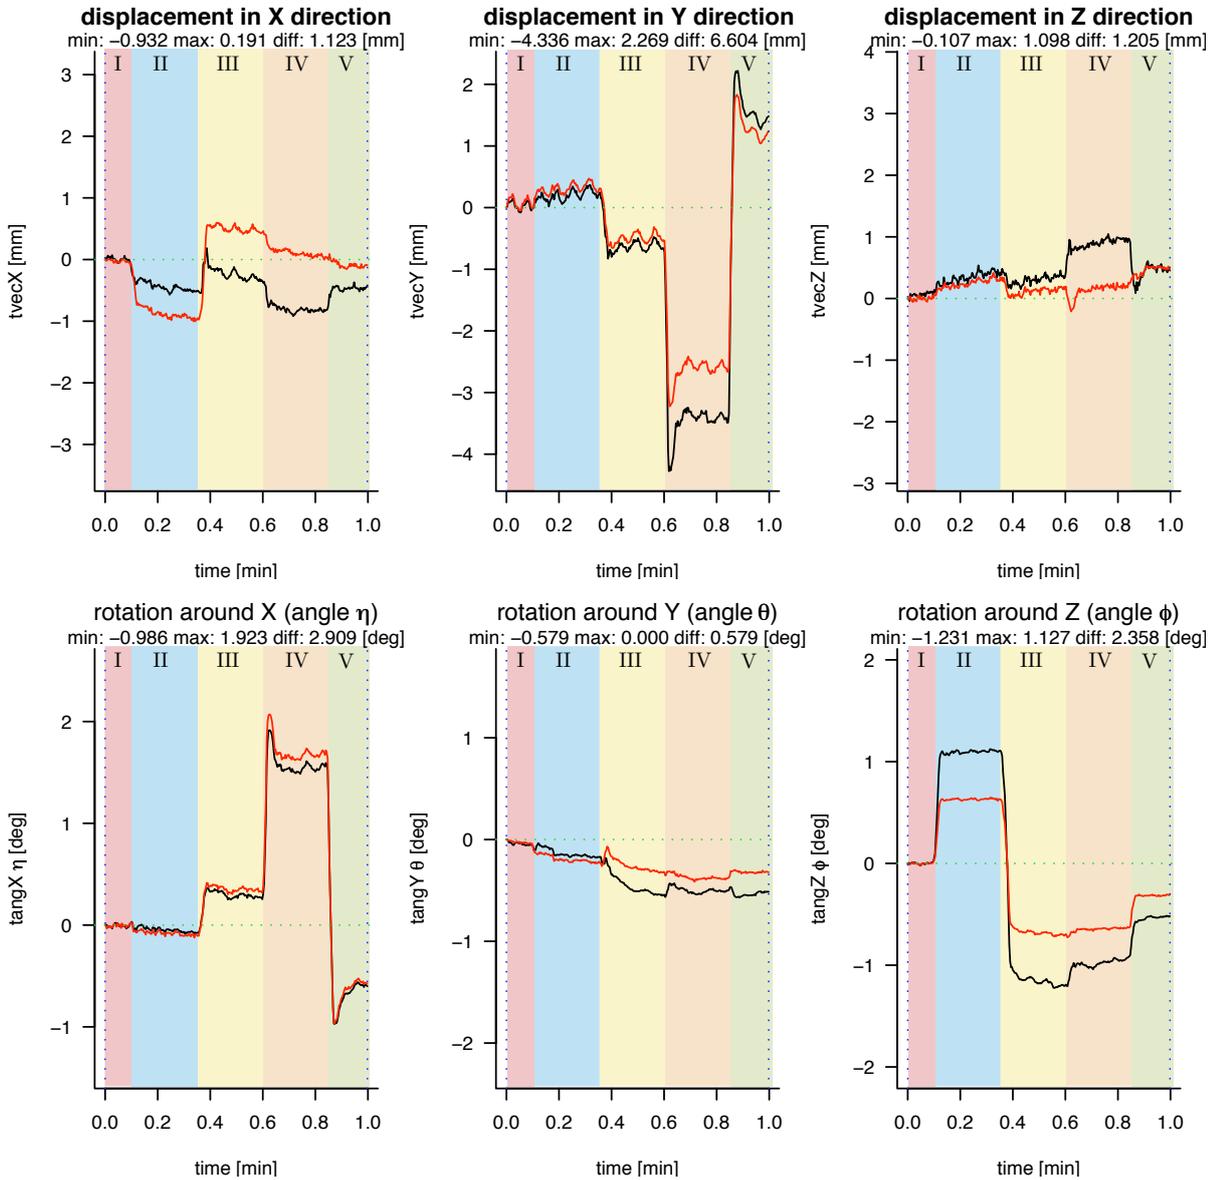
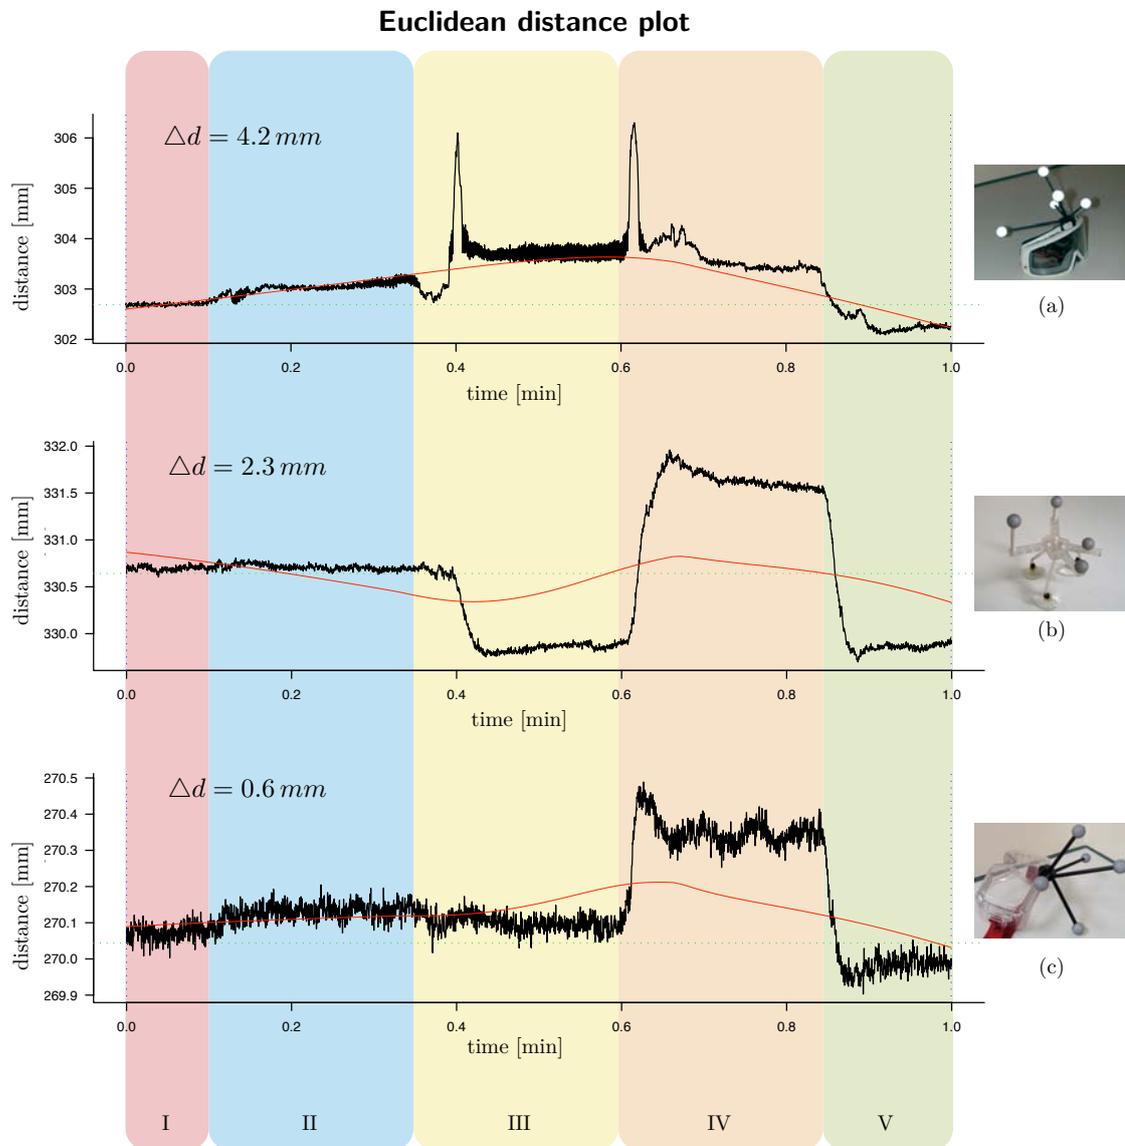


Figure 3.7:

Results of the safety goggles head target evaluation: All six motion parameters of the sampled motion are shown in individual plots. Here the motion of the *safety goggles head target* (red line) and the *Maxilla fixed reference target* (black line) are shown. In addition, the plots are shown in respect to the different movement phases (I-V) during the acquisition.

**Figure 3.8:**

Results of the Euclidean distance-based motion target evaluation: Test measurements were subdivided into five distinct motion phases (I-V) (cf. section 2.3.2). Here the Euclidean distance between the attached reference target (Maxilla fixed) and the individual motion target (a-c) over time is plotted (black line). In addition, a red line shows the general trend of the motion difference. At the top row (a), the results of the *ski goggles-based target* are shown, at the centre (b) the results of the *Os frontale fixed head target* and at the bottom row (c) the results of the evaluation of the *safety goggles-based motion target*. In addition, the maximum Euclidean distance variance Δd has been calculated. As shown, the variance of target c suggests the best motion transferability.

3.2.3 Motion Analysis

As a result of the motion analysis method outlined in section 2.3.3, a software package has been developed using the ‘R’ programming language (`analyze_trk`). This package not only manages the identification and output of ‘significant’ motion according to a fixed threshold. Moreover, as a fundamental part of this analysis, the generation of a clinical report document on patient motion has been developed. Part of this report document are plots and statistical data which are presented on three pages. The result of this report document, as discussed in section 2.3.3, will be illustrated on example data in the following section.

Raw Motion Report Page

The raw motion parameters (translations, rotations) of the recorded motion are presented on the first page of the report document (cf. figure 3.9). The six motion parameters of the reference target, which is used to verify the general stability of the motion tracking, are shown at the top of the page (part a). Similar to that, the raw motion parameters, but converted into the PET coordinate system, of the patient target are shown at the bottom (part b). Both motion plots feature statistical outputs on the minimum and maximum value of each motion parameter, as well as an output of the difference between the maximum and minimum value.

Sphere Analysis Report Page

The second page of the report document summarises the data on the position-relative motion analysis as discussed in section 2.3.3. The statistical data of the sphere motion analysis are illustrated in three main parts on the page. The first part contains a range plot and distance frequency plots for the sphere motion data. An example of these plots is shown in figure 3.10.

Both plots are based on the actual distances of the intersection points on the virtual grid from their original position. In the range plot the distance of an intersection point at a specific time is plotted on the ordinate. Therefore, the grey area in the range plot defines the limits of all distances found on the virtual patient head. In addition, as discussed in section 2.3.3, the range plot also contains signs where motion has been found to be ‘significant’ according to a predefined threshold. The two frequency plots (red and blue) on the bottom of the figure illustrate the distribution of the distances within the percentiles of 10% to 90% for the ‘worst’ and ‘best’ point on the sphere. As shown, all plots also contain an output of absolute values so that the motion magnitude distribution can be read directly from these plots.

In the second part of the sphere analysis report page the distance fluctuation and mean distance is shown in tiled image maps. Figure 3.11 illustrates this in more detail on example data in plot (c) and (d). Together with anatomic labels the image maps illustrate the distribution of the distances found on the surface of the virtual sphere. The same figure also shows a screenshot of the interactive 3D sphere graphics display as discussed in section 2.3.3.

Motion Summary Report Page

On a third and final report page, all information of the motion magnitude analysis is summarised. As shown in figure 3.12 this is accomplished by a plain text document. The raw motion values

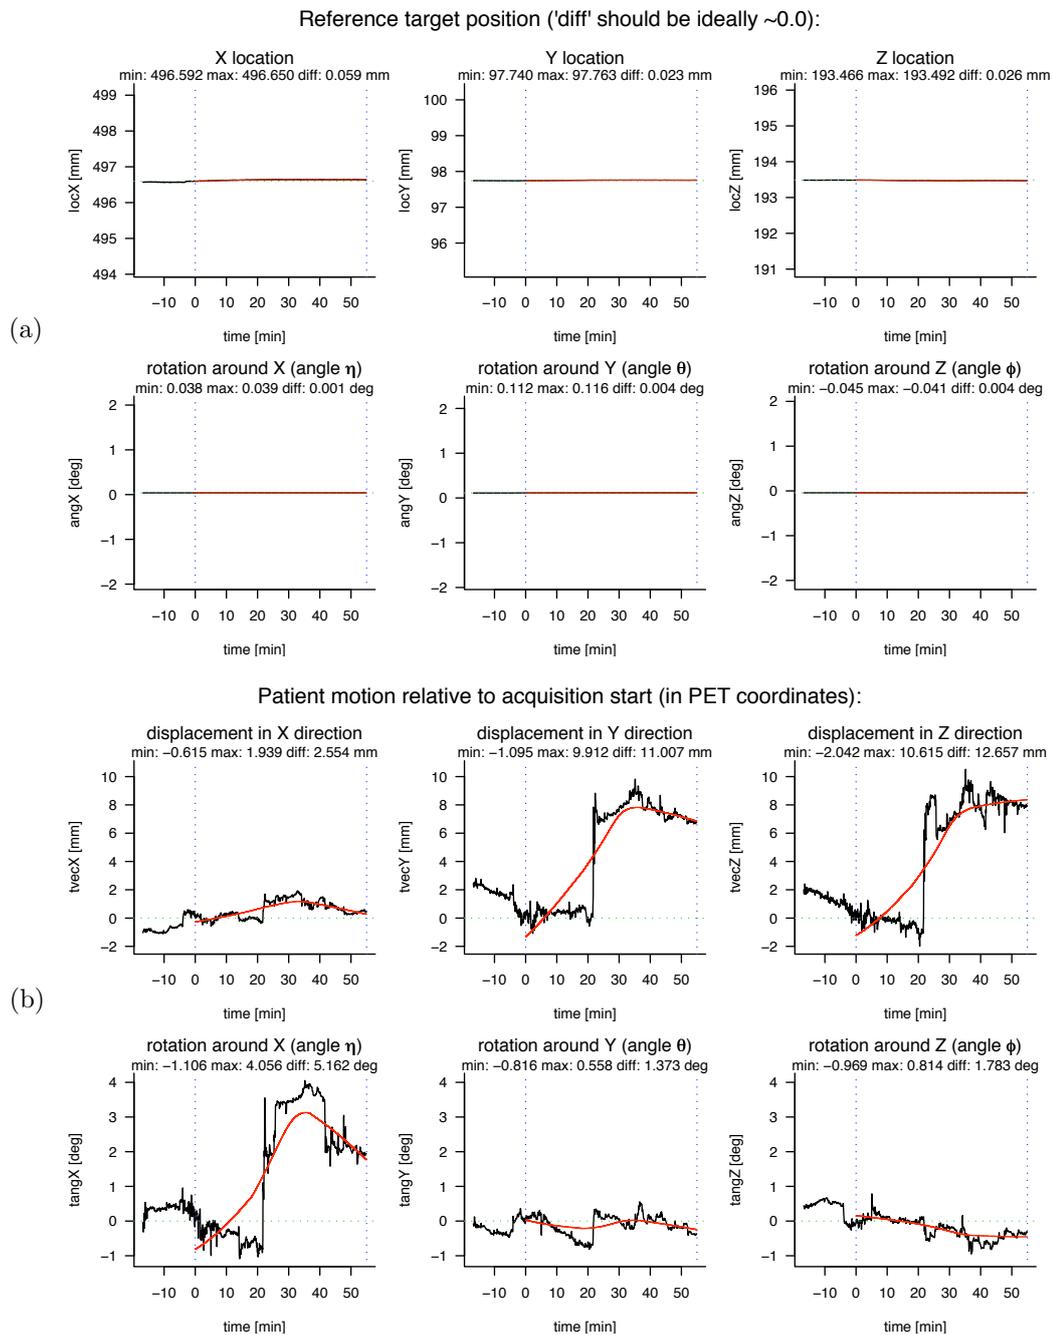


Figure 3.9:

Raw Motion Report Page: Example plots of the raw motion of a reference target (a) and the patient target are shown (b). These plots are part of a report document generated with each motion tracking. The motion information on the fixed reference target is included to verify the general stability of the motion tracking. In addition, the patient motion is illustrated with its *six degrees of freedom*, including statistical information on the minimum (min) and maximum (max) movement, as well as on the difference between max and min (diff) of each motion parameter.

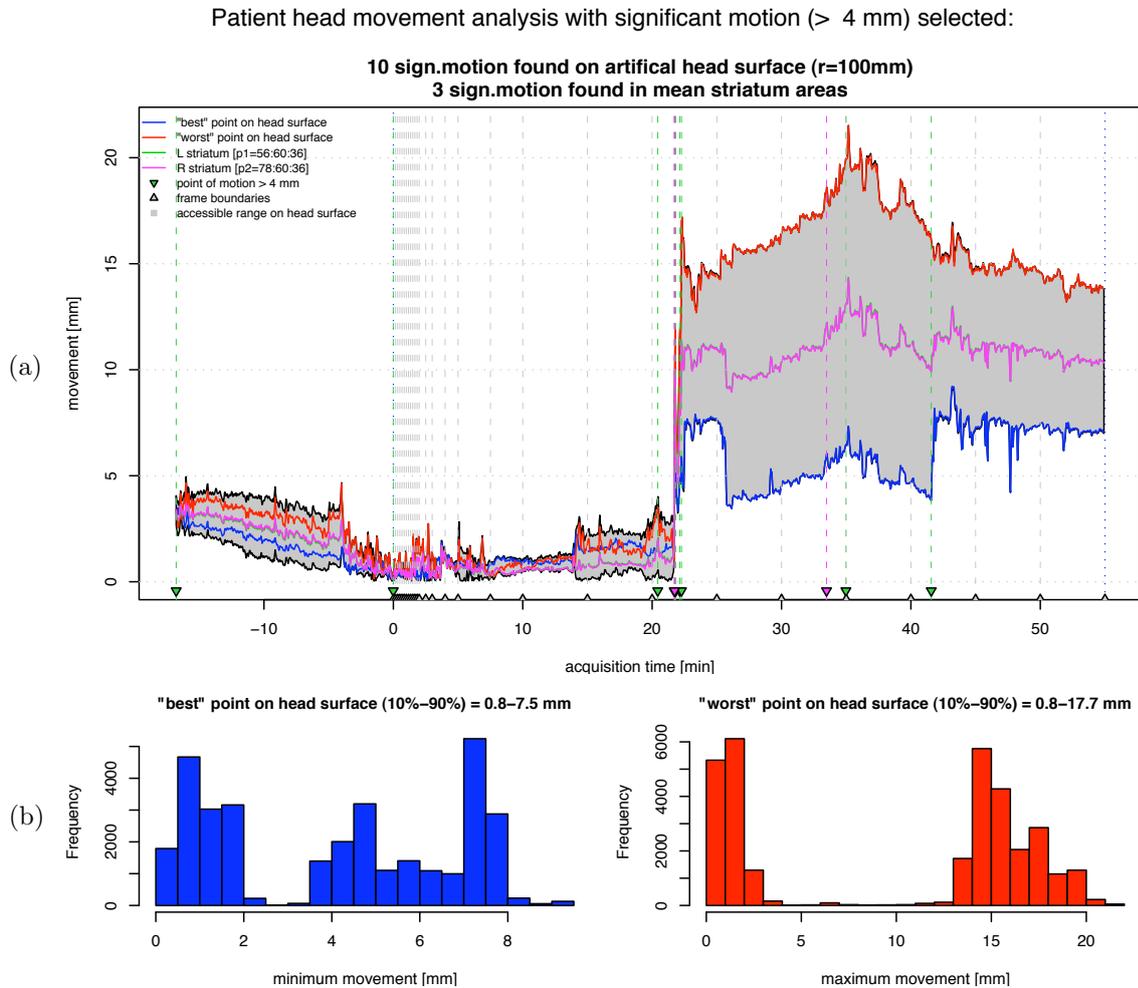


Figure 3.10:

Example of a motion analysis report page showing details on the sphere analysis: At the upper row (a) a range plot of the calculated Euclidean distances of the sphere analysis method of section 2.3.3 is shown. Each tracked grid point on the imaginary sphere corresponds to a single horizontal line within the gray area of the plot. In addition, the red line corresponds to the coordinate on the sphere with the highest motion, whereas the blue line corresponds to the point on the sphere with the lowest one. In addition, green vertical lines and small triangles mark 'significant' motions in time relative to the set threshold. In the lower row (b), frequency plots on motion of the 'best' (blue) and 'worst' (red) point on the sphere are shown.

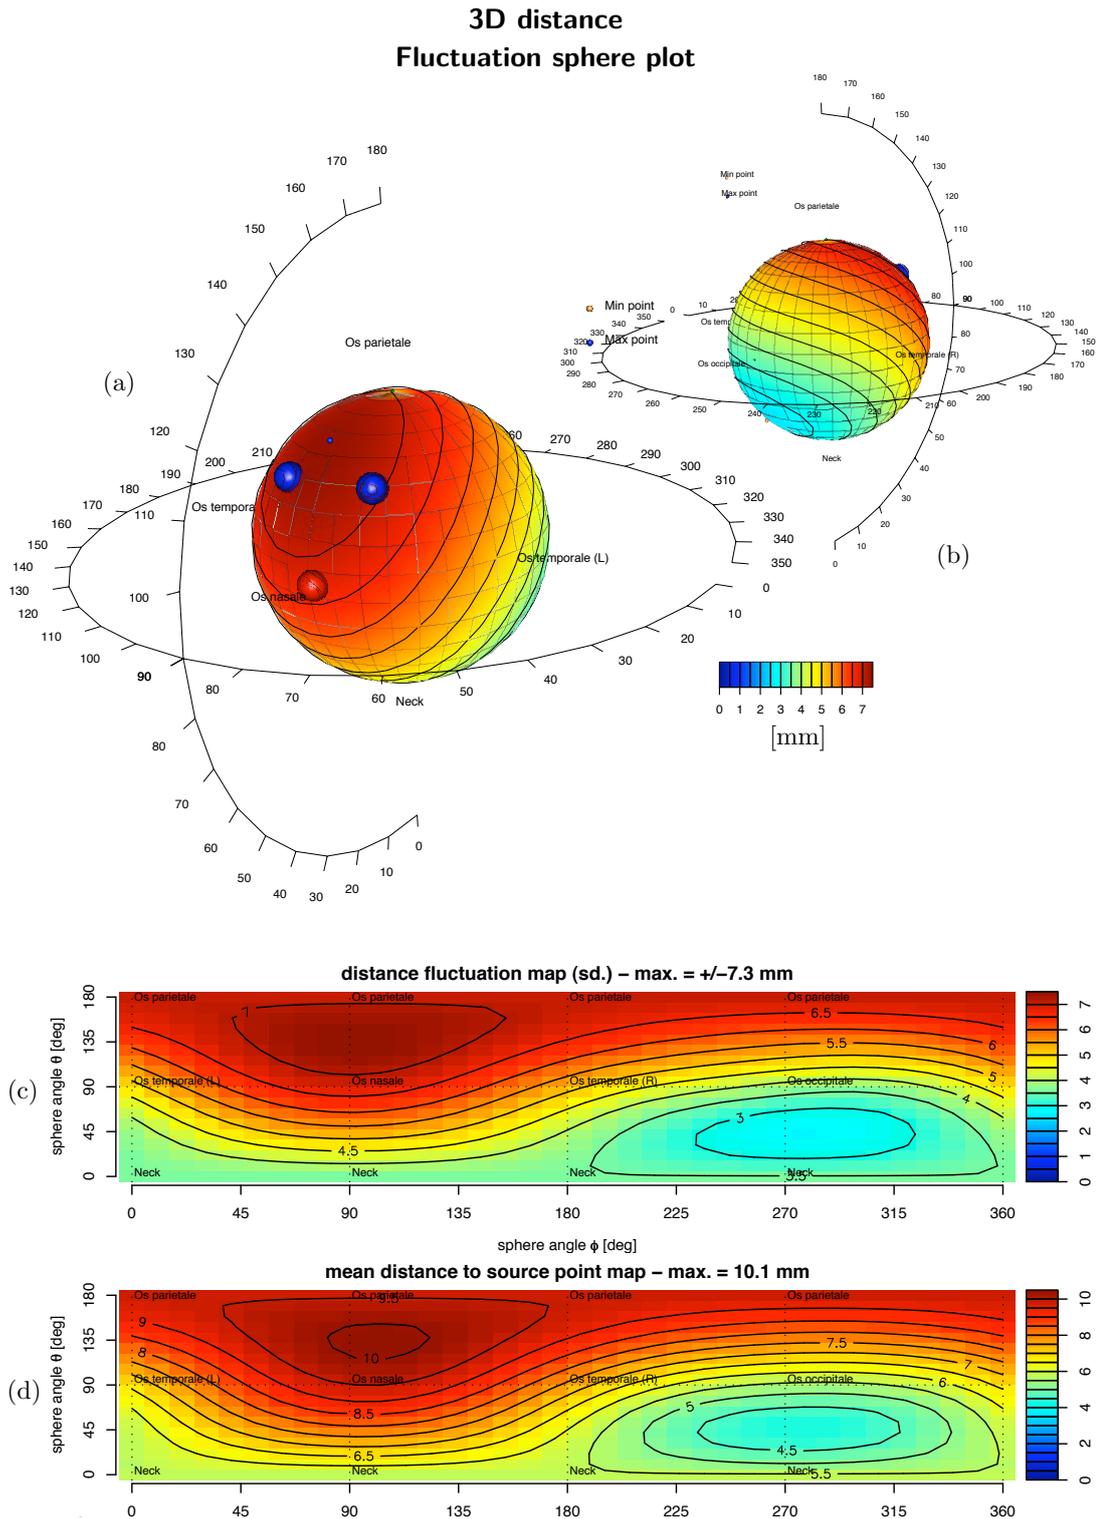
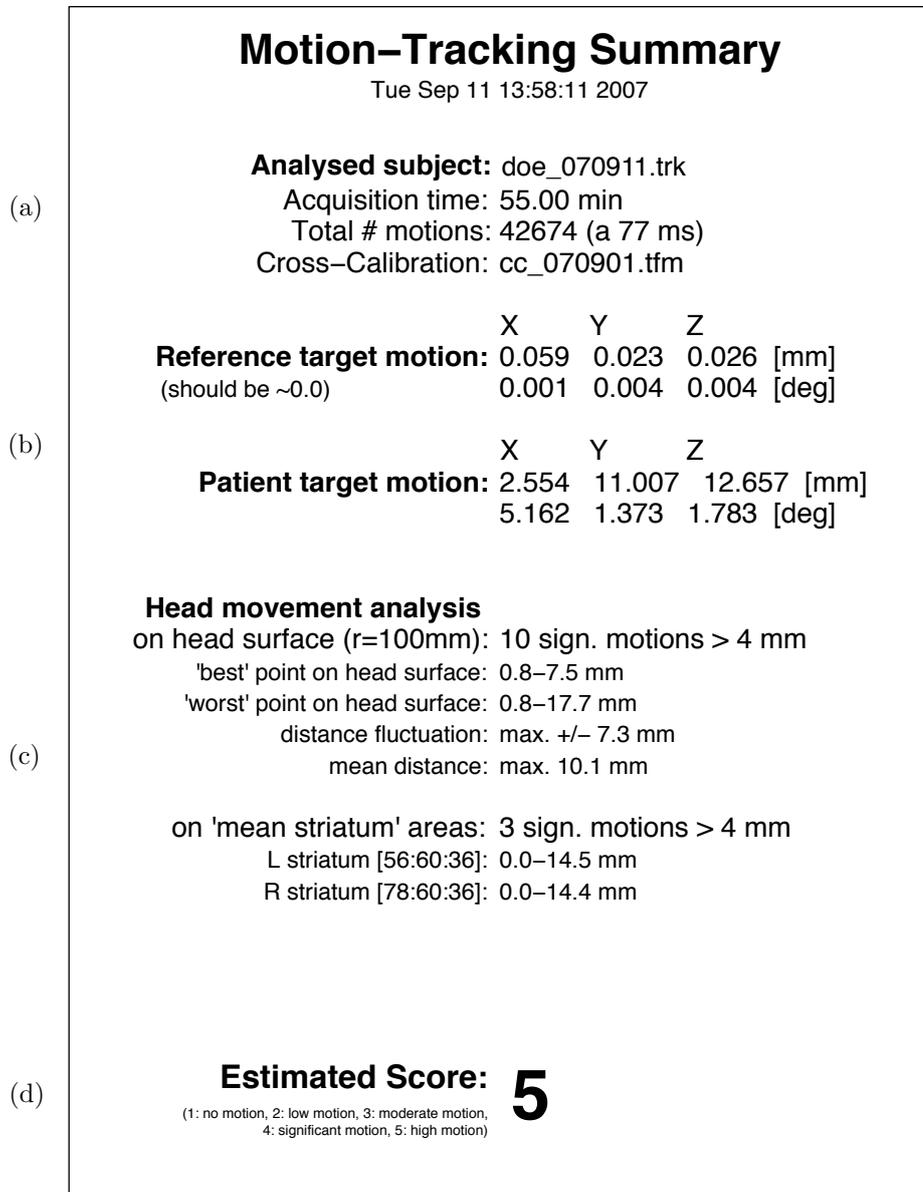


Figure 3.11:

Detailed sphere analysis report page: In accordance to the motion analysis shown in figure 3.10, additional plots on motion intensive areas on the imaginary sphere are generated. Plot (a) and (b) show screenshots of the front and back side of an interactive 3D sphere plot with the sphere surface coloured according to the amount of distance fluctuation of each point on the sphere. As a part of the standard report document, plot (c) shows the same distance fluctuation information, but via a 2D image map. Furthermore, an additional plot (d) is also part of the report and illustrates the mean distance of each sphere surface point from its origin at acquisition start time.

**Figure 3.12:**

Motion Tracking Summary page: As a summary of the results of the motion analysis (cf. figure 3.9-3.11) the report document contains a final page with summed information on the identified patient motion. Apart from general patient information in part (a), the raw motion parameters of the stability target as well as the parameters on the patient target are summed up in part (b) of the document. Furthermore, the results, such as the values on the 'best' and 'worst' point of the sphere analysis are shown in part (c). As a final summary of this information a 'motion score' value is also part of the summary page.

of the reference and head motion target, and the statistical parameters of the sphere analysis are grouped into four different sections of this page. As already explained in section 2.3.3, this summary page also contains the output of a ‘motion score’. This score has been added to facilitate the assessment of the motion magnitude, thus not having to go into details of all the different statistical parameters of the analysis. Based on the distance distribution of the ‘worst’ point on the imaginary head surface, the score divides the magnitude of motion into five values. The relation of the score value and the distance distribution is shown in table 2.2.

3.3 Event-Driven Motion Compensation

Apart from a general speed optimisation of the event-driven motion compensation algorithms, three main parts of the motion compensation have been improved. In the following three sections the results of the implementation and optimisation of the *normalisation correction*, *LOR discretisation correction* as well as of the *Out-of-FOV correction* are shown.

3.3.1 Normalisation Correction

After having developed and implemented the required methods to load the normalisation data from a standard ECAT normalisation file (cf. section 2.4.1), ten test measurements have been performed. To verify the correct calculation of the normalisation factors per LOR, the test measurements have been performed in list-mode and normalised at different phases of the image calculation process.

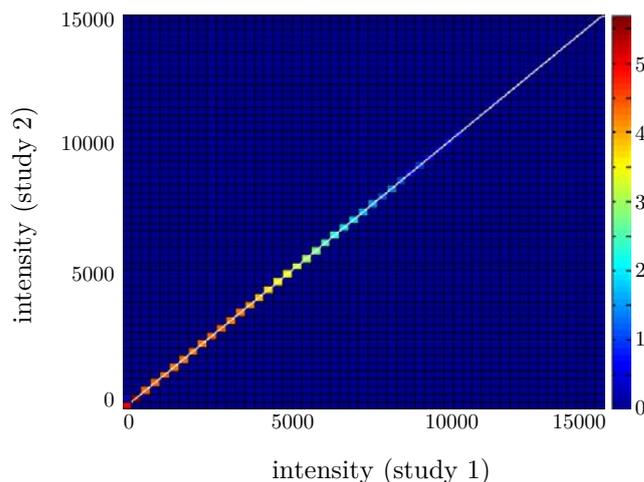


Figure 3.13:

Intensity correlation plot of two normalised data sets: One data set was normalised during the image reconstruction (study 1) and the other one was already normalised during the sinogram creation (study 2). This plot shows, that our normalisation correction is able to incorporate the normalisation data accurately. In fact, the quantitative output shows that the two images have a difference in the voxel intensities of lower than 0.1%.

An example of one voxel-based intensity correlation plot of the ten test acquisitions is shown in figure 3.13. The quantitative analysis of two differently normalised data sets shows, that the deviation in the differently normalised data sets of the test acquisitions is always lower than 0.1%. This is in accordance with the same type of voxel-based intensity correlation analysis, as shown in section 3.1.3.

3.3.2 LOR Discretisation Correction

An artificial list-mode data stream has been generated and used for the evaluation of the optimised LOR discretisation correction method as shown in section 2.4.2.

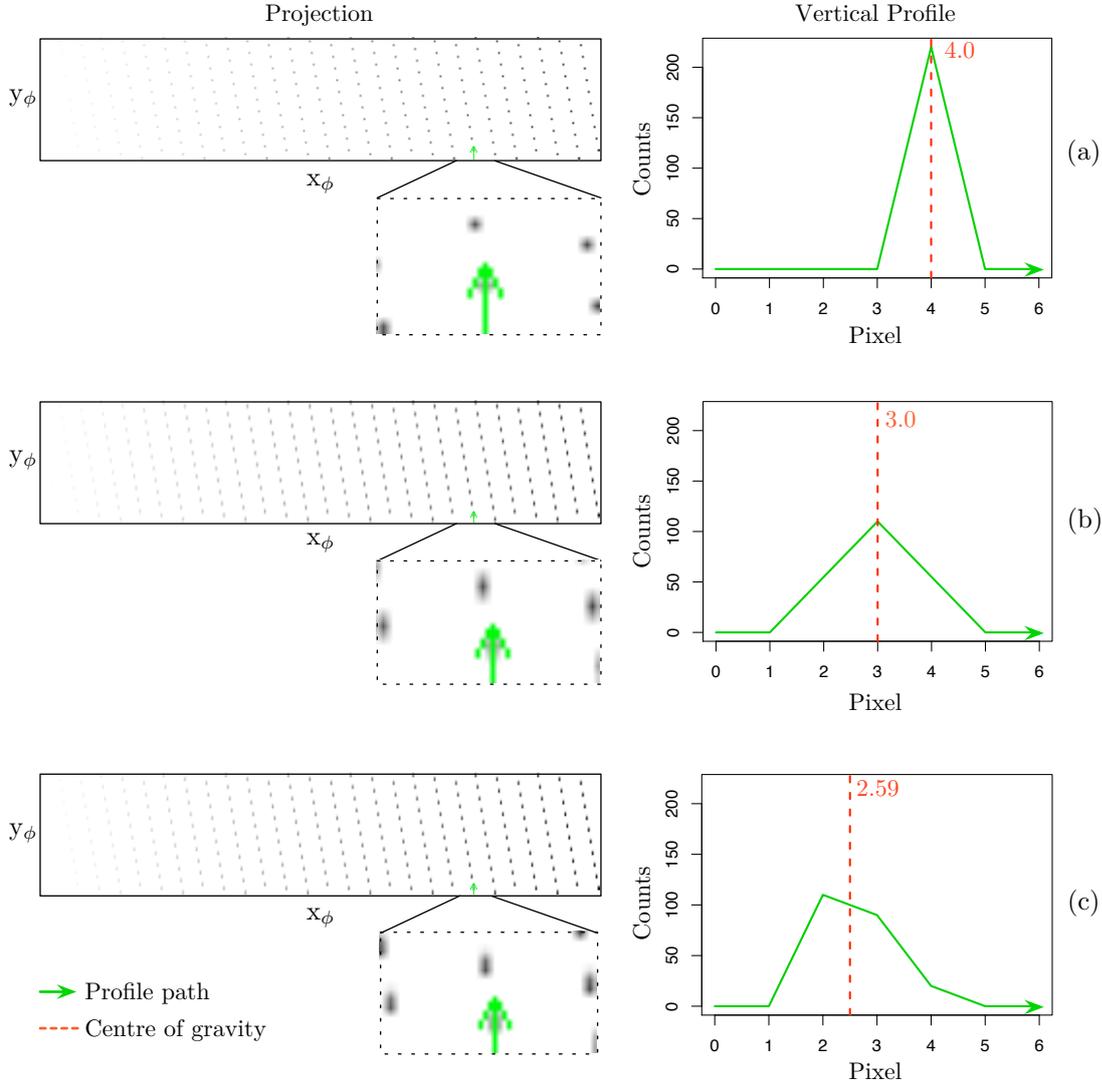


Figure 3.14:

Results of the evaluation of the LOR discretisation correction: On the left-hand side of the figure projection images with a zoomed area are shown. In row (a) the projection of the unmoved data (cf. figure 2.24) with a vertical profile (green line and arrow) through one sinogram bin is shown. In row (b) and (c) the same list-mode data set has been moved along the Z-axis (y_ϕ) so that the reoriented LORs overlap certain bins. In row (b) the image and the profile is shown where the LOR discretisation has been disabled. In contrast, row (c) illustrates that our LOR discretisation optimisation is able to apply better weighting factors to the amount of overlap found on the detector surface (here visible due to the slightly unbalanced profile).

After having used the event-driven motion compensation with both, the optimised and non-optimised discretisation correction, the data has been sorted into a standard sinogram. To classify if the optimised LOR discretisation is able to improve the motion compensation, profiles have been drawn in the sinogram along the direction of the performed motion (y_ϕ). The results of these user defined profiles are shown in figure 3.14.

In a first test the unmoved and patterned list-mode data have been sorted into a sinogram and analysed before actually applying the motion compensation. By drawing a profile along one of the filled bins we selected one bin in the projection view of the sinogram and were able to identify 220 counts at pixel number 4 along y_ϕ (cf. part (a) of figure 3.14). After having applied the motion compensation along y_ϕ by $1.5\times$ of the voxel width, we generated two additional data sets. One data set has been generated with the non-optimised LOR discretisation correction. The profile analysis of this data set, as shown in part (b) of the same figure, illustrates the count distribution of the same sinogram bin. The 220 counts, which had been put into one single pixel in the unmoved data set, have been distributed throughout 3 pixels (55 counts in pixel 2, 110 counts in pixel 3 and 55 counts in pixel 4). According to these values the centre of the count distribution is located at pixel 3. In contrast, as shown in part (c) of the same figure, the distribution of the counts is different in the data set where the optimised LOR discretisation has been activated. In this data set the discretisation correction has been applied with a weighting raster of 10×10 (cf. figure 2.22). As shown in the profile plot, the 220 counts in the unmoved data set are now again distributed throughout 3 pixels along y_ϕ . However, due to the additional weighting of our optimised LOR discretisation correction, the count distribution was found to be 110 counts in pixel 2, 90 counts in pixel 3 and 20 counts in pixel 4. According to this count distribution, the centre of the count rate weighted average is located at pixel position 2.59. This shows a better concordance with the applied movement of $1.5\times$ of the voxel width.

3.3.3 Out-of-FOV Correction

The optimised Out-of-FOV correction method, as shown in section 2.4.3, requires not only changes to the event-driven motion compensation algorithms, but also to the motion analysis routines discussed in section 2.3.3.

Table 3.6:

OFOV scaling factor distribution: The scaling factor calculation has been statistically analysed. For six selected studies (cf. table 2.3), the distribution of the OFOV factors for the non-optimised and optimised OFOV correction are shown. In addition, in case the optimised image also showed a qualitative improvement this has been marked in the table accordingly.

	non-optimised			optimised			qualitative improvement
	mean \pm std.dev	min.	max.	mean \pm std.dev	min.	max.	
(1)	26.81 \pm 73.38	1.00	1123.60	1.22 \pm 0.41	1.00	4.25	✓
(2)	6.86 \pm 15.74	1.00	595.24	1.56 \pm 0.37	1.00	2.51	✓
(3)	4.10 \pm 17.32	1.00	375.00	1.18 \pm 0.19	1.00	1.46	✓
(4)	9.63 \pm 31.93	1.00	321.54	1.16 \pm 0.17	1.00	1.87	✓
(5)	16.58 \pm 31.69	1.01	92.99	2.23 \pm 1.75	1.01	5.38	✓
(6)	7.21 \pm 25.58	1.04	194.93	2.29 \pm 1.95	1.00	19.92	✓

To evaluate the performance of the optimised OFOV correction method, the data of six PET acquisitions have been analysed. As shown in the evaluation part of section 2.4.3, these examinations were performed in list-mode with the motion tracking system monitoring the patient's head motion during the acquisition. During the processing of the event-driven motion compensation the calculated OFOV factors have been stored in a data file and then analysed

statistically. Table 3.6 summarises the distribution of the calculated OFOV factors in the six example studies with statistical parameters of the distribution of these values per study.

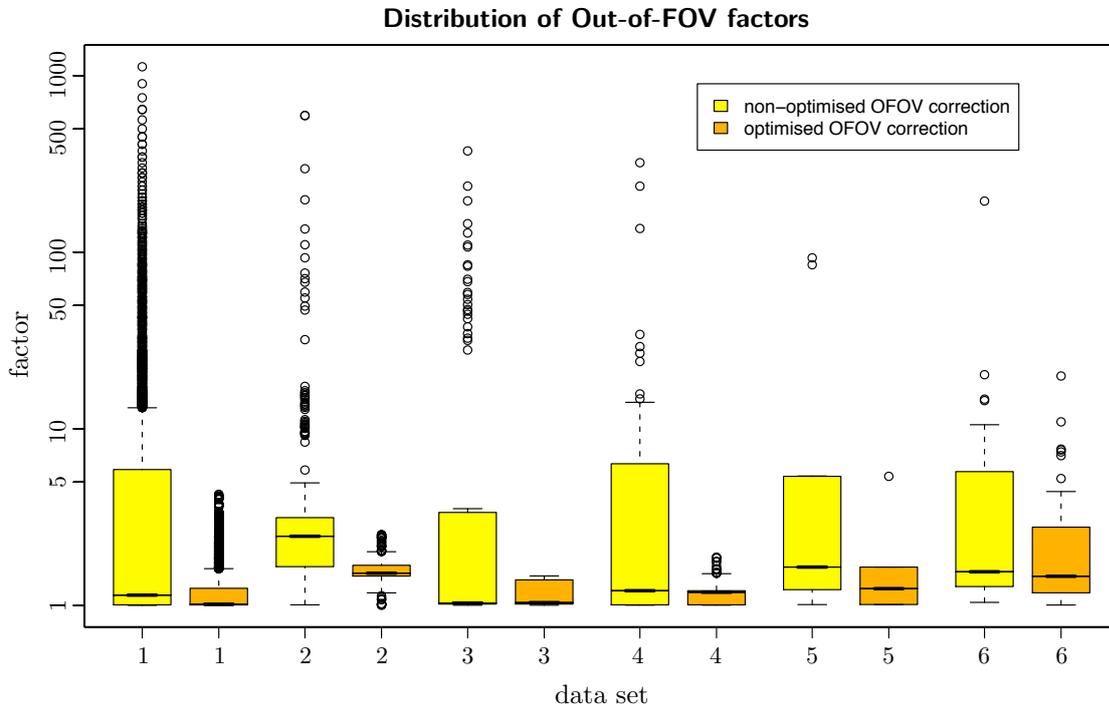


Figure 3.15:

Box plots of a quantitative analysis on the OFOV factors of six data sets: The distribution of the scaling factors in the standard OFOV corrected data sets (yellow) shows that they are generally higher than those found in the optimised OFOV corrected versions (orange). The outliers in the box plots of the standard OFOV corrected data sets support the qualitative observation of a higher likeliness of image artefacts. In contrast, the reduced amount of outliers in the optimised data sets suggest an improvement.

In addition, the statistical parameters extracted from the analysis of the OFOV factors have been illustrated via box plots. The output of these box plots on the six studies is shown in figure 3.15. The result of this analysis shows that the calculated scaling factors range from 1.0 to 1123.60 for the data sets that have been motion compensated with the activated non-optimised OFOV correction method. In contrast, for the optimised OFOV correction method the factors were found to range between 1.0 and 19.92. Also the number of outliers in the box plot analysis of the non-optimised vs. optimised data sets shows a reduction.

Furthermore, to evaluate the qualitative improvement of the images as a result of the optimised OFOV processing, an image comparison has been performed. As shown in the rightmost column of table 3.6, the application of the optimised OFOV method has been rated a ‘success’ if a subjective comparison shows a reduction of image artefacts. In figure 3.16 an example of such a comparison (on study 2) is shown. It illustrates the successful reduction of image artefacts due to the application of the optimised OFOV correction method in one plane of the motion compensated image data set.

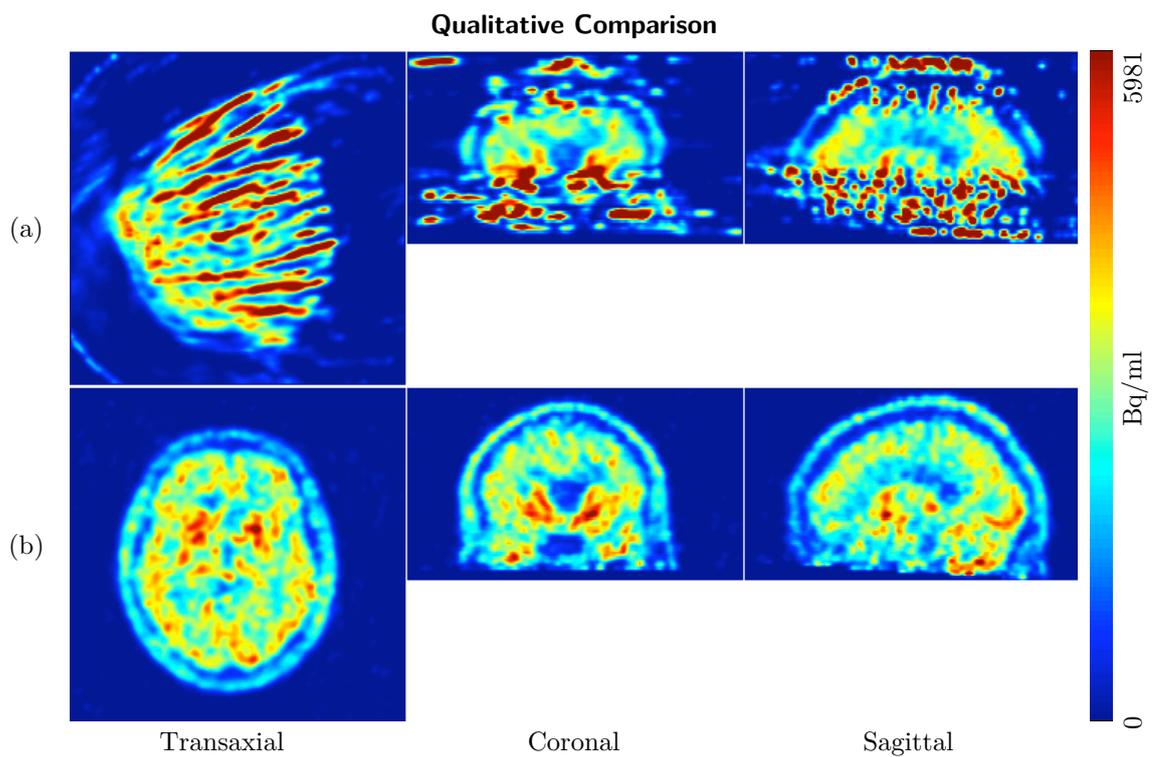


Figure 3.16:

Results of the optimisation of the Out-of-FOV correction method: Images on a single frame of the dynamic PET study number 2 are shown (5 min frame @ 30 min p.i. with 171 MBq of [^{18}F]DOPA). In row (a), image artefacts are visible due to the over-scaling of the standard OFOV correction. Row (b) shows the same data, but processed with the optimised OFOV minimising the count scaling and thus reducing image artefacts.

3.4 Clinical Integration

A number of six command-line-based software tools, as well as two graphical applications have been developed. As discussed in section 2.5, these tools are implementing the aforementioned methods from the acquisition of list-mode including the final image reconstruction of the motion compensated data. The main focus of these tools has been to implement these methods so that they are applicable for clinical use.

The two graphical user interfaces `lmacq` and `trkacq` are used to process the acquisition related tasks of the motion compensation. Figure 3.17 shows screenshots of these applications.

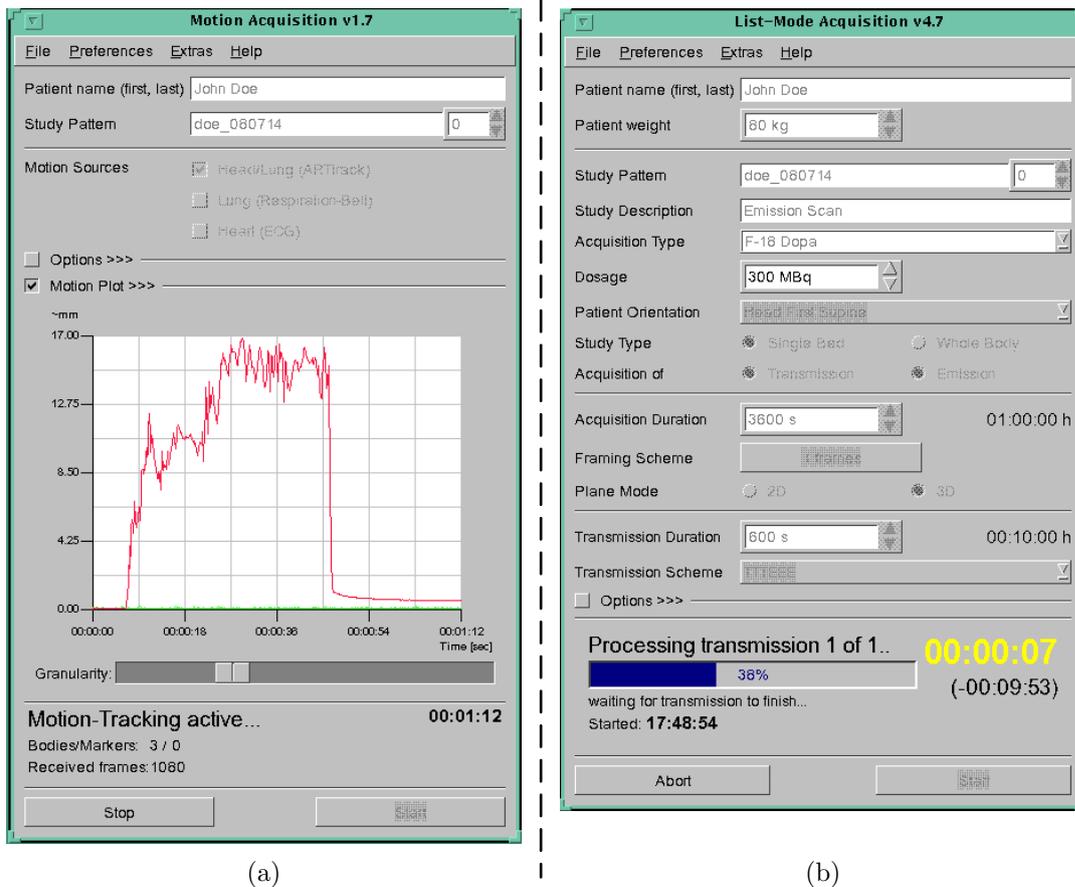


Figure 3.17:

Screenshots of the graphical user interfaces: On the left-hand side the `trkacq` interface is shown with a plot of the patient motion. During the acquisition, this plot is constantly updated so that motion of the patient becomes immediately visible. In addition, all patient-relevant and acquisition-relevant settings are configured through the `lmacq` user interface shown on the right-hand side. Upon acquisition start, commands to set up the HR⁺ PET scanner for list-mode processing are executed in the background. In addition, the progress of the list-mode acquisition is displayed and elements exist which allow for aborting of the acquisition.

Both of these interfaces provide graphical elements for the definition of acquisition related and patient related information. In addition, the `trkacq` application provides a live view of the motion where an estimated magnitude of the motion is plotted during a motion acquisition in progress. Both applications are started directly from a generic ECAT protocol block, as shown

in figure 3.18. This figure also demonstrates the process of how these applications receive the initial patient information from the ECAT console (black arrows).

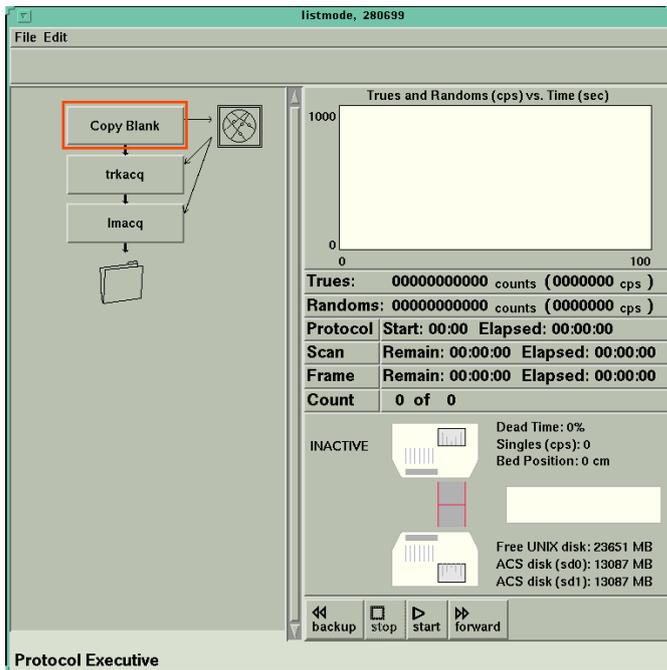


Figure 3.18:

Screenshot of the custom ECAT protocol: The graphical user interfaces `trkacq` and `lmacq` are executed from a custom ECAT protocol block. When the protocol is executed, the graphical interfaces shown in figure 3.17 are presented to the user and allow for the list-mode set up of the PET scanner as well as for the set up of the motion tracking device.

As soon as the list-mode acquisition is finished the coincidence data will be either sorted via `lmSorter` or motion compensated via `lmMover`. Both of these tools are command-line oriented and have been optimised for the operation on multiprocessor computer systems.

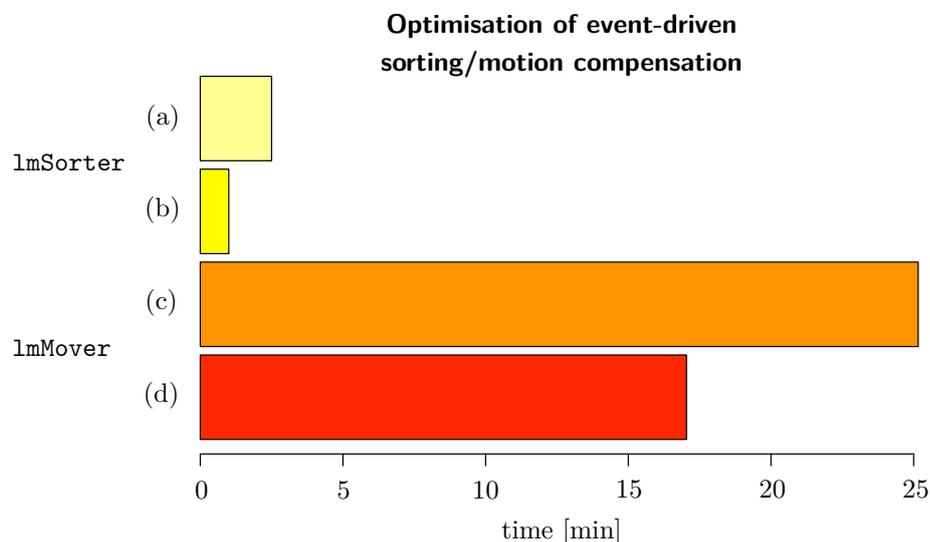


Figure 3.19:

Results of the Optimisation of list-mode sorting/motion compensation: After optimisation of the list-mode software, a speed comparison has been performed. Due to the multi-processor optimisation the sorting software (`lmSorter`) has been found to be sped up by $2.5\times$ – see (a) vs. (b). In comparison, the implementation of the event-driven motion compensation (`lmMover`) has been found to be sped up by $1.5\times$ as shown in the (c) vs. (d) bar plot.

As a result of this optimisation, the required time for sorting a list-mode data stream has been reduced. After the optimisation, the sorting of a typical list-mode study (via `lmSorter`)

of, e.g., ≈ 3 GB of data (55 min, 27 frames, 3D mode, 300 MBq, ^{18}F]DOPA) has been found to take $\approx 1:00$ min in comparison to $\approx 2:30$ min – a speed-up of 2.5 (on a 8×2.4 GHz multiprocessor system). Regarding the multiprocessor optimisation of the motion compensation tool (*lmMover*), the optimised implementation was found to process the data of the same list-mode study $\approx 1.5 \times$ faster (≈ 25 min vs. ≈ 17 min). Those results are summarised in figure 3.19.

```

langner @ uranus — ssh — bash — ttys018 — 103x19
ssh ssh
lmMover 0.8.7 - a multithreaded Qt4-based list-mode event mover
(Jun 19 2008) Copyright (c) 2006-2008 by Jens Langner / www.fzd.de

Applying Movement correction of list-mode data to 'lmMover_test.S'
Using 10 parallel processing units.

  Frame | CPU time/estim | REAL time/estim | ETA
16/27 ( 77%)| 2:27:48/ 3:13:12| 18:30/ 24:10| 5:40
17 [ 86.7%] M1 *****
18 [ 84.0%] M1 *****
19 [ 42.0%] M1 *****
20 [ 42.7%] M1 *****
21 [ 40.0%] M1 *****
22 [ 38.7%] M1 *****
23 [ 41.0%] M1 *****
24 [ 30.7%] M1 *****
25 [ 30.7%] M1 *****
26 [ 29.3%] M1 *****

```

Figure 3.20:

*Screenshot of the command-line interface (*lmMover*):* During execution, the progress of the event-driven motion compensation can be monitored. The above figure demonstrates the processing of an example data set. Here the parallel processing of 10 frames is shown. Together with progress information on each frame, the overall *estimated time of arrival* (ETA) is shown when the processing will be finished.

For the user interface of the *lmSorter* and *lmMover* application, a unified command-line interface has been developed. A screenshot of the execution of this interface is shown in figure 3.20. The live progress for each acquisition frame is illustrated by horizontal asterisks. In addition, the overall status of the processing is shown at the top with an output of the *estimated time of arrival* (ETA) when the processing is finished.

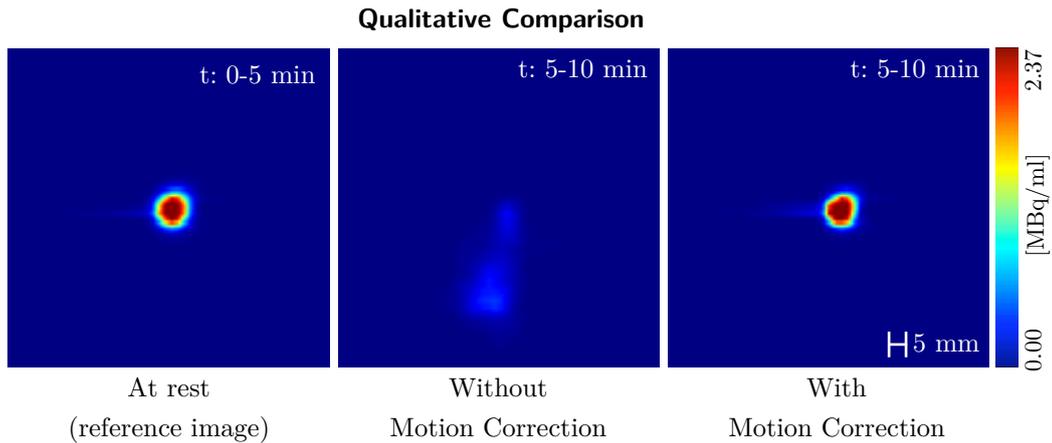
3.5 Clinical Evaluation

In two evaluations the qualitative and quantitative influence of the motion compensation in a clinical environment has been analysed. The following sections outline the results of the evaluations presented in section 2.6.

I. Point Source Evaluation

After having performed the point source test acquisition three images have been reconstructed for a qualitative evaluation of the motion compensation. These images are shown in figure 3.21.

The images demonstrate that the performed motion (cf. figure 2.31) has a severe influence on the quality of the PET data from the phase of the acquisition where the motion has taken place (5 min to 10 min). In addition, the corrected image on the right-hand side shows that the motion compensation is capable of restoring the data and allows to generate an image with no visible motion artefacts.

**Figure 3.21:**

Qualitative Point Source Evaluation Results: The images are generated from different phases of a test acquisition with a moving point source attached to a motion tracking target. On the left-hand side an image of the unmoved phase (0–5 min) of the acquisition is shown. In contrast, the centre image shows the typical blurring due to the motion performed in the last half (5–10 min) of the test acquisition. Finally, by comparing the left-hand and right-hand side images no apparent motion artefacts are visible anymore.

To support this qualitative result of the motion compensation evaluation, an additional quantitative analysis of the effective image resolution has been carried out. As discussed in section 2.6, the image resolution has been determined by using an already existing tool (`psfFit`) which is able to calculate the reconstructed image resolution (FWHM) for spherical objects. Table 3.7 lists the results of the application of this tool for the reference image and the motion compensated image.

Table 3.7:

Quantitative Point Source Evaluation Results: The results of an analytic determination of the image resolution shows that because of the motion compensation the reconstructed image resolution (FWHM) increases. However, considering the resolution of the PET scanner itself, the increase of the FWHM by ≈ 1 mm does not represent a severe reduction of the image resolution in general.

	Resolution (FWHM)	Std. Error
At rest (reference image)	5.46 mm	0.02
Motion Corrected	6.62 mm	0.01

II. Patient Data Evaluation

As shown in section 2.6, a qualitative and quantitative comparison of an uncorrected and motion compensated patient data set has been performed. Out of fifteen motion tracked patient examinations a single one has been selected. It serves for a detailed analysis of the accuracy of the motion compensation method regarding the qualitative, but also quantitative influence on PET data.

Firstly, the ability of the motion compensation to restore the image quality in case of motion has been analysed. Figure 3.22 shows a qualitative comparison of the uncorrected and motion compensated image of a selected [^{18}F]DOPA patient data set.

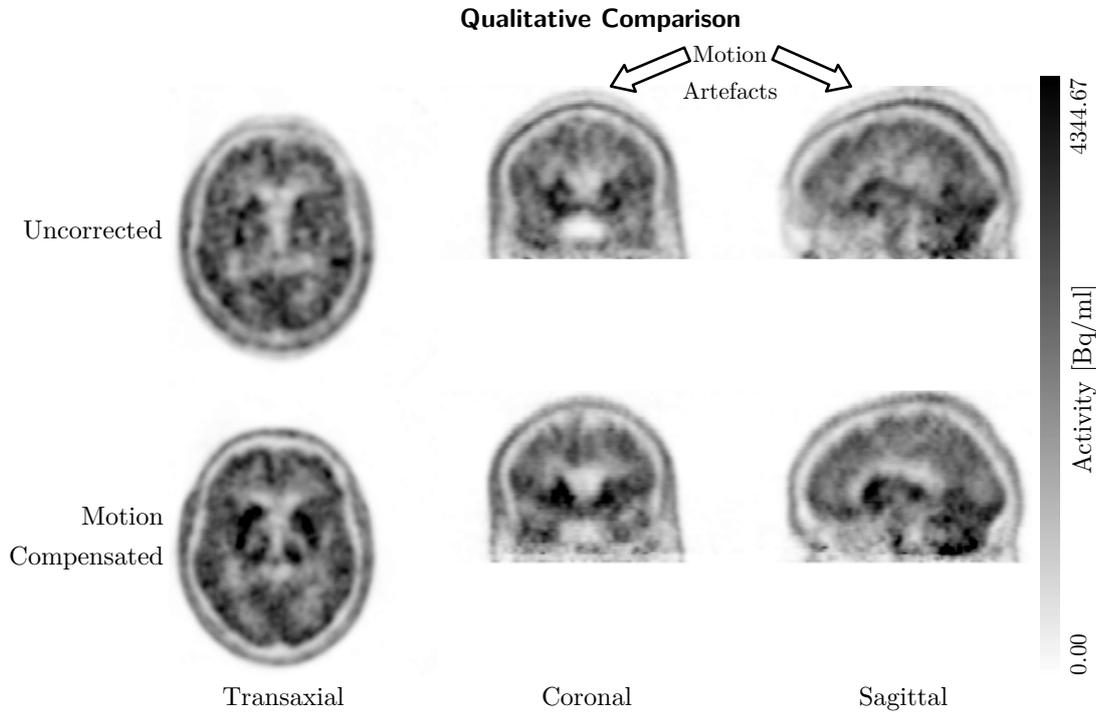


Figure 3.22:

Qualitative Comparison Results: For a qualitative evaluation of the accuracy of the motion compensation, the PET images of an uncorrected and corrected data set of a clinical [^{18}F]DOPA brain acquisition (3D list-mode, 171 MBq, 55 min, 27 frames) have been compared. The top row shows the image data of the uncorrected data set with visible motion artefacts at the skull, as a result of the patient's motion. In contrast, the bottom row shows the same cross section but after the event-based motion correction has been applied. Apart from the reduction of artefacts, the image contrast is also improved.

In the top row of the figure a single cross section of the uncorrected data set is shown. In addition to a general motion blurring it also shows severe image artefacts at the area of the patient's skull visible in the Coronal and Sagittal slices, as a result of the patient's motion. The same cross section is shown in the bottom row of the figure, but after having applied the event-driven motion compensation method as discussed in section 2.4.

In addition to this figure, which demonstrates the improvement in the image quality due to the motion compensation, a quantitative analysis has been performed. For this analysis nine elliptical ROIs have been placed within the volume data, according to the aforementioned ROI delineation protocol (cf. section 2.6). In a first test the activities in these ROIs have been analysed by calculating a *time-activity-curve* (TAC) per ROI. Figure 3.23 shows the TACs for the two ROIs delineating the *Nucleus caudatus* (ncr and ncl) in the selected study. In addition, the TACs for the reference ROI in both, the uncorrected and corrected data set are shown. Referring to the results shown in these TACs, the uncorrected data set shows a progression of the tracer accumulation over time with a maximum difference of $\approx 30\%$ found in the TAC values of the right *Nucleus caudatus* (ncr).

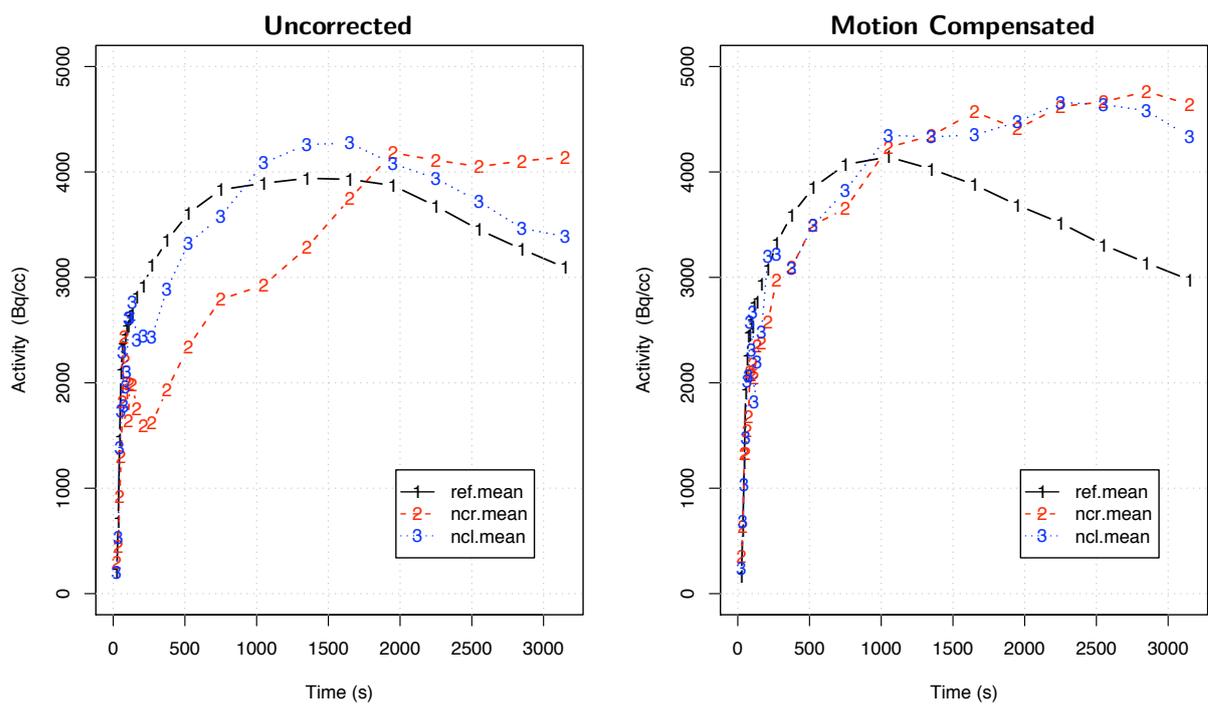


Figure 3.23:

Time-Activity-Curve Evaluation: To analyse the tracer accumulation over time TACs have been generated from the delineated ROIs. Here the TACs of the reference ROI as well as on the ROIs delineating the *Nucleus caudatus* are shown. The plot on the left-hand side shows a large difference in the tracer accumulation of the left and right *Nucleus caudatus* (ncr/ncl) and has been generated from the uncorrected data set. In contrast, the plot on the right-hand side shows the TACs of the motion compensated data set. Due to the motion compensation the right TACs show a better concordance to the expected progression of a tracer accumulation of [^{18}F]DOPA.

To support the difference in the tracer kinetics of the uncorrected and corrected data set, a *Patlak* analysis has been performed. As discussed in section 2.6, the R_0k_3 parameter has been determined for each ROI of the data sets. Table 3.8 and figure 3.24 summarise the values, including the information on the respective ROI volumes as well as on the norm values according to the aforementioned quantitative tracer kinetics protocol.

Table 3.8:

Tracer Quantification Results: The value of the R_0k_3 parameter for the nine delineating ROIs along the striatum has been determined via a *Patlak* analysis. Here a list of the calculated R_0k_3 parameters for the uncorrected and corrected data set is shown. For a comparison the norm values for a healthy patient collective are given [113].

ROI	volume [ml]	Uncorrected	Compensated	norm
		$R_0k_3^{mean}$ [1/min]	$R_0k_3^{mean}$ [1/min]	$R_0k_3^{mean} \pm 1\text{ sd.}$
ncr	0.524	0.015 59	0.012 74	0.0166 \pm 0.0010
ncl	0.524	0.004 60	0.010 54	0.0161 \pm 0.0022
pr	1.573	0.001 89	0.009 47	0.0160 \pm 0.0014
pl	1.573	0.004 06	0.009 51	0.0167 \pm 0.0011
pra	0.524	0.004 25	0.009 38	
prm	0.524	0.003 25	0.010 99	
prp	0.524	0.002 24	0.008 75	
pla	0.524	0.009 48	0.008 85	
plm	0.524	-0.000 79	0.007 99	
plp	0.524	0.001 33	0.010 24	

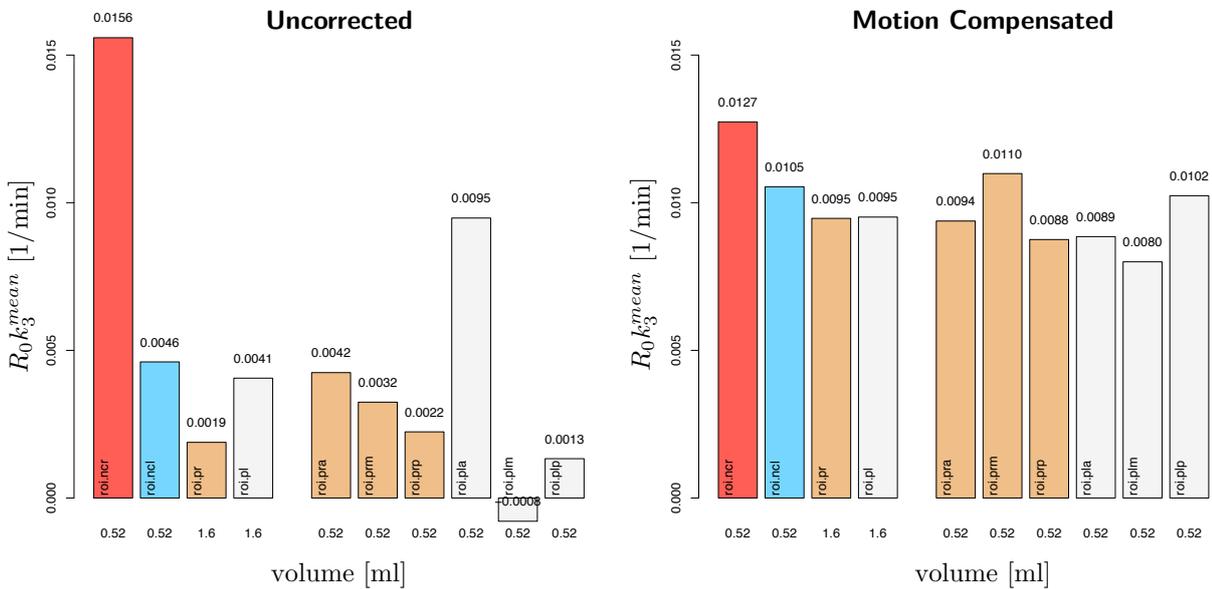


Figure 3.24:

Tracer Kinetics Results Plot: The R_0k_3 values listed in table 3.8 for the motion corrected and non motion corrected data set.

Parametric images have been generated to analyse the distribution of the R_0k_3 parameter – not only within the delineated ROIs. These three-dimensional images are generated from the original PET images and contain the value of the R_0k_3 parameter for every image voxel. Figure 3.25

shows the output of the parametric images of the uncorrected and motion compensated data set on a single cross section.

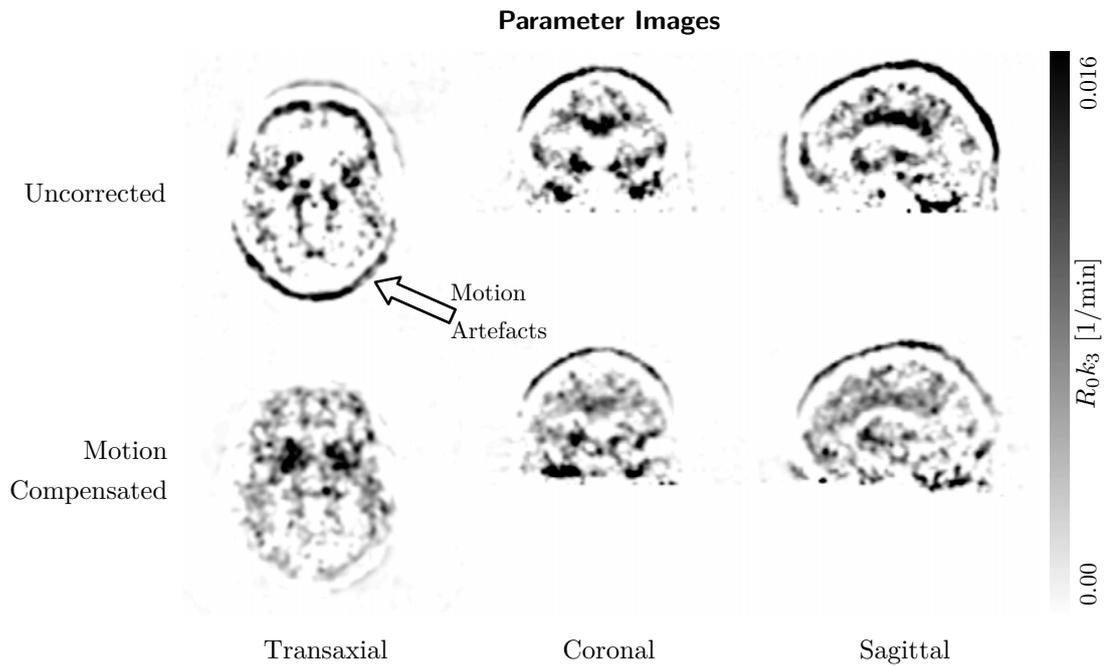


Figure 3.25:

Parametric Image Evaluation: Parametric images of the R_0k_3 values have been generated as part of the quantitative tracer kinetics analysis. The images of the uncorrected and corrected data set illustrate the spatial distribution of the R_0k_3 parameter. Apart from a typical R_0k_3 distribution in the striatal areas of a brain, the high values along the skull illustrate typical artefacts due to motion. In contrast, the motion compensated parametric image shows less motion artefacts and thus points out a successful motion compensation.

4 Discussion

A scientist's aim in a discussion with his colleagues is not to persuade, but to clarify.

(Leo Szilard)

4.1 Optimised List-Mode Data Acquisition

Our methods for an optimised list-mode acquisition with an ECAT Exact HR⁺ scanner suggest that important limitations of this PET scanner type have been resolved. In the following sections the implications and the results of our list-mode acquisition methods will be discussed.

4.1.1 Shared Storage Device Acquisition

The results of our development of a method using a shared RAID system for the access to the acquisition data of an ECAT Exact HR⁺ PET scanner show a considerable increase in data access speed. This hardware-based list-mode acquisition overcomes long-term limitations of the *acquisition control system* (ACS2) hardware of the HR⁺ system. The achieved increase in speed, for example, of $140 \times$ (in comparison to the ACS2's own capabilities) results in a list-mode data transfer requiring only ≈ 1 minute in contrast to the previous ≈ 2.5 hours for a typical amount of ≈ 4 GB of acquired data (one hour 3D acquisition with [¹⁸F]FDG @ ≈ 330 MBq). Consequently, the previously unacceptable duration for a clinical access to list-mode data has been successfully reduced to an amount that allows for the clinical application of list-mode with this scanner type.

Furthermore, the possibility to immediately transfer all acquired data off the ACS2's storage directly after the acquisition improves the general stability of the ACS2 system and therefore of all kind of PET acquisitions with this scanner. A parallel access, for example, to data stored on the ACS2's storage device is usually not possible without compromising the overall stability of the ACS2 system. This issue is resolved by the immediate transfer and the transparent data access through our *shared storage device* solution. Additionally, other standard tools (e.g. visualisation tools such as the *ECAT Volume Viewer*) accessing newly acquired data also profit from these new access methods. The acquired data become available on the ECAT workstation directly after an acquisition is finished. This enables these tools to access acquisition data more efficiently. Even for large amounts of acquisition data (e.g. sinograms with a large number of frames) this results in an instant access to the requested data after an acquisition, and thus is also interesting for conventional *histogram-mode* acquisitions.

All this has become available due to the implementation of the VXEXT file system for the Linux operating system. Extensive tests of the new file system implementation showed that a fast and reliable access to the data is possible without interfering with any usual operations of the standard PET gantry hardware or software. As a consequence of that development and to

support the ongoing development of such a custom file system implementation, the source code has been published within the Linux developer community under an open-source license [27, 28].

The only shortcoming of the shared storage solution is that the data access of the ACS2 and the Linux system have to be synchronised. This, however, does not represent a real problem because the ECAT system can be queried remotely if an acquisition is finished. Therefore, as soon as an acquisition is finished the Linux system will be instructed to transfer the acquisition data from the storage device. Afterwards, the ECAT system is instructed to delete the data via a remote procedure call. The space on the storage device is then immediately freed.

Therefore, our *shared storage device acquisition* method, in combination with the file system driver, provides new ways for a transparent access to coincidence data sampled with an ECAT Exact HR⁺ PET scanner. This fast data access method is a unique possibility to arrange for list-mode acquisition within a clinical environment, especially for list-mode acquisitions with a generally large amount of data. Moreover, even for a standard HR⁺ system which is used in histogram-mode only, our shared storage solution improves important aspects of the PET scanner environment. In fact, it improves the stability of the ACS2 and generally speeds up the access to any acquisition data stored on the ACS2.

4.1.2 DAQ-based Data Acquisition

As an extension of the aforementioned *shared storage device acquisition*, our *DAQ-based data acquisition* method allows for accessing list-mode data without having to wait until the PET examination is finished. Moreover, this hardware-based method allows to acquire and process list-mode data in real-time. The DAQ method allows for directly acquiring and immediately processing list-mode data on an external workstation while the actual PET acquisition is still in progress. This represents a new way of data acquisition for ACS2-based PET systems. By using the DAQ solution, an ECAT Exact HR⁺ can be operated in conventional histogram-mode, while a concurrently running Linux system acquires and processes the raw list-mode data from the external data bus of the ACS2. In addition to an immediate access to the coincidence data, this has another positive effect on the stability of the ACS2 hardware in general. Previously, the ACS2 had to be reconfigured and rebooted explicitly for list-mode processing. Now that a transparent hardware-based sampling of the coincidence data is possible, the HR⁺ can be operated constantly in histogram-mode while the raw list-mode data is transparently sampled via the DAQ hardware. This allows for the use of the conventional data processing environment within the PET system instead of having to use custom shell scripts to prepare the ACS2 for list-mode acquisition.

Although our DAQ solution requires a direct interaction with the ACS2 hardware, an interference with the PET scanner's own operations is avoided. In fact, our hardware adapter works fully electrically isolated (*galvanic decoupling*) because of its optoelectronic components. This facilitates the passive sampling and forwarding of the raw coincidence data without any risk for the standard hardware of the PET scanner. This is particularly helpful in clinical environments where regulations or maintenance agreements deny a modification of the standard PET acquisition protocols.

In addition, by being able to acquire a sinogram and list-mode file in parallel, a physician may base his first clinical evaluation on the sinogram data which is generated in conventional histogram-mode. Then, in a second step the physician may take the list-mode data into consideration, after having applied certain post-processing methods such as an event-driven motion compensation or data reframing. This also potentially increases the examination flexibility of a PET facility because it allows for providing both, list-mode and histogram-mode acquisitions, in parallel and without having to switch between the acquisition protocols. Such a DAQ-based list-mode acquisition might serve as a new approach for facilitating potential list-mode-based image reconstruction methods [36, 70, 71]. Due to the immediate sampling of the coincidence data, for example, a real-time image reconstruction becomes a tangible goal.

As for a potential routine use of the DAQ method, the results of the test measurements suggest that a real-time acquisition of list-mode data is possible without interfering with the standard operations of the PET scanner. Only a statistically negligible amount of coincidences are lost during an acquisition due to the limitations of the DAQ cards that are used in our method. One minor obstacle, however, is the fact that during a transparent DAQ-based acquisition which is sampling the coincidence data from the PET scanner in histogram-mode, the ACS2 hardware does not emit any time events. Therefore, we have developed a method to emit artificial time words at the very time the list-mode data is registered by the DAQ software. In comparative tests we have further verified the correct operation of that method and could show that the artificially inserted time words are in line with the time words the ACS2 usually inserts during a normal list-mode acquisition.

4.1.3 Event-based Attenuation Measurement

For an accurate quantitative analysis of PET data, the measurement of the attenuation is essential. However, like for an emission, the acquisition of the attenuation is also susceptible to motion artefacts. While a wide variety of motion compensation approaches exist and are documented to correct patient motion during the emission phase of a PET acquisition, a time resolved motion compensation during the transmission phase has not been documented yet [13, 14, 16, 35–37, 39, 79, 82, 92, 120–123]. This is mainly due to limitations of the hardware and software of PET scanners which in fact are also true for the ECAT Exact HR⁺ scanner model.

Our analyses and development of methods to achieve a list-mode acquisition of transmission data with an HR⁺ show, that this is possible by minor adaptations of the protocols and the corresponding list-mode sorter software. We have been able to acquire transmission data in list-mode by extracting the radioactive sources during a list-mode acquisition on a standard HR⁺. As a result of the evaluation of this method, differences in the recorded list-mode data format of a transmission scan have been identified. To incorporate these differences into our methods, an existing list-mode sorter software (`lmSorter`) has been modified to be able to generate ECAT compliant transmission sinograms out of the slightly different list-mode data. To evaluate the correctness of this method, test measurements on the general stability of the new method and a quantitative voxel-wise comparison have been carried out. This comparison shows, that the transmission data recorded in conventional histogram-mode and list-mode do

not show any differences greater than 3.5% in the coincidence data and is therefore applicable for clinical use.

However, the new acquisition method has one limitation: a list-mode acquisition of transmission data is only possible in 2D mode with the septa rings extracted. This makes it impossible to apply an event-driven motion compensation (cf. section 1.2.3). Nevertheless it still allows, for the application of certain other types of motion compensations. A list-mode acquisition of transmission data can, for example, be used to split the coincidence data based on the magnitude of the patient motion. Motion compensation methods like, e.g. *multiple acquisition framing* (MAF) or *gated acquisition* can then be applied to 2D sinogram data [13, 82].

In addition, the shown possibility to acquire list-mode-based transmission data may help in developing more advanced image reconstruction and motion compensation methods. In that case, data on each transmission coincidence could be directly incorporated in the algorithms instead of having to extract them from compressed transmission sinograms as it is today.

4.2 Motion Quantification

As outlined in section 1.2.2, there are a number of different methods to quantify the motion of a patient. A comparison of the analysed quantification methods is summarised in table 4.1. It shows the advantages and disadvantages of the individual quantification methods and is used as the basis for the discussions in the following sections.

Table 4.1:

Motion Quantification Method Comparison: While external systems like optical tracking devices generally tend to output more accurate and reliable motion data, methods like an image-based *coregistration* are easily available and inexpensive.

Method	Δ Advantages	∇ Disadvantages	References
Radioactive Markers	+ easily available	- radiation exposure - data interaction	[79]
Video-based Methods	+ inexpensive	- low precision - no direct motion data	[17, 80]
Image-based Methods (<i>e.g. coregistration, SPM</i>)	+ inexpensive + fast	- inter-frame only - late frames only	[81, 82]
Analogue Motion Sensors (<i>e.g. respiration belt</i>)	+ inexpensive	- low precision - relative motion only	[12]
Optical Tracking Device <i>active tracking</i>	+ high resolution	- expensive - marker synchronisation	[86]
Optical Tracking Device <i>passive tracking</i>	+ high resolution + high sampling rate	- expensive	[14, 15], [85, 124]

4.2.1 Cross-Calibration

A calibration between the PET scanner's own coordinate system and the coordinate system of the motion quantification device is necessary for the motion compensation of PET data.

Our *cross-calibration* method (cf. section 2.3.1) suggests the calculation of a spatial transformation matrix which is used during analyses on motion data and during motion compensation.

When applied to motion data, this matrix transforms the data of the motion tracking system into the PET scanner's coordinate system via a single spatial transformation. This coordinate transformation makes it possible to apply motion analyses and motion compensation methods which usually require to have the data available in the same coordinate system as the one where the PET scanner is measuring the patient.

For a convenient calculation of this *cross-calibration* transformation matrix, a software package has been developed. As shown in section 3.2.1, this software package allows for a semi-automatic calculation of the matrix. By providing graphical user interface elements, the motion tracking data as well as image data from a cross-calibration PET scan are loaded interactively. In addition, our software implements the aforementioned algorithms to calculate the cross-calibration matrix by matching the identified positions of a minimum set of four markers in both coordinate systems. Finally, by having matched the coordinates of the markers properly in each coordinate system the cross-calibration matrix is calculated and saved into a custom matrix file format. This software simplifies the complex processing of a cross-calibration, especially in terms of a clinical application of motion tracking.

In test acquisitions, the principle and the accurate operation of the cross-calibration method and its involved algorithms have been evaluated. The results show, that due to the calculated cross-calibration matrix, the motion data of an arbitrary moved point source has been transformed successfully from one coordinate system to another. In a comparative test with an uncalibrated motion data set of the same point source, the motion compensated image demonstrates that due to the cross-calibration matrix the point source which previously vanished becomes visible again. This is an indicator that the cross-calibration method is able to match the two coordinate system sufficiently. However, as this type of evaluation requires a working event-driven motion compensation method, which is to be discussed in a further section of this thesis, a direct statement on the validity of the cross-calibration cannot be given. Still, a plausibility of the cross-calibration can be seen from the comparison of the calibrated and uncalibrated test images (cf. figure 3.4).

As a cross-calibration implies an attenuation measurement with the PET scanner of several minutes in duration, it cannot be performed right before each clinical motion tracking acquisition. However, as the tracking system is usually located at a fixed position within both coordinate systems, a re-calibration is only required in case one of the two systems is relocated. Moreover, as such manual relocation is only rarely the case, a re-calibration is usually performed only in regular maintenance intervals (e.g. once a month). This ensures a general measurement consistency as a motion tracking system is usually subject to stability or drifting irregularities over time.

In our PET facility we have established such a regeneration of the cross-calibration with the regular PET scanner maintenance. Over the years of usage of the motion tracking system, a monthly re-calibration proved to be sufficient to assure a good motion capturing. During such a regular re-calibration of the coordinate systems the motion tracking system itself is also re-calibrated. Furthermore, the aforementioned evaluation test with a constantly moving point source is repeated with every re-calibration and thus part of the regular tracking system maintenance. As a final test of the motion tracking maintenance, this point source evaluation

provides an intuitive quality control which allows to value the quality of the current tracking system set up.

4.2.2 Motion Target

The type of the used motion target usually depends on the actual motion quantification method. In accordance to the comparison of the motion quantification methods shown in section 1.2.2 and table 4.1, three head motion targets for the ARTtrack optical tracking device have been developed.

The evaluation of the motion targets (cf. section 3.2.2) show, that the *safety goggles*-based head motion target has the best motion transferability. This is the result of the motion registered with each target compared to the motion registered with a reference target fixed at the *maxilla* of a volunteer. The results shown in the motion target distance comparison plot (cf. figure 3.8) illustrate clearly, that the safety goggles target is the most stable of the three candidates. This is an interesting result, because the *Os frontale*-based head motion target was assumed to be the most reliable and accurate head target within our clinical environment. In a direct cooperation with the PET group of the *Forschungszentrum Jülich/Germany*, this target has been developed with the goal to replace the ski goggles-based head motion target in our own clinical PET environment. A similar *Os frontale*-based target has in fact already been evaluated by the group in Jülich as part of a Master's thesis [107]. The results of this thesis suggest, that an *Os frontale* attached motion target produces the most reliable results. In a similar quantitative evaluation, the group in Jülich found the best correlation between a maxilla-fixed reference target and an *Os frontale* attached target. To verify the results of this thesis, a similar head target has been developed and was given to our facility. However, as shown in the results of section 3.2.2, we were not able to support this observation by our own tests.

The discrepancy in the results of the evaluation of an *Os frontale*-based head motion target may be due to the different tracking systems at the PET group in Jülich and our group. While the group in Jülich used a *Polaris* motion tracking system [88] with a total number of two optical cameras installed at a fixed distance to each other, an *ARTtrack* motion tracking system [87] with a total number of four freely positionable cameras has been used during our evaluations. In addition to differences in the technical specifications of these tracking systems, the evaluation performed at Jülich has been focused on finding a replacement target for experimental examinations mainly. In contrast, the evaluation at our facility focused on the actual usability of the targets in a clinical environment. This difference has the following implications. In our clinical environment the patient compliance is a very important aspect in addition to the required motion transferability of the motion target itself. Furthermore, the target has to be easily attachable on the patient's head as well as allowing for a fast and uncomplicated removal.

Therefore, in our evaluation of an *Os frontale*-based head target, the attachment of the target via adhesive tapes proved to be critical for clinical use. As figure 2.17 and the results in table 3.5 show, it does not only take a long time to attach such a target onto the patient head, but the removal also proved to be difficult. Here the goggles-based targets showed a clear advantage. If a patient has hairs which cover the area of the *Os frontale*, an attachment via adhesive tapes is particularly complicated or even impossible. In contrast, a goggles-based target proved to

be attachable more easily on most kinds of human heads and hairstyles. Furthermore, during long acquisitions the attachment via adhesive tapes is critical because of the fading grip of these tapes, especially when in contact with skin. Here, the attachment via elastic straps around the patient head showed more reliability during long acquisitions. In addition, the subjective rating of the patient compliance of each target shows that the safety goggles target is preferred. However, the two types of goggles-based head targets also showed a difference not only during the quantitative motion transferability comparison. During the evaluations, the safety goggles head target was experienced more comfortable than the ski goggles target. This was also verified in measurements with regular patients. A possible reason for this subjective observation may be the plastics frame, which is softer for the safety goggles. In addition, the replaceable foamed material at the inner lining of the safety goggles add to extra comfort and fit for the patient.

Therefore, not only the quantitative analysis of the motion transferability suggests the safety goggles-based head motion target to be the most appropriate target for clinical use. Also the subjective evaluation of the targets according to their clinical fit on patients shows that the safety goggles target is best suited in a clinical environment. Furthermore, after integration of the safety goggles-based head motion target in clinical practice, the results during routine operation support this observation. In fact, no motion measurements had to be aborted so far because of patients feeling uncomfortable with the motion target. This had previously happened a few times with ski goggles-based motion target.

4.2.3 Motion Analysis

It is problematic to draw an unambiguous conclusion concerning the magnitude of the patient motion from raw motion tracking data only. This has been pointed out in previous sections of this thesis (cf. section 2.3.3). However, it has also been shown that an assessment of the magnitude of patient motion can be achieved by a quantitative statistical analysis of the motion data (cf. section 3.2.3).

Our methods are using a statistics toolkit ('R') for performing analyses on the raw motion data. This toolkit is used to calculate parameters for the evaluation of the motion magnitude. This is done by projection of the motion data on the surface of an imaginary patient head (a sphere with a diameter approximating a 'mean' patient head diameter). The spatial transformation and the evaluation of selected points on the surface of the virtual sphere provide parameters which allow for direct assessment of the motion magnitude. Instead of having to rely on the raw motion parameters (i.e. rotation and translation parameters) which have a combinatorial relation, our methods provide a better classification of the motion magnitude for a specific ROI (e.g. the patient head).

One main result of our motion analysis method is the calculation of statistical values for the detailed analysis on the patient head motion. As discussed in section 2.3.3 and shown in figure 3.9 to 3.12, our head motion analysis is based on the constant calculation of the distances of points on a virtual sphere after having spatially transformed them with each recorded motion data sample of the tracking system. As a result of this method, a motion sample is classified as 'significant' by applying a single threshold to all calculated distances on the virtual sphere. As soon as one of the distances on the sphere exceeds this threshold, the respective motion

sample is considered to be ‘significant’ and used during motion compensation. In addition, the distribution of the calculated distances of the ‘worst’ and ‘best’ point on the sphere are illustrated in frequency plots. Apart from information on the frequency of a distance value, the inspection of these frequency plots allows to classify the nature of the motion. If a frequency plot shows, for example, a generally high frequency in all channels of the plot, this is a degree for high image blurring (i.e. *continuous drifting motion*). In contrast, if a frequency plot shows only high frequencies of distances in certain distinct areas, the motion can be classified as *one-time motion*. This means, the patient moved only once or in distinct steps during the acquisition. As a consequence of this observation different motion compensation methods may be applied.

Furthermore, the two-dimensional image maps (cf. figure 3.11) of the sphere point distances provide information on the actual location of motion intensive areas. In one image map (part c) the standard deviation of the distances is illustrated for each grid point. This corresponds to the fluctuation of the distances of a sphere point which is in fact another sign for image blurring. In a second image map (part d) the mean distance for each sphere point is illustrated. In addition, it allows to identify the points on the sphere which have been moved away from their initial position farther than others during acquisition.

The image maps and the range plot on the distance distribution of the sphere points are a tool for a detailed assessment of the motion data. In addition, the implementation of an interactive output of the motion analysis data via a three-dimensional sphere plot provide an interesting new approach for the assessment of head motion. The sphere plots visualise the areas of significant motion directly. Furthermore, as these plots contain labels on the actual anatomic areas, the area on the patient head on which the motion had most effect can be identified. Therefore, even the reason for the occurred motion can be identified. If the motion distribution on the imaginary patient head indicates, for example, that the patient has moved most in the area of the *Os temporale* (left/right), this could be a sign of an insufficient fixation of the patient.

For clinical use, the statistical parameters as well as the aforementioned motion plots have been included in a report document. This document is automatically generated in the ‘pdf’ file format and thus can be integrated in every patient examination as a means of quality control document. The final *estimated motion score* shown on the motion summary page of this document essentially indicates the actual magnitude of motion the patient has performed during the acquisition. Based on this score the other pages provide an overview with detailed parameters for various aspects. The medical staff can identify, for example, the motion intensive areas on the patient by looking at the image maps. Similarly the time-relative plots on the patient motion illustrate the times where the patient moved most. This is very useful for a clinical evaluation of the magnitude of patient motion and its potential impact on a PET examination.

4.3 Event-Driven Motion Compensation

An event-driven motion compensation requires more than just methods for a spatial reorientation of LORs. Post-correction methods are required to compensate aspects which are based on specific details in the geometry and data structure of PET systems in general. This was already shown in the previous chapters. These post-correction methods are a very important part of an

event-driven motion compensation. Therefore their proper implementation and optimisation is essential and thus is discussed in the following sections.

4.3.1 Normalisation Correction

The evaluation of our normalisation correction method (cf. section 3.3.1) shows, that the extraction and calculation of the efficiencies of each γ -detector are correct. This is supported by the results illustrated in figure 3.13. The voxel-based intensity correlation plot shows a clear concordance between two list-mode data sets which have been normalised before and during the image reconstruction. The minimal differences in the voxel intensities of these data sets of $< 0.1\%$ can be neglected and thus support the correctness of our normalisation correction.

In our event-driven motion compensation the detector sensitivities are extracted from the standard ECAT normalisation files. Then normalisation correction factors $f_{(i,j,i',j')}$ are calculated for every LOR (cf. section 2.4.1). As a result the motion compensation generates sinograms which allow the application of the standard normalisation coefficients when the image data are reconstructed. This facilitates later use of the motion compensated sinogram data as it can be handled in the same way as non-motion compensated data – which usually requires a normalisation file during image reconstruction.

4.3.2 LOR Discretisation Correction

The results of the evaluation of our optimised LOR discretisation correction show an improved assignment of counts because of the LOR overlap weighting. As illustrated in figure 3.14, the count distribution in part (c) is more in line with the applied movement of $1.5\times$ of the voxel width. In fact, compared to the results of part (b), the centre of the count distribution is at pixel 2.59, rather than at pixel 3. This indicates a more accurate assignment of the counts because the counts in the unmoved data set (part a) are located at pixel 4. Therefore, counts are distributed more accurately to the expectable bins in case of our optimised LOR discretisation. For movements which result in a reorientation of LORs along the Z-axis, the optimised LOR discretisation correction particularly provides a more accurate assignment of counts to the individual sinogram bins. This is because the original discretisation correction [15] did not implement an overlap weighting along the Z-axis due to speed considerations. Our new method, however, uses a grid for the intersection overlap calculation and thus the number of evaluations on the 10×10 grid is always constant. Therefore only a minimal amount of time is required per LOR reorientation. This is substantiated by the overall speed of the event-driven motion compensation, as shown in figure 3.19.

While the evaluation of the optimised LOR discretisation correction has been performed in one direction only (y_ϕ), the results are transferable to all other directions as well. The evaluation of our Out-of-FOV correction method have in fact been performed with the optimised discretisation correction and thus support the aforementioned results. In addition, the clinical evaluation of the motion compensation (cf. section 4.5) also shows, that our optimised correction is capable of producing accurate results even for data with motion resulting in arbitrary overlaps of LORs with γ -detector surfaces.

4.3.3 Out-of-FOV Correction

As shown in the previous chapters, an OFOV correction is an essential part of an event-driven motion compensation. If the lost counts caused by the Out-of-FOV events are not compensated, the PET images will turn out with inhomogeneity. Apart from major artefacts in the images this also affects the possibility for an accurate tracer quantification of motion compensated PET data. However, a complete compensation of the lost counts is not possible due to the nature of the OFOV problem. In fact, the OFOV correction method, as discussed in [15], is based on the scaling of counts relative to the amount of time that an LOR was found to be outside the FOV. In practice, however, this scaling of counts relative to one fixed point in time can lead to additional image artefacts, as shown in section 2.4.3.

Our OFOV correction method uses an optimised approach. It also scales the registered counts relative to the OFOV times of each LOR, but is largely based on the avoidance of these events as well as on the reduction of the calculated scaling factors. Within a two-fold process, the amount of OFOV events and thus the number of large OFOV scaling factors is reduced by motion compensation to frame-wise reference positions rather than to a single acquisition-wide position. This accounts for the fact that during long acquisitions patients are usually located farther away than immediately after the acquisition start.

As shown in the results section 3.3.3 the new optimised OFOV correction method calculates reduced OFOV scaling factors and thus reduces the amount of image artefacts. This image improvement is not just the result of a pure reduction of the OFOV scaling factors, but also due to the lower magnitude of the motion compensation. Therefore, this scaling factor optimisation reduces the probability of potential OFOV-based image artefacts, but it does not eliminate it completely. Thus, there is still possibility of OFOV factors resulting in an overestimation of counts if the patient moves near the borders of an acquisition frame. This observation is largely based on the nature of the scaling factor calculation ratio, as shown in equation 2.6. In addition, our OFOV correction method is still not capable of compensating for lost counts in situations where the subject has moved immediately after acquisition start and thus no counts have been registered for a specific LOR at all ($t_{OFOV(i,j)} \cong t$). In fact, this indicates a limitation of our OFOV correction and thus leaves room for future improvement.

However, our optimised OFOV correction method generates clearly improved PET images for a quantification of the tracer kinetics in most of the OFOV situations. This is especially true for neurological examinations where not the whole FOV is occupied and the main ROI is usually near the centre of the FOV (e.g. the striatum). Within this ROI the lost OFOV events therefore only have a minimal impact on the PET data as OFOV artefacts mostly occur on the fringe of the FOV. This is also supported by the results of our statistical analysis on the distribution of OFOV factors, as shown in table 3.6. The quantitative values, but also the qualitative observation of the motion compensations performed with the optimised OFOV method show a clear scaling factor minimisation over the non-optimised OFOV correction method.

4.4 Clinical Integration

For the integration of our methods in a clinical environment different software tools have been developed. Those tools have been presented in section 2.5 with screenshots and a speed comparison shown in section 3.4.

The tools especially with a graphical user interface (`lmacq` and `trkacq`) represent an important improvement for the successful clinical integration of an event-driven motion compensation. As shown in figure 2.30 and 3.18, these tools are directly invoked from the standard PET scanner user interface. Together with the automatic processing of the list-mode data by either `lmSorter` or `lmMover` the clinical use of an event-driven motion compensation is facilitated considerably. The usually complex process right from acquiring the data until the final post-processing of the motion compensated image is dramatically simplified due to the automatic processing of the necessary tools. In fact, previously all the different steps had to be executed manually. Now the whole process of acquiring list-mode and applying a motion compensation is performed fully transparently for the medical staff.

Furthermore, because of the multiprocessor optimisation of our list-mode processing software tools (`lmSorter` and `lmMover`), the additional time and effort for using list-mode as a standard acquisition mode is not an obstacle anymore. The optimised data access, the automatic processing and the speed optimisations shown in figure 3.1 and 3.19 reduced the time needed for handling and processing of list-mode data to an acceptable period of 15 to 20 additional minutes. This compares to times when it took more than 24 hours, including the time needed to access the list-mode data. As the uncorrected histogrammed data set is already available at the time the motion compensation is processed, the reduced amount of time is therefore easily sustainable in clinical use. Moreover, due to the supply of statistical data on the patient motion (i.e. `analyze_trk`) right after an acquisition, a physician is able to evaluate the influence of the patient motion on the PET image data immediately without having to wait until the motion compensated data set is available. This facilitates the examination flexibility considerably.

4.5 Clinical Evaluation

The evaluation of the accuracy and correctness of the motion compensation (cf. section 3.5) shows, that our event-driven compensation methods are able to improve the image quality. Moreover, in two evaluations we have analysed the qualitative, but also the quantitative improvements based on experimental and clinical data.

I. Point Source Evaluation

The first evaluation has been performed as a *proof-of-concept* test of our event-driven motion compensation in general. As shown in figure 3.21, images of different phases of an acquisition with a radioactive point source in the FOV have been generated. The motion compensation is able to restore the quality of the image data. Compared to the image which has been reconstructed from the list-mode data without motion (reference image), the motion compensated image clearly restores the spherical structure of the imaged point source. This is also supported

by comparison of the motion compensated image to the uncorrected image with the typical and visible blurring.

The reconstructed image resolution in the motion compensated image and the reference image has been calculated to further support the image improvement with a quantitative analysis. The results in table 3.7 show, that the motion compensated image of the point source has a calculated image resolution of approximately 6.6 mm. Compared to the image resolution of ≈ 5.6 mm in the image of the acquisition phase at rest, this is a decrease in resolution of ≈ 1 mm. Considering that the effective resolution of the HR⁺ PET scanner is specified to be between 5 mm to 6 mm [125], the reconstructed image resolution in the motion compensated image versus the point source image at rest still shows an overall good performance of our motion compensation.

II. Patient Data Evaluation

In the second phase of the clinical evaluation an event-driven motion compensation has been performed on the list-mode data of a selected clinical [¹⁸F]DOPA brain acquisition. As shown in section 2.6 and 3.5, patient data have been selected by an analysis of the motion data. The accuracy of the motion compensation has then been assessed by qualitative and quantitative tests.

As shown in figure 3.22, the qualitative results demonstrate, that our optimised event-driven compensation is able to ‘motion compensate’ the clinical patient data accurately. The comparison of the corrected images versus the uncorrected images shows a clear reduction of visible motion artefacts, especially at the top of the Coronal and Sagittal slices. Furthermore, the motion compensated data illustrate, that the typical image blurring has been reduced and the overall image contrast has improved. Therefore, important anatomic brain structures (e.g. the striatum) are more clearly visible after the motion compensation. This is an important improvement for the evaluation of such motion affected PET data. In fact, the given patient data would have been of no clinical use without a motion compensation. The magnitude and nature of the patient motion, especially within the regions-of-interest (e.g. *Nucleus caudatus*), make it almost impossible to evaluate the patient data according to the usual clinical questions of such a [¹⁸F]DOPA examination. Thus it would have been necessary to repeat the PET acquisition, which however is often impossible in clinical routine.

To emphasise the qualitative improvements due to the motion compensation, a quantitative evaluation has been carried out. Within this quantitative evaluation *time-activity-curves* (TAC) have been calculated for delineated ROIs of the striatum (cf. figure 3.23). A comparison of the TACs have shown, that due to the motion compensation the radioactivity in the ROIs of the *Nucleus caudatus* has a better concordance with the typical progression of the tracer dynamics of [¹⁸F]DOPA. The activity values for the right *Nucleus caudatus* in particular show a large difference ($\approx 30\%$) between the uncorrected and corrected PET data set. In fact, while the TACs of the uncorrected data set show a large asymmetrical tracer accumulation between the right and left *Nucleus caudatus*, the motion compensated data set shows a clear reduction of this asymmetry. Furthermore, the TACs of the motion compensated data set show a higher tracer uptake at early times of the acquisition. In addition, the longer tracer accumulation in the *Nucleus caudatus* over time is in better concordance with a usual TAC found in a data set

without any motion. Furthermore, the more typical TAC progression of the motion compensated reference ROI is another indication that the motion compensated data set should be preferred in clinical use.

In addition to the TAC analysis, table 3.8 and figure 3.24 illustrate the results of a performed *Patlak* analysis [118, 119]. Similar to the TACs, the R_0k_3 values of the left and right *Nucleus caudatus* also show an asymmetry in the tracer kinetics of the uncorrected data set. The respective values of the motion compensated data set, however, show a clear reduction of this asymmetry. To further emphasise the high uncertainty in the uncorrected values, the R_0k_3 value of the left *Putamen medius* (plm) shows an obvious incorrect negative value. Under the assumption that this ROI has been positioned accurately enough, such negative values generally suggest the existence of motion artefacts. The reduction of the large asymmetry and the restored value of the plm in particular substantiate that there are less motion related artefacts in the motion compensated data. Furthermore, the R_0k_3 value of the right *Nucleus caudatus* seems to suggest an almost ‘normal’ dopamin metabolism in the uncorrected data set. However, the large difference between the ncr and ncl values, the negative plm value as well as the anamnesis of the patient show that the motion compensated data set produces more accurate quantitative results.

This is also supported by the parametric image comparison results shown in figure 3.25. In these images the R_0k_3 parameter has been calculated for every voxel of the original PET image. Again, the uncorrected data set shows typical artefacts as a result of the patient motion. In the corrected data set, however, there is an almost typical distribution of the R_0k_3 values within brain structures, whereas the parametric image of the uncorrected data clearly exhibits artefacts because of high R_0k_3 values at the skull of the patient. In addition, the R_0k_3 values in the area of the left *Nucleus caudatus* also support the incorrect asymmetric difference between ncl and ncr (cf. table 3.8 and figure 3.23) and thus clearly favours the motion compensated data set for a clinical evaluation.

Summary

Don't fear failure so much that you refuse to try new things. The saddest summary of a life contains three descriptions: could have, might have, and should have.

(Louis E. Boone)

Patient motion during a medical examination has a severe influence on the usability of the data of a PET acquisition. It does not only result in artefacts in the reconstructed images (e.g. motion blurring), it also makes a precise delineation of anatomic structures difficult. This is particularly critical if a *region-of-interest* (ROI) delineation is performed for a quantification of the tracer accumulation in these regions. Therefore, if motion occurs, it is not only the image quality that is degraded. Also the metabolic information derived from the PET examination is biased by patient motion. As the quantification of the metabolism *in vivo* is one of the strengths of PET and because the resolution of new scanners is constantly increasing, patient movements are becoming more and more crucial.

For the compensation of the influence of patient motion a number of methods have already been discussed in the existing literature. While there exists a variety of different motion compensation methods, they all have in common that they compensate either for non-rigid or rigid motion. Non-rigid (cardiac and respiratory) motion is mostly dominant in whole-body and cardiac PET but absent in brain PET imaging. Because we concentrated on the compensation of motion for brain PET imaging, we therefore considered the development of an event-based compensation method for applying corrections for rigid body motion only.

We have demonstrated that one prerequisite for an accurate motion compensation is the acquisition of PET data in list-mode. However, many PET scanners do not support the acquisition of list-mode in a clinical environment. In section 2.2 we have presented two methods which allow for the acquisition of list-mode with the widely used ECAT Exact HR⁺ PET scanner type. A clinical acquisition of list-mode has been accomplished by applying modifications and enhancements to the hardware of the *acquisition control system* (ACS2) of this PET scanner. As shown in section 2.2.1, the use of a *shared storage device* in combination with a Linux-based computer system accelerates the access to acquisition data. Apart from providing faster access to list-mode data, this method also increases the general stability of clinical acquisitions performed with this scanner type.

Furthermore, we have developed an additional acquisition method which uses digital acquisition cards (DAQ) to gain real-time access to the coincidence data. These cards are put into a Linux system and read the coincidence data directly from an external data bus of the ACS2. With the help of an optocoupler-driven adapter card this solution requires only minimal modifications of the original ACS2 hardware. This adapter isolates the DAQ acquisition from the

internal circuits of the ACS2, which is a prerequisite for a clinical application. Our results show (cf. section 3.1), that due to the optimised list-mode acquisition methods the data of a typical acquisition can be accessed after ≈ 1 minute already instead of several hours. In addition, the alternative DAQ method can acquire and process the list-mode data in real-time. These new acquisition methods represent a major improvement for the development of further list-mode-based methods in PET. The faster access in general and the real-time access to list-mode data in particular facilitates the development of new methods such as, e.g., a potential real-time image reconstruction, considerably.

Many of the motion compensation methods in the literature suggest the use of a motion tracking system for recording patient motion. Our study shows that a passive optical tracking system which is normally used for commercial motion capturing (e.g., creating computer animations) can be used in clinical PET. By regularly performing a special *cross-calibration* measurement and by using a graphical application, a transformation matrix is calculated which provides direct coordinate mapping between the two systems via a single spatial transformation. This is substantiated by the evaluation of our cross-calibration method, as shown in section 3.2.1.

When using an optical tracking system a *motion target* is usually attached to the area that is to be tracked. For our tracking of head motion we have therefore developed and evaluated three different motion targets which are attachable to the patient's head. The evaluation of these targets (cf. section 3.2.2) demonstrates that in contrast to the initial expectations a motion target which consists of *safety goggles* shows the best results. We could show that this target is able to transfer the motion with a high accuracy and provides a good patient compliance at the same time.

In addition, we have demonstrated that a quantitative pre-processing of the motion data is essential to assure a fast compensation of motion. Motion thresholds are usually applied to correct only for 'significant' motions during a motion compensation. We have shown that in contrast to other approaches, with a separate threshold for each raw motion parameter (three translations, three rotations), our method provides more intuitive criteria for assessing the significance of motion. By applying the raw motion data to the surface of an imaginary patient head, the combinatorial connection between the raw motion parameters is clarified and a direct *motion distance* value is calculated over time. This means that only a single threshold is applied and provides a direct test of the motion magnitude instead of using the six ambiguous raw motion parameters separately. Part of this motion analysis is also the generation of a clinical report document (cf. figure 3.2.3). This routinely generated document contains plots and information on the patient's motion over time. A summary page of this document contains information which is also to be used as a quality control for the clinical evaluation of patient motion. It includes a classification of the patient motion according to a grading system with marks from one (no motion) to five (high motion).

After establishing the aforementioned prerequisites, we have developed an optimised event-driven motion compensation method. As explained in section 1.2.3, this type of compensation spatially transforms each *line-of-response* (LOR) according to transformation matrices calculated throughout the aforementioned motion analysis. A special attribute of our optimised event-driven motion compensation is the incorporation of three post-correction methods which

are necessary because of geometric and data storage considerations of a PET system in general. One of these methods corrects the differences in the detector sensitivities when reorienting LORs in the FOV. The results of this optimised *normalisation correction* show (cf. section 3.3.1), that an accurate and direct correction of the detector sensitivities within the motion compensation process is feasible and facilitates the clinical use of such a motion compensation.

Another optimised post-correction method is the *LOR discretisation correction*. This correction method has been explained in section 2.4.2. It considers an LOR as a volume confined by the limits of two detectors instead of the usual approximation as a simple line between these detectors. After spatial transformation of this cuboid LOR volume, the intersection area of the volume with the detectors on the PET ring is calculated. The overlap of the intersection area with the known boundaries of each detector is used to calculate weighting factors for a proportional assignment of the count of an LOR to new detector pairs (cf. figure 2.23). In section 3.3.2 we have shown, that this detailed overlap weighting improves the accurate assignment of counts considerably and thus the accuracy of the motion compensation in general.

The final post-correction method is an optimised correction for lost counts due to motion compensated LORs falling outside the FOV. As shown in section 2.4.3, the *Out-of-FOV correction* scales the counts of a specific LOR according to the time this LOR was found to be outside the FOV. However, we could show that this pure scaling normally leads to additional image artefacts due to the scaling of the statistical error in the measured counts. Therefore, our OFOV correction calculates these scaling factors relative to optimised reference positions. As shown in section 3.3.3, our method minimises the scaling factors and thus the number of image artefacts considerably. This has many advantages over other event-driven motion compensation methods – especially with regard to the additional OFOV-based image artefacts. They have been one of the main obstacles for clinical use of an event-driven motion compensation.

The development of software tools for the integration of the aforementioned methods in clinical PET has been an essential part of this thesis (cf. section 2.5). The standard PET scanner environment does usually not provide the necessary tools for the clinical acquisition and processing of list-mode data. Software tools are missing for, e.g., the processing of a list-mode acquisition as well as for the processing of a motion tracking acquisition. Apart from the development of acquisition related tools, we have therefore developed a complete software package with seven tools for the processing of list-mode and motion tracking data in general. These list-mode processing tools, such as the implementation of the aforementioned event-driven motion compensation, have been also optimised for multiprocessor systems. The results in section 3.4 show that because of the increase in speed, the additional time required for a list-mode or motion compensation processing is no longer an obstacle for clinical use.

Within the final steps of our study we have evaluated the accuracy and the general functioning of the aforementioned methods in our clinical environment. As shown in section 2.6, two main evaluations have been performed for a qualitative and quantitative analysis of these methods. In the first evaluation we tested the general *proof-of-concept* of the event-driven motion compensation via a motion affected point source examination. The reconstructed images show a clear restoration of the image quality after motion correction (cf. section 3.5). Furthermore, the quantitative results show that the reconstructed image resolution is reduced only marginally

due to residual inaccuracies in the motion correction process itself. Relative to the nominal resolution of the PET scanner, however, this minor loss in resolution is negligible because of the general improved quality of the motion corrected image. This is also supported by our second evaluation with data from a clinical patient examination (cf. section 2.6). In addition to a similar improvement in image quality, this evaluation shows that our motion compensation improves the accuracy of a quantitative tracer kinetics analysis. A quantitative analysis based on an *irreversible reference tissue two compartment model*, shows (section 3.5) that the motion compensation affects the values on the tracer accumulation of [^{18}F]DOPA considerably. In fact, we could show that our event-driven motion compensation is not only able to produce accurate results, but also to restore patient data which otherwise would have not been clinically usable because of the patient motion.

Given the final evaluation of the aforementioned methods, the motion tracking and motion analysis have been established as standard methods for monitoring the patient's head motion in our clinical environment. In addition, the optimised list-mode acquisition and the event-driven motion compensation method have been integrated in the standard protocols of our PET scanner so that a motion corrected data set is automatically generated for every brain acquisition. This clinical integration shows that the motion tracking of the patient, in addition to the motion analysis, already helps considerably to classify patient motion into 'significant' or 'non-significant'. Furthermore, the additional possibility to compare a potentially motion affected PET data set to a motion corrected data set has increased the examination flexibility of the physicians at our facility. The importance of list-mode processing and the need for motion compensation methods in PET is also shown by the fact that we have been able to integrate and test our methods at two other institutions throughout the world. The fast access to list-mode data via the aforementioned *shared storage device* solution has been implemented at the PET centre of the *Forschungszentrum Jülich*, Germany in 2006 [106] as well as at the PET centre of the *Columbia University*, New York, U.S.A. in 2008 [126]. Furthermore, our list-mode-based software tools, i.e., the list-mode sorter (`lmSorter`) and event-driven motion compensation tool (`lmMover`) are currently used within official collaborations with these institutions.

Even if the discussed methods improve the PET imaging in case of patient motion considerably, a combination of our motion compensation with an LOR-based image reconstruction would be a possible future improvement. This would also simplify certain aspects of our motion compensation. The *normalisation correction*, for example, could then be performed directly in the image reconstruction process. However, the problem of a potentially over-scaling of sinogram bins still exists because of the *Out-of-FOV correction*. To finally resolve this problem, it might be addressed directly in a potential list-mode-based image reconstruction by virtually expanding the FOV [127]. In addition, a straight forward transfer of the algorithms to different PET systems is possible because of the portable implementation of our motion analysis and event-driven motion compensation algorithms. This could be of particular interest because of recent development in the field of high resolution PET. Also recent developments of multi-modality systems (e.g. PET/MRI systems) show that the motion of a patient is an increasing dilemma in medical imaging. Therefore, motion compensation solutions, such as the one presented in this thesis, are more and more important for precise insights into the human body.

Theses

1. Patient motion severely affects the quantitative accuracy of *positron emission tomography* (PET). It also reduces the image quality by causing image artefacts and blurring.
2. An *event-driven motion compensation* provides the most accurate correction of patient motion. It is superior to other compensation techniques such as *gated acquisition* or *multiple acquisition framing*.
3. The functionality to access the *list-mode* data of a PET scanner which is provided by manufacturers is often severely limited. This hinders the application of event-driven methods, especially for clinical use.
4. The developed *shared storage device acquisition* method speeds up the access to acquisition data of an ECAT Exact HR⁺ PET scanner considerably. It also solves existing stability problems with the *acquisition control system* (ACS2) of this PET scanner type.
5. With the *DAQ-based data acquisition* method *list-mode* data can be acquired and processed in real-time. This allows to operate an HR⁺ in conventional *histogram-mode* while transparently acquiring and processing list-mode data in parallel.
6. Optical tracking systems provide a high accuracy for the tracking of patient motion. They are superior to conventional motion quantification methods which assess the patient motion directly in the image data (e.g. *radioactive markers*).
7. With the use of an external tracking device a method for the coordinate mapping between the tracking system and PET system is required. The *cross-calibration* method is capable of mapping the patient motion from the tracking coordinate system to the PET coordinate system with high accuracy.
8. A *safety goggles-based motion target* provides high accuracy for the tracking of head motion. It combines good motion transferability with good patient compliance and has been well accepted in clinical routine.
9. A *motion analysis* method has been developed which enables an accurate assessment of the magnitude of three-dimensional head motion by explicit computation of Euclidean distances between corresponding points. It is superior to the usual approach of analysing the six transformation parameters (translations/rotations) separately.
10. The automatic generation of a *motion report document* enables the clinical use of the motion analysis. It combines all relevant data and allows an easy assessment of the magnitude of the patient motion.

11. An *event-driven motion compensation* has been developed which is superior to previous approaches. It has been optimised for speed and resolves shortcomings in the post-corrections which previously prevented its clinical use.
12. Clinical use of the event-driven techniques is possible but requires adequate software tools and graphical user interfaces.
13. The optimised motion compensation method is capable of substantially improving the image quality of severely motion affected clinical studies. It also restores the ability of quantitative evaluation of these studies.

List of Figures

1.1	Positron-electron annihilation	7
1.2	PET data registration flow-chart	8
1.3	Photomultiplier tube (PMT)	9
1.4	Detector block and ring layout	10
1.5	Scattered coincidences	12
1.6	Random coincidences	13
1.7	2D/3D mode, septa rings	14
1.8	Interleaving technique illustrated	14
1.9	One-dimensional projection of a 2D object	15
1.10	Correlation between two-dimensional projections and oblique sinograms	16
1.11	List-mode data format	17
1.12	3D central-section-theorem	19
1.13	Comparison of analytic (FBP) and iterative reconstruction (OSEM)	21
1.14	Motion susceptibility illustrated	22
1.15	Motion quantification nomenclature	24
1.16	Analogue motion sensors	27
1.17	Difference between active and passive optical motion tracking	28
1.18	Optical motion tracking body/target	28
1.19	Motion compensation via ‘Gating’	30
1.20	Motion compensation via ‘Multiple Acquisition Framing’ (MAF)	31
1.21	Event-driven motion compensation algorithm	33
1.22	The anatomy of a human skull	35
1.23	Medical imaging plane nomenclature	35
2.1	ECAT Exact HR ⁺ PET scanner – overview	37
2.2	ECAT Exact HR ⁺ PET scanner – inner part	38
2.3	ECAT Exact HR ⁺ PET scanner – coordinate system	38
2.4	ECAT Exact HR ⁺ PET scanner – interleaving implementation	39
2.5	ECAT Exact HR ⁺ PET scanner – attenuation measurement	40
2.6	ARTtrack motion tracking system	41
2.7	Hardware limitations of the acquisition control system (ACS2)	42
2.8	Shared storage device acquisition	43
2.9	Data access to shared storage device	44
2.10	Schematic outline of ACS2 hardware	46
2.11	Parallel data acquisition using DAQ cards	47
2.12	DAQ hardware adapter card	47

2.13	ECAT list-mode command flow-chart	50
2.14	Cross-calibration method illustrated	53
2.15	Maxilla-based reference motion target	57
2.16	Ski goggles-based head target	57
2.17	Os frontale-based head target	58
2.18	Safety goggles-based head target	59
2.19	Motion target evaluation protocol	60
2.20	Position-relative motion analysis method illustrated	61
2.21	LOR discretisation problem illustrated	66
2.22	LOR discretisation post-correction illustrated	67
2.23	Optimised LOR intersection calculation	68
2.24	LOR discretisation correction – Evaluation	69
2.25	Out-of-FOV problem illustrated – I	71
2.26	Out-of-FOV problem illustrated – II	72
2.27	Out-of-FOV problem illustrated – III	73
2.28	Out-of-FOV optimised motion analysis	75
2.29	Out-of-FOV optimised motion compensation	76
2.30	Clinical Integration – software tools workflow	79
2.31	Clinical Evaluation – performed point source movement	82
2.32	Clinical Evaluation – patient movement	84
2.33	Clinical Evaluation – ROI delineation procedure	84
3.1	Shared Storage Device Results – data transfer speed comparison	88
3.2	Event-based attenuation measurement – intensity correlation plot results	90
3.3	Graphical user interface – trkCalib	91
3.4	Cross-calibration method evaluation results	92
3.5	Motion Target Evaluation Results – ski goggles target	94
3.6	Motion Target Evaluation Results – Os frontale target	95
3.7	Motion Target Evaluation Results – safety goggles target	96
3.8	Motion Target Evaluation Results – summary	97
3.9	Motion Analysis Report Document – raw motion data	99
3.10	Motion Analysis Report Document – position-relative data	100
3.11	Motion Analysis Report Document – sphere analysis	101
3.12	Motion Analysis Report Document – summary	102
3.13	Normalisation correction results	103
3.14	LOR discretisation correction results	104
3.15	Out-of-FOV correction results – I	106
3.16	Out-of-FOV correction results – II	107
3.17	Graphical user interface - trkacq/lmacq	108
3.18	ECAT protocol screenshots	109
3.19	Optimisation results – sorting/motion Compensation	109
3.20	Command-line interface - lmMover	110
3.21	Point Source Evaluation – qualitative results	111

3.22 Patient Data Evaluation – qualitative results	112
3.23 Patient Data Evaluation – quantitative TAC plot	113
3.24 Patient Data Evaluation – quantitative R_0k_3 plot	114
3.25 Patient Data Evaluation – parametric images	115

List of Tables

1.1	Motion Type Classification	23
2.1	List-mode Processing Command Overview	49
2.2	Motion Analysis Scores	64
2.3	Out-of-FOV Evaluation Study Parameters	77
2.4	Patient Examinations Parameters	85
3.1	Shared Storage Device Evaluation Results	87
3.2	DAQ-based Data Acquisition Evaluation Results	89
3.3	Event-based Attenuation Measurement Evaluation Results	90
3.4	Cross-Calibration Evaluation Results	92
3.5	Motion Target Evaluation Results	93
3.6	OFOV scaling factor distribution	105
3.7	Quantitative Point Source Evaluation Results	111
3.8	Tracer Quantification Results	114
4.1	Motion Quantification Method Comparison	120

Abbreviations

A/D	<i>Analogue/Digital</i>
ACS2	<i>Acquisition Control System 2</i>
APD	<i>Avalanche photodiode</i>
CAPP	<i>Clinical Applications Programming Package</i>
CCD	<i>Charge-Coupled Device</i>
CT	<i>Computed Tomography</i>
DAQ	<i>Digital Acquisition</i>
DC	<i>Direct Current</i>
DOPA	<i>Dihydroxyphenylalanine</i>
ECG	<i>Electrocardiogram</i>
EM	<i>Expectation-Maximization</i>
ETA	<i>Estimated Time of Arrival</i>
FAT	<i>File Allocation Table</i>
FBP	<i>Filtered-backprojection</i>
FDG	<i>Fluorodeoxyglucose (2-fluoro-2-deoxy-D-glucose)</i>
FOV	<i>Field-of-View</i>
FT	<i>Fourier transform</i>
FTP	<i>File Transfer Protocol</i>
FWHM	<i>Full Width at Half Maximum</i>
LE	<i>Lower Energy</i>
LOR	<i>Line-of-Response</i>
MAF	<i>Multiple Acquisition Framing</i>
ML-EM	<i>Maximum-Likelihood Expectation-Maximization</i>
MRI	<i>Magnetic Resonance Imaging</i>

OFOV	<i>Out-of-FOV</i>
OMFD	<i>3-O-methyl-6-(18)F-fluoro-L-dopa</i>
OSEM	<i>Ordered Subset Expectation-Maximization</i>
PCI	<i>Peripheral Component Interconnect</i>
PET	<i>Positron Emission Tomography</i>
PMT	<i>Photomultiplier tube</i>
PSF	<i>Point Spread Function</i>
PVC	<i>Polyvinyl Chloride</i>
RAID	<i>Redundant Array of Independent Disks</i>
ROI	<i>Region-of-Interest</i>
ROVER	<i>ROI Visualization, Evaluation and Registration</i>
RTP	<i>Radiotherapy Planning</i>
SCSI	<i>Small Computer System Interface</i>
SiPM	<i>Silicon photomultiplier</i>
SMA	<i>Simple Moving Average</i>
SNR	<i>Signal-To-Noise Ratio</i>
SPM	<i>Statistical Parametric Mapping</i>
SVD	<i>Singular Value Decomposition</i>
TAC	<i>Time-Activity-Curve</i>
TTL	<i>Transistor-Transistor-Logic</i>
UDP	<i>User Data Protocol</i>
VME	<i>Versa Modular Eurocard</i>
VSF	<i>Virtual Sorter Board</i>

Bibliography

- [1] R. Standke, “Technische Grundlagen der 18 F-Fluorodeoxyglukose-Positronen-emissionstomographie-Diagnostik,” *Acta Medica Austriaca*, vol. 29, no. 5, pp. 149–155, 2002. [1](#)
- [2] B. Pichler, M. Judenhofer, C. Catana, J. Walton, M. Kneilling, R. Nutt, S. Siegel, C. Claussen, and S. Cherry, “Performance Test of an LSO-APD Detector in a 7-T MRI Scanner for Simultaneous PET/MRI,” *Journal of Nuclear Medicine*, vol. 47, no. 4, pp. 639–647, 2006. [1](#)
- [3] R. Grazioso, N. Zhang, J. Corbeil, M. Schmand, R. Ladebeck, M. Vester, G. Schnur, W. Renz, and H. Fischer, “APD-based PET detector for simultaneous PET/MR imaging,” *Nuclear Inst. and Methods in Physics Research, A*, vol. 569, no. 2, pp. 301–305, 2006. [1](#)
- [4] J. van den Hoff, “Principles of quantitative positron emission tomography,” *Amino Acids*, vol. 29, no. 4, pp. 341–353, 2005. [2](#), [85](#)
- [5] M. Hebert, V. Lowe, J. Hoffman, E. Patz, and M. Anscher, “Positron Emission Tomography in the Pretreatment Evaluation and Follow-Up of Non-Small Cell Lung Cancer Patients Treated with Radiotherapy: Preliminary Findings.” *American Journal of Clinical Oncology*, vol. 19, no. 4, p. 416, 1996. [2](#)
- [6] M. Schulte, D. Brecht-Krauss, M. Werner, E. Hartwig, M. Sarkar, P. Keppler, and J. Kotzerke, “Evaluation of neoadjuvant therapy response of osteogenic sarcoma using FDG PET,” *Journal of Nuclear Medicine*, vol. 40, no. 10, pp. 1637–1643, 1999. [2](#)
- [7] T. Denecke, B. Rau, K. Hoffmann, B. Hildebrandt, J. Ruf, M. Gutberlet, M. Hünnerbein, R. Felix, P. Wust, and H. Amthauer, “Comparison of CT, MRI and FDG-PET in response prediction of patients with locally advanced rectal cancer after multimodal preoperative therapy: Is there a benefit in using functional imaging?” *European Radiology*, vol. 15, no. 8, pp. 1658–1666, 2005. [2](#)
- [8] L. Oehme, B. Beuthien-Baumann, P. Bühler, J. Langner, C. Pöttsch, F. Hofheinz, and J. van den Hoff, “Untersuchung des Einflusses von Kopfbewegungen auf die Quantifizierung von FDOPA-PET-Studien,” in *Nuklearmedizin 2007; 45th Annual DGN Conference, Hannover*, vol. 46: 49-64. German Society of Nuclear Medicine, April 2007, p. A118 (P108). [2](#)
- [9] V. Bettinardi, R. Scardaoni, M. Gilardi, G. Rizzo, D. Perani, E. Paulesu, G. Striano, F. Trulzi, and F. Fazio, “Head Holder for PET, CT, and MR Studies.” *Journal of Computer Assisted Tomography*, vol. 15, no. 5, p. 886, 1991. [2](#), [5](#)
- [10] S. Shrawder, G. Lapin, C. Allen, N. Vick, and D. Groothuis, “A New Head Holder for Reducing Axial Movement and Repositioning Errors during Physiological CT Imaging.” *Journal of Computer Assisted Tomography*, vol. 18, no. 2, p. 329, 1994. [2](#), [5](#)
- [11] R. Sweeney, R. Bale, M. Vogele, M. Nevinny-Stickel, A. Bluhm, T. Auer, G. Hessenberger, and P. Lukas, “Repositioning Accuracy: Comparison of a Noninvasive Head Holder With Thermoplastic Mask for Fractionated Radiotherapy and a Case Report,” *International Journal of Radiation Oncology Biology Physics*, vol. 41, pp. 475–484, 1998. [2](#), [5](#)

- [12] U. Just, “Listmode-basierte Verfahren zur Erfassung und Korrektur von Herz- und Atemzyklus-korrelierten Bewegungen in der Positronen-Emissions-Tomographie,” Ph.D. dissertation, Medizinische Fakultät Carl Gustav Carus der Technischen Universität Dresden, 2007. [2](#), [27](#), [120](#)
- [13] Y. Picard and C. Thompson, “Motion correction of PET images using multiple acquisition frames,” *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 137–144, 1997. [2](#), [5](#), [29](#), [119](#), [120](#)
- [14] L. Tellmann, R. Fulton, K. Bente, I. Stangier, O. Winz, U. Just, H. Herzog, and U. Pietrzyk, “Motion correction of head movements in PET: realisation for routine usage,” *Nuclear Science Symposium Conference Record, 2003 IEEE*, vol. 5, 2003. [2](#), [6](#), [119](#), [120](#)
- [15] P. Bühler, U. Just, E. Will, J. Kotzerke, and J. van den Hoff, “An accurate method for correction of head movement in PET,” *IEEE Transactions on Medical Imaging*, vol. 23, no. 9, pp. 1176–1185, 2004. [2](#), [4](#), [5](#), [32](#), [67](#), [70](#), [120](#), [125](#), [126](#)
- [16] P. Bloomfield, T. Spinks, J. Reed, L. Schnorr, A. Westrip, L. Livieratos, R. Fulton, and T. Jones, “The design and implementation of a motion correction scheme for neurological PET,” *Physics in Medicine and Biology*, vol. 48, no. 8, pp. 959–978, 2003. [2](#), [5](#), [119](#)
- [17] D. Murray, A. Basu, M. Detwiler, and B. Richmond, “Motion tracking with an active camera,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 5, pp. 449–459, 1994. [5](#), [25](#), [120](#)
- [18] H. Strasburger and D. Klenk, “Opto-Electronic Belts for Recording Respiration in Psychophysiological Experimentation and Therapy,” *Psychophysiology*, vol. 20, no. 2, pp. 230–238, 1983. [5](#)
- [19] Y. Picard and C. Thompson, “Digitized video subject positioning and surveillance system for PET,” *IEEE Transactions on Nuclear Science*, vol. 42, no. 4 Part 1, pp. 1024–1029, 1995. [5](#)
- [20] B. Lopresti, A. Russo, W. Jones, T. Fisher, D. Crouch, D. Altenburger, and D. Townsend, “Implementation and performance of an optical motion tracking system for high resolution brain PET imaging,” *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, pp. 2059–2067, 1999. [5](#)
- [21] M. Daube-Witherspoon, Y. Yan, M. Green, R. Carson, K. Kempner, and P. Herscovitch, “Correction for motion distortion in PET by dynamic monitoring of patient position,” *Journal of Nuclear Medicine*, vol. 31, no. 5, p. 816, 1990. [5](#)
- [22] M. Menke, M. Atkins, K. Buckley, and V. TRIUMF, “Compensation methods for head motion detected during PET imaging,” *IEEE Transactions on Nuclear Science*, vol. 43, no. 1, pp. 310–317, 1996. [5](#)
- [23] J. Langner, P. Bühler, U. Just, C. Pötzsch, E. Will, and J. van den Hoff, “Optimized List-Mode Acquisition and Data Processing Procedures for ACS2 Based PET Systems,” *Z. Med. Phys.*, vol. 16, no. 1, pp. 75–82, 2006. [6](#), [81](#)
- [24] J. Langner, C. Pötzsch, E. Will, and J. van den Hoff, “Entwicklung von Methoden zur echtzeitfähigen Listmodeakquisition an ACS2-basierten PET-Scannern,” in *Nuklearmedizin 2006; 44th Annual DGN Conference, Berlin*, vol. 44 (25). German Society of Nuclear Medicine, April 2006, p. A50. [6](#)

- [25] J. Langner, P. Bühler, U. Just, C. Pötzsch, E. Will, and J. van den Hoff, "Optimized List-Mode Acquisition and Data Processing Procedures for ACS2 Based PET Systems," *Nuclear Science Symposium Conference Record, 2006 IEEE*, November 2006. 6
- [26] J. Langner, C. Pötzsch, P. Bühler, and J. van den Hoff, "Entwicklung von Methoden zum beschleunigten Zugriff auf Akquisitionsdaten ACS2-basierter PET-Scanner," in *Nuklearmedizin 2005; 43th Annual DGN Conference, Basel*, vol. 43. German Society of Nuclear Medicine, April 2005, p. A161. 6
- [27] J. Langner, "A reverse engineered implementation of the VxWorks 5.2+ extended DOS filesystem for Linux kernels 2.6+," 2005. [Online]. Available: http://www.jens-langner.de/vxext_fs/ 6, 44, 118
- [28] J. Langner, "Dateisysteme: FUSE, VXEXT und JFFS," *Linux Magazin*, vol. 06/05, p. 20, Jun. 2005, ISSN 1432-640X. 6, 118
- [29] J. Langner, L. Oehme, C. Pötzsch, S. Dittrich, B. Beuthien-Baumann, and J. van den Hoff, "Methods for an Automatic Analysis of Motion Tracking Data in PET," in *EANM Annual Conference, Munich, Germany*. European Association of Nuclear Medicine, October 2008. 6
- [30] J. Langner, L. Oehme, B. Beuthien-Baumann, and J. van den Hoff, "Quantitative Auswertemethoden für die Beurteilung von Patientenbewegungen in der klinischen PET," in *Nuklearmedizin 2007; 45th Annual DGN Conference, Hannover*, vol. 46: 49-64. German Society of Nuclear Medicine, April 2007, p. A34 (V78). 6
- [31] J. Langner, "Development of a Parallel Computing Optimized Head Movement Correction Method in Positron Emission Tomography," Master's thesis, University of Applied Sciences Dresden and Research Center Dresden-Rossendorf, 2003. [Online]. Available: <http://www.jens-langner.de/ftp/MScThesis.pdf> 6, 8, 32, 52, 57, 67
- [32] J. Langner, P. Bühler, C. Pötzsch, and J. van den Hoff, "Nutzung von Mehrprozessorsystemen für den routinefähigen Einsatz listmode-basierter Bewegungskorrektur in der PET," in *Nuklearmedizin 2004; 42th Annual DGN Conference, Rostock*, vol. 43. German Society of Nuclear Medicine, March 2004, p. V189. 6
- [33] J. Langner, S. Dittrich, C. Pötzsch, B. Beuthien-Baumann, and J. van den Hoff, "Reduzierung des Out-of-Field-of-View Einflusses auf die Event-basierte Bewegungskorrektur von Hirnuntersuchungen in der PET," in *Nuklearmedizin 2008; 46th Annual DGN Conference, Leipzig*, vol. 47: 65-85. German Society of Nuclear Medicine, April 2008, p. A59 (V136). 6
- [34] J. Langner, H. Mölle, S. Dittrich, E. Will, and J. van den Hoff, "Event-by-Event Attenuation Measurement for ACS2-Based PET Systems," *Nuclear Science Symposium Conference Record, 2008 IEEE*, October 2008. 6
- [35] K. Thielemans, S. Mustafovic, and L. Schnorr, "Image reconstruction of motion corrected sinograms," *Nuclear Science Symposium Conference Record, 2003 IEEE*, vol. 4, 2003. 6, 119
- [36] F. Lamare, M. Ledesma Carbayo, T. Cresson, G. Kontaxakis, A. Santos, C. Cheze LeRest, A. Reader, and D. Visvikis, "List-mode-based reconstruction for respiratory motion correction in PET using non-rigid body transformations," *Physics in Medicine and Biology*, vol. 52, no. 17, pp. 5187–5204, 2007. 6, 17, 119

- [37] A. Rahmim, P. Bloomfield, S. Houle, M. Lenox, C. Michel, K. Buckley, T. Ruth, and V. Sossi, "Motion compensation in histogram-mode and list-mode EM reconstructions: beyond the event-driven approach," *IEEE Transactions on Nuclear Science*, vol. 51, no. 5 Part 2, pp. 2588–2596, 2004. [6](#), [119](#)
- [38] A. Rahmim, K. Dinelle, J. Cheng, M. Shilov, W. Segars, S. Lidstone, S. Blinder, O. Rousset, H. Vajihollahi, B. Tsui *et al.*, "Accurate Event-Driven Motion Compensation in High-Resolution PET Incorporating Scattered and Random Events," *IEEE Transactions on Medical Imaging*, vol. 27, no. 8, pp. 1018–1033, 2008. [6](#)
- [39] S. Woo, H. Watabe, Y. Choi, K. Kim, C. Park, P. Bloomfield, and H. Iida, "Sinogram-based motion correction of PET images using optical motion tracking system and list-mode data acquisition," *IEEE Transactions on Nuclear Science*, vol. 51, no. 3, pp. 782–788, 2004. [6](#), [119](#)
- [40] C. D. Anderson, "The positive electron," *Phys. Rev.*, vol. 43, no. 6, pp. 491–494, Mar 1933. [7](#)
- [41] O. Warburg, "Über den Stoffwechsel der Carcinomzelle," *Die Naturwissenschaften*, vol. 12, no. 50, pp. 1131–1137, 1924. [7](#)
- [42] O. Warburg, "The metabolism of tumors," *London: Constable*, vol. 6, 1930. [7](#)
- [43] J. Volker, H. Hodge, H. Wilson, and S. Van Voorhis, "The absorption of fluorides by enamel, dentin, bone and hydroxyapatite as shown by the radioactive isotope," *Journal of Biological Chemistry*, vol. 134, no. 2, pp. 543–548, 1940. [7](#)
- [44] R. Brooks and G. Chiro, "Principles of computer assisted tomography (CAT) in radiographic and radioisotopic imaging," *Physics in Medicine and Biology*, vol. 21, no. 5, pp. 689–732, 1976. [7](#), [18](#)
- [45] T. Ido, C. Wan, J. Fowler, and A. Wolf, "Fluorination with molecular fluorine. A convenient synthesis of 2-deoxy-2-fluoro-D-glucose," *The Journal of Organic Chemistry*, vol. 42, no. 13, pp. 2341–2342, 1977. [7](#)
- [46] A. van Baardwijk, B. Baumert, G. Bosmans, M. van Kroonenburgh, S. Stroobants, V. Gregoire, P. Lambin, and D. De Ruyscher, "The current status of FDG-PET in tumour volume definition in radiotherapy treatment planning," *Cancer Treatment Reviews*, vol. 32, no. 4, pp. 245–260, 2006. [8](#)
- [47] U. Nestle, S. Kremp, A. Schaefer-Schuler, C. Sebastian-Welsch, D. Hellwig, C. Rube, and C. Kirsch, "Comparison of Different Methods for Delineation of 18 F-FDG PET-Positive Tissue for Target Volume Definition in Radiotherapy of Patients with Non-Small Cell Lung Cancer," *Journal of Nuclear Medicine*, vol. 46, no. 8, p. 1342, 2005. [8](#)
- [48] G. Leunens, J. Menten, C. Weltens, J. Verstraete, and E. van der Schueren, "Quality assessment of medical decision making in radiation oncology: variability in target volume delineation for brain tumours." *Radiotherapy and Oncology*, vol. 29, no. 2, pp. 169–75, 1993. [8](#)
- [49] J. Humm, A. Rosenfeld, and A. Guerra, "From PET detectors to PET scanners," *European Journal of Nuclear Medicine and Molecular Imaging*, vol. 30, no. 11, pp. 1574–1597, 2003. [9](#)
- [50] B. Pichler, B. Swann, J. Rochelle, R. Nutt, S. Cherry, and S. Siegel, "Lutetium oxyorthosilicate block detector readout by avalanche photodiode arrays for high resolution animal PET," *Physics in Medicine and Biology*, vol. 49, no. 18, pp. 4305–4319, 2004. [9](#)

- [51] P. Dokhale, R. Silverman, K. Shah, R. Grazioso, R. Farrell, J. Glodo, M. McClish, G. Entine, V. Tran, and S. Cherry, "Performance measurements of a depth-encoding PET detector module based on position-sensitive avalanche photodiode read-out," *Physics in Medicine and Biology*, vol. 49, no. 18, pp. 4293–4304, 2004. 9
- [52] A. Otte, J. Barral, B. Dolgoshein, J. Hose, S. Klemin, E. Lorenz, R. Mirzoyan, E. Popova, and M. Teshima, "A test of silicon photomultipliers as readout for PET," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 545, no. 3, pp. 705–715, 2005. 9
- [53] N. Tomic, C. Thompson, and M. Casey, "Investigation of the "Block Effect" on spatial resolution in PET detectors," *IEEE Transactions on Nuclear Science*, vol. 52, no. 3 Part 1, pp. 599–605, 2005. 9
- [54] S. Cherry, H. Tornai, C. Levin, S. Siegel, and E. Hoffman, "A comparison of PET detector modules employing rectangular and round photomultiplier tubes," *IEEE Transactions on Nuclear Science*, vol. 42, no. 4, pp. 1064–1068, 1995. 9
- [55] P. Valk, D. Bailey, D. Townsend, and M. Maisey, *Positron Emission Tomography: Basic Science and Clinical Practice*. Springer-Verlag, 2003. 10, 21
- [56] M. Casey, H. Gadagkar, and D. Newport, "A component based method for normalization in volume PET," *3rd International Meeting on fully 3D Reconstruction in Nuclear Medicine and Radiology, Aix-les-Bains*, pp. 67–71, 1995. 11
- [57] M. Phelps, E. Hoffman, S. Huang, and M. Ter-Pogossian, "Effect of positron range on spatial resolution," *Journal of Nuclear Medicine*, vol. 16, no. 7, pp. 649–652, 1975. 11
- [58] L. Geworski, S. Müller, and C. Reiners, "Einfluß der Energiefenstereinstellung und der Phantomgeometrie auf die Streustrahlung bei der Positronen-Emissions-Tomographie." *Medizinische Physik*, pp. 302–303, 1993. 12
- [59] B. Bendriem and D. Townsend, *The Theory and Practice of 3D PET*. Kluwer Academic, 1998. 12, 14, 18, 19, 67
- [60] A. Werling, "Modellbasierte Korrektur der Streustrahlung in der Positronen-Emissions-Tomographie," Ph.D. dissertation, Deutsches Krebsforschungszentrum Heidelberg, 2001. 13
- [61] C. Watson, M. Casey, C. Michel, and B. Bendriem, "Advances in scatter correction for 3D PET/CT," *Nuclear Science Symposium Conference Record, 2004 IEEE*, vol. 5, 2004. 13
- [62] P. Ell and B. Holman, *Computed emission tomography*. Oxford University Press, 1982. 13
- [63] E. Hoffman, S. Huang, M. Phelps, and D. Kuhl, "Quantitation in positron emission computed tomography: 4. Effect of accidental coincidences," *Journal of Computer Assisted Tomography*, vol. 5, no. 3, pp. 391–400, 1981. 13
- [64] C. Thompson, S. James, and N. Tomic, "Under-sampling in PET scanners as a source of image blurring," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 545, no. 1-2, pp. 436–445, 2005. 14
- [65] P. Edholm, L. Hellstrom, and B. Jacobson, "Transverse tomography with incoherent optical reconstruction," *Physics in Medicine and Biology*, vol. 23, pp. 90–99, 1978. 15

- [66] F. Natterer, *The Mathematics of Computerized Tomography*. New York: Wiley, 1986. [17](#)
- [67] T. Nichols, "ECAT Exact HR⁺ List Mode Format," 1998, technical note, unpublished. [17](#)
- [68] U. Just, "Sammlung der Einzelereignisse einer PET-Kamera, Listmodemessung und Sortieren nach beliebigen Kriterien," Master's thesis, Hochschule für Technik und Wirtschaft Dresden (FH), 2000. [17](#)
- [69] W. Jones, M. Running, L. Draughn, and J. Reed, "Advanced hardware architecture for on-line data acquisition in clinical 3-D PET: smart DRAM PCI Card for 14 M event/sec histogramming across terabytes of DRAM," *Nuclear Science Symposium Conference Record, 2004 IEEE*, vol. 6, 2004. [17](#)
- [70] D. Kadrmas, "LOR-OSEM: statistical PET reconstruction from raw line-of-response histograms," *Physics in Medicine and Biology*, vol. 49, no. 20, pp. 4731–4744, 2004. [17](#), [119](#)
- [71] A. Reader, S. Ally, F. Bakatselos, R. Manavaki, R. Walledge, A. Jeavons, P. Julyan, S. Zhao, D. Hastings, and J. Zweit, "One-Pass List-Mode EM Algorithm for High-Resolution 3-D PET Image Reconstruction Into Large Arrays," *IEEE Transactions on Nuclear Science*, vol. 49, no. 3, p. 693, 2002. [17](#), [119](#)
- [72] A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988. [18](#)
- [73] L. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Transactions on Medical Imaging*, vol. 1, no. 2, pp. 113–122, 1982. [20](#)
- [74] A. Dempster, N. Laird, and D. Rubin, "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 39, no. 1, pp. 1–38, 1977. [20](#)
- [75] H. Hudson and R. Larkin, "Accelerated image reconstruction using ordered subsets of projection data," *IEEE Transactions on Medical Imaging*, vol. 13, no. 4, pp. 601–609, 1994. [20](#)
- [76] H. Barrett, D. Wilson, and B. Tsui, "Noise properties of the EM algorithm. I. Theory," *Physics in Medicine and Biology*, vol. 39, no. 5, pp. 833–846, 1994. [20](#)
- [77] E. Veklerov and J. Llacer, "Stopping rule for the MLE algorithm based on statistical hypothesis testing," *IEEE Transactions on Medical Imaging*, vol. 6, no. 4, pp. 313–319, 1987. [20](#)
- [78] K. McLeish, D. Hill, D. Atkinson, J. Blackall, and R. Razavi, "A study of the motion and deformation of the heart due to respiration," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1142–1150, 2002. [23](#)
- [79] A. Naum, M. Laaksonen, H. Tuunanen, V. Oikonen, M. Teräs, J. Kempainen, M. Järvisalo, P. Nuutila, and J. Knuuti, "Motion detection and correction for dynamic 15 O-water myocardial perfusion PET studies," *European Journal of Nuclear Medicine and Molecular Imaging*, vol. 32, no. 12, pp. 1378–1383, 2005. [25](#), [119](#), [120](#)
- [80] C. Mayer and S. Pietzsch, "Modellbasierte Mimikererkennung in Videobildern mittels Optical Flow," Technische Universität München, Fakultät für Informatik, Tech. Rep., 2006. [25](#), [120](#)

- [81] K. Friston, “Statistical parametric mapping,” *Functional neuroimaging: Technical foundations*, vol. 8, pp. 79–93, 1994. 26, 120
- [82] M. Dawood, N. Lang, X. Jiang, and K. Schäfers, “Lung Motion Correction on Respiratory Gated 3-D PET/CT Images,” *IEEE Transactions on Medical Imaging*, vol. 25, no. 4, 2006. 26, 119, 120
- [83] Vernier Software and Technology, “Respiration monitor belt.” [Online]. Available: <http://www.vernier.com/probes/rmb.html> 26
- [84] AD Instruments, “Piezo respiratory belt transducer.” [Online]. Available: <http://www.adinstruments.com/products/product.php?id=MLT1132> 26
- [85] R. Beach, P. Bruyant, P. Fulton, R. Gennert, M. Nadella, and S. King, “Feasibility of stereo-infrared tracking to monitor patient motion during cardiac SPECT imaging,” *IEEE Transactions on Nuclear Science*, vol. 51, no. 5, pp. 2693–2698, 2004. 27, 120
- [86] Inition Corporation, “Motion capturing / 3d tracking.” [Online]. Available: <http://www.inition.co.uk> 27, 120
- [87] Advanced Realtime Tracking GmbH, “ARTtrack Optical Tracking System.” [Online]. Available: <http://www.ar-tracking.de/> 27, 28, 41, 122
- [88] Northern Digital Inc. (NDI), “Polaris optical tracking system.” [Online]. Available: <http://www.ndigital.com/medical/polarisfamily.php> 27, 122
- [89] A. Montgomery, K. Thielemans, M. Mehta, F. Turkheimer, S. Mustafovic, and P. Grasby, “Correction of Head Movement on PET Studies: Comparison of Methods,” *Journal of Nuclear Medicine*, vol. 47, no. 12, p. 1936, 2006. 29
- [90] M. Dawood, F. Büther, N. Lang, O. Schober, and K. Schäfers, “Respiratory gating in positron emission tomography: A quantitative comparison of different gating schemes,” *Medical Physics*, vol. 34, p. 3067, 2007. 29
- [91] J. Ashburner and K. Friston, “Multimodal Image Coregistration and Partitioning—A Unified Framework,” *Neuroimage*, vol. 6, no. 3, pp. 209–217, 1997. 32
- [92] R. Fulton, I. Nickel, L. Tellmann, S. Meikle, U. Pietrzyk, and H. Herzog, “Event-by-event motion compensation in 3D PET,” *Nuclear Science Symposium Conference Record, 2003 IEEE*, vol. 5, 2003. 32, 119
- [93] Wikipedia, “Human skull.” [Online]. Available: http://en.wikipedia.org/wiki/Human_skull 35
- [94] J. Wickett, P. Vernon, and D. Lee, “In vivo brain size, head perimeter, and intelligence in a sample of healthy adult females,” *Personality and individual differences*, vol. 16, no. 6, pp. 831–838, 1994. 34, 61
- [95] R. Pearl, “On the correlation between intelligence and the size of the head,” *Journal of Comparative Neurology and Psychology*, vol. 16, no. 3, pp. 189–199, 1906. 34, 61
- [96] Y. Mrabet, “Wikimedia - Human anatomy planes.” [Online]. Available: http://commons.wikimedia.org/wiki/Image:Human_anatomy_planes.svg 35
- [97] *ECAT Exact HR⁺ ACS2 Service Manual*, Part Number 35 49 743 Rev. 01 ed., Siemens Medical Systems, Inc., Knoxville, TN, 1996. 39

- [98] M. Dahlbom, J. Reed, and J. Young, "Implementation of true continuous bed motion in 2-D and 3-D whole-body PET scanning," *IEEE Transactions on Nuclear Science*, vol. 48, no. 4 Part 2, pp. 1465–1469, 2001. 40
- [99] *SCSI-2 Standard Specification*, ANSI Std. X2.131-1994, 1994. 43
- [100] Microsoft Inc. (2008, March) FAT32 File System Specification. [Online]. Available: <http://www.microsoft.com/whdc/system/platform/firmware/fatgen.mspx> 44
- [101] *VME Standard Specification*, IEEE Std. 1014-1987, IEC821, 1987. 46
- [102] Spectrum Mikroelektronik, GmbH, "The MI-70xx Digital Acquisition Cards," 2008. [Online]. Available: <http://www.spec.de/> 47
- [103] *Specialized Hardware Developments*, Seiler IT Services, Ungerstrasse 3, 30451 Hannover, Germany, 2005. 48
- [104] G. Golub and W. Kahan, "Calculating the singular values and pseudo-inverse of a matrix," *SIAM J. Numer. Anal.*, vol. 2, no. 3, pp. 205–224, 1965. 53, 54
- [105] C. Pötzsch, "Entwicklung volumetrischer Verfahren für die Positronen-Emissions-Tomographie zur quantitativen Bewertung und Abgrenzung von Raumforderungen bei onkologischen Fragestellungen," Ph.D. dissertation, Medizinische Fakultät Carl Gustav Carus der Technischen Universität Dresden, 2007. 52, 83
- [106] Forschungszentrum Jülich, "Institute of Neurosciences and Biophysics," 2008. [Online]. Available: <http://www.fz-juelich.de/> 58, 134
- [107] M. Weyler, "Instrumentelle und methodische Untersuchungen zur Bewegungskorrektur bei Hirn-PET-Messungen," Master's thesis, Fachhochschule Aachen and Forschungszentrum Jülich, 2007. 58, 122
- [108] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2008, ISBN 3-900051-07-0. [Online]. Available: <http://www.R-project.org> 60
- [109] J. Kenney, *Mathematics of statistics*. New York: D. Van Nostrand, 1939. 60
- [110] Qt Software, Nokia Corporation, "Qt Cross-Platform Application Framework," 2008. [Online]. Available: <http://trolltech.com/> 80
- [111] R. Kessler, J. Ellis, and M. Eden, "Analysis of emission tomographic scan data: limitations imposed by resolution and background," *Journal of Computer Assisted Tomography*, vol. 8, no. 3, pp. 514–522, 1984. 83
- [112] F. Hofheinz, S. Dittrich, C. Pötzsch, and J. van den Hoff, "Effects of cold sphere walls in PET phantom measurements on the volume reproducing threshold," *Physics in Medicine and Biology*, vol. submitted, 2008. 83
- [113] V. Holthoff-Detto, J. Kessler, K. Herholz, H. Bonner, U. Pietrzyk, M. Wurker, M. Ghaemi, K. Wienhard, R. Wagner, and W. Heiss, "Functional effects of striatal dysfunction in Parkinson disease," *Archives of Neurology*, vol. 54, no. 2, pp. 145–150, 1997. 84, 85, 114
- [114] L. Oehme, B. Beuthien-Baumann, F. Hofheinz, M. Wolz, C. Pötzsch, J. van den Hoff, and J. Kotzerke, "Vergleich einer 2D- und 3D-basierten ROI-Auswertung beim F-18-DOPA-Hirn-PET in einem Normalkollektiv," in *Nuklearmedizin 2008; 46th Annual DGN Conference, Leipzig*, vol. 47: 65-85. German Society of Nuclear Medicine, April 2008, p. P64. 84, 85

- [115] F. Vingerhoets, B. Snow, M. Schulzer, S. Morrison, T. Ruth, J. Holden, S. Cooper, and D. Calne, "Reproducibility of Fluorine-18-6-Fluorodopa Positron Emission Tomography in Normal Human Subjects," *Journal of Nuclear Medicine*, vol. 35, no. 1, pp. 18–24, 1994. [84](#), [85](#)
- [116] ABX advanced biochemical compounds, Radeberg, Germany, "ROVER – ROI Visualization, Evaluation and Image Registration." [Online]. Available: <http://www.abx.de/rover/> [85](#)
- [117] R. Carson, V. Cunningham, R. Gunn, J. van den Hoff, A. Lammertsma, K. Leenders, W. Müller-Schauenburg, and A. Willemsen, *PET Pharmacokinetic course manual*, Nuclear Science Symposium Conference, 2008 IEEE, Dresden, Germany ed., 2008. [85](#)
- [118] C. Patlak, R. Blasberg, and J. Fenstermacher, "Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data." *Journal of Cerebral Blood Flow & Metabolism*, vol. 3, no. 1, pp. 1–7, 1983. [86](#), [129](#)
- [119] C. Patlak and R. Blasberg, "Graphical evaluation of blood-to-brain transfer constants from multiple-time uptake data. Generalizations." *Journal of Cerebral Blood Flow & Metabolism*, vol. 5, no. 4, pp. 584–90, 1985. [86](#), [129](#)
- [120] F. Lamare, T. Cresson, J. Savean, C. Le Rest, A. Reader, and D. Visvikis, "Respiratory motion correction for PET oncology applications using affine transformation of list mode data," *Physics in Medicine and Biology*, vol. 52, no. 1, pp. 121–140, 2007. [119](#)
- [121] R. Fulton, L. Tellmann, U. Pietrzyk, O. Winz, I. Stangier, I. Nickel, A. Schmid, S. Meikle, and H. Herzog, "Accuracy of motion correction methods for PET brain imaging," *Nuclear Science Symposium Conference Record, 2004 IEEE*, vol. 7, 2004. [119](#)
- [122] A. Rahmim, "Advanced Motion Correction Methods in PET," *Iran Journal of Nuclear Medicine*, vol. 13, no. 24, p. 1, 2005. [119](#)
- [123] R. Fulton, S. Meikle, S. Eberl, J. Pfeiffer, and C. Constable, "Correction for head movements in positron emission tomography using an optical motion-tracking system," *IEEE Transactions on Nuclear Science*, vol. 49, no. 1, pp. 116–123, 2002. [119](#)
- [124] D. Hu, C. Hayden, M. Casey, and Z. Burbar, "Stereo computer vision system for measuring movement of patient's head in PET scanning," *Nuclear Science Symposium Conference Record, 2004 IEEE*, vol. 5, 2004. [120](#)
- [125] L. Adam, J. Zaers, H. Ostertag, H. Trojan, M. Bellemann, and G. Brix, "Performance evaluation of the whole-body PET scanner ECAT EXACT HR following the IEC standard," *IEEE Transactions on Nuclear Science*, vol. 44, no. 3 Part 2, pp. 1172–1179, 1997. [128](#)
- [126] Columbia University, New York, U.S.A., "Columbia Kreitchman PET Center," 2008. [Online]. Available: <http://www.columbiapet.org/> [134](#)
- [127] J. Scheins, M. Axer, U. Pietrzyk, and H. Herzog, "Virtual scanner geometry for iterative 3D PET reconstruction adaptable to complex and irregular scanner topologies," *Nuclear Science Symposium Conference Record, 2007 IEEE*, vol. 5, 2007. [134](#)

Hiermit erkläre ich, dass ich die der Medizinischen Fakultät Carl Gustav Carus der Technischen Universität Dresden zur Promotion eingereichte Dissertation mit dem Titel

Event-Driven Motion Compensation in Positron Emission Tomography: Development of a Clinically Applicable Method

im Institut für Radiopharmazie des Forschungszentrums Dresden-Rossendorf unter der Leitung von Prof. Dr. hum. biol. habil. Dr. rer. nat. Jörg van den Hoff

ohne sonstige Hilfe selbst durchgeführt und bei der Abfassung der Dissertation keine anderen als die dort angeführten Hilfsmittel und Quellen benutzt habe. Ich habe diese Dissertation bisher an keiner anderen Stelle zum Zwecke eines Promotions- oder anderen Prüfungsverfahrens eingereicht. Ich habe bisher an keiner in- oder ausländischen Fakultät ein Gesuch um Zulassung zur Promotion eingereicht.

Dresden, November 2008

A handwritten signature in blue ink, appearing to read 'Jens', with a stylized flourish extending to the right.

Jens Langner