



Technische Universität München
Fakultät für Informatik




IDP - Interdisciplinary Project (Physics)

Respiratory Motion Estimation:
Tests and Comparison of Different Sensors

Michael Riedel



Technische Universität München
Fakultät für Informatik



IDP - Interdisciplinary Project (Physics)

Respiratory Motion Estimation:
Tests and Comparison of Different Sensors

Michael Riedel

Advisor: Prof. Dr. Peter Böni

Supervision: Prof. Dr. Nassir Navab, Axel Martinez-Möller,
Ralph Bundschuh

Date: April 12th, 2006

Abstract

Respiratory motion has negative impact on some medical imaging techniques, in particular on Positron Emission Tomography (PET). This is due to the physical limitations of PET, which lead to long acquisition times of several minutes. For high quality PET images of the thorax area, respiratory motion needs to be measured and corrected. This can be done by respiratory gating, a method dividing the respiratory cycle into phases of similar thorax expansion, called gates. Separate images are then reconstructed from the PET data acquired during each such gate.

Respiration gates can be calculated automatically from a respiratory signal, which has to be measured by an external sensor. The purpose of this work was to evaluate four respiration sensors based on different physical principles: an elastic belt measuring thorax circumference, a spirometer measuring air-flow from the lungs, a thermometer measuring the temperature of respired air, and a stereo infrared camera system tracking thorax motion. Focus was put on the camera system, which delivers raw motion data of markers attached to the thorax, and can therefore be used to estimate motion more generally. A feasible setup of the markers and the camera is described, and several processing steps applied to the raw tracking data are explained. This includes interpolation of temporarily occluded markers and calibration of marker locations with respect to the PET scanner. Methods to digest respiratory signals from the multi-dimensional tracking data are also described.

The respiratory signals acquired with all the sensors were compared both visually and statistically. All sensors were able to track both regular breathing, and breathing artifacts like breath-holds or speaking. Though based on different physical effects, all respiration curves were basically suited for respiratory gating. However, only the elastic belt and the camera system used in this work, were reliable enough for clinical use, and were tested in a preliminary study with cardiac patients undergoing PET. Gated PET images, reconstructed using the respective respiratory signals, showed respiration-induced motion of the heart clearly, though its magnitude was only in the region of half a centimeter.

The camera system also allowed quantification of respiratory motion on the patient's surface, which turned out have a similar magnitude. Measurements of translative body motion not related to respiration showed that it was usually even smaller, and therefore has little impact on PET imaging. However, more general measurements of body motion are discussed, since non-rigid motion could be much more significant. Only if its characteristics are known, can such body motion be corrected in the future.

In conclusion, respiratory gating can improve the quality of PET imaging, no matter what sensor it is based on. However, the camera system is too cumbersome for respiratory gating alone, and other respiration sensors, like the elastic belt are more advisable for clinical use. Yet, the camera system should be used to study motion of a larger number of patients, to estimate its impact on PET imaging.

Contents

1. Introduction	1
1.1. Overview of Medical Imaging with PET	2
1.1.1. Physical Basics of PET	3
1.1.2. PET Scanner Technology	6
1.1.3. Image Reconstruction	10
1.1.4. Limitations of Clinical PET Scanners	13
1.2. Objectives of this IDP	16
1.3. Respiratory Motion Correction in Medical Imaging	17
1.3.1. Gating Systems in Clinical Use	17
1.3.2. Current Research	18
2. Methods for Testing Respiration Sensors	21
2.1. Use of Respiration Sensors in the Context of PET	23
2.2. Anzai Belt	26
2.3. PMM Spirometer	28
2.4. BioVet Temperature Sensor	29
2.5. ART Stereo Infrared Camera	31
2.5.1. Principles of 3D Tracking with ART	33
2.5.2. Clinical Setup for Tracking Respiratory Motion	37
2.5.3. Coordinate System Calibration	39
2.5.4. Processing of Raw ART Camera Data	42
2.5.5. Extraction of Respiration Curves	45
3. Evaluation of Resulting Sensor Data	49
3.1. Comparison of the Sensors	49
3.1.1. Qualitative Comparison	50
3.1.2. Quantitative Comparison	54
3.1.3. Comparison of Gated Images	58
3.2. Motion Estimation Beyond Respiratory Gating	59
4. Conclusion	63
4.1. Outcome	63
4.2. Future Work	64
4.3. Acknowledgments	66
A. ART Camera Software	67

B. Signal Analysis Software	69
Bibliography	73

1. Introduction

There are various medical imaging techniques used for diagnosis in everyday clinical practice. While *Positron Emission Tomography (PET)* is not the most widespread one, it is quite an interesting technology. As a method of nuclear medicine, it uses radioactive isotopes, which can be embedded into various chemical compounds. PET can then detect the distribution of the tracers in the living organism, and visualize it as a 3D image. Hence it provides functional information, whereas many other medical imaging techniques only provide structural information.

At the department of nuclear medicine at *Klinikum München Rechts der Isar (MRI)*, the medical center of *Technische Universität München (TUM)*, several PET scanners are in clinical use. Additionally, an interdisciplinary group of scientists is steadily working on the improvement of PET diagnostics. This includes optimizing protocols for PET examinations, developing radioactive tracers, enhancing data processing and image reconstruction software as well as researching concepts for next generation PET scanners. Many of these efforts aim towards a higher quality of PET images and/or shorter acquisition times.

Alongside the rather low spatial resolution, long acquisition times of typically several minutes are the major problem of PET imaging. Any movement of the patient will blur the resulting images substantially. However, some movements just cannot be avoided during a PET acquisition, particularly respiratory motion. When acquiring images of a patient's thorax, the result shows a blurred overlay of all phases of the respiration cycle, rather than a sharp snapshot.

Therefore, methods to avoid respiration artifacts in PET images are currently being evaluated at the department. *Respiratory gating* is a promising method, that has been proposed for this. Here, the respiratory cycle is divided into several recurring phases, called gates. Then, for each gate, a separate image is reconstructed from the PET data, as shown in figure 1.1. This method implies that the progression of respiration is known. Hence, a patient undergoing PET has to be monitored with appropriate respiration sensors. The purpose of this IDP¹ was the evaluation of different respiration sensors and their applicability for respiratory gating in the context of PET.

This introduction first gives an overview of PET instrumentation in section 1.1. This also explains why respiratory gating is necessary to handle respiration artifacts. Readers already familiar with PET, or not interested in the details, may consider skipping this section. The subsequent sections of this introduction define the scope of this IDP more closely. This includes an outline of the methods described in chapter 2

¹An IDP (inter-disciplinary project) is a part of the computer-science studies at TUM. It is a work related to the respective minor subject, physics in this case.

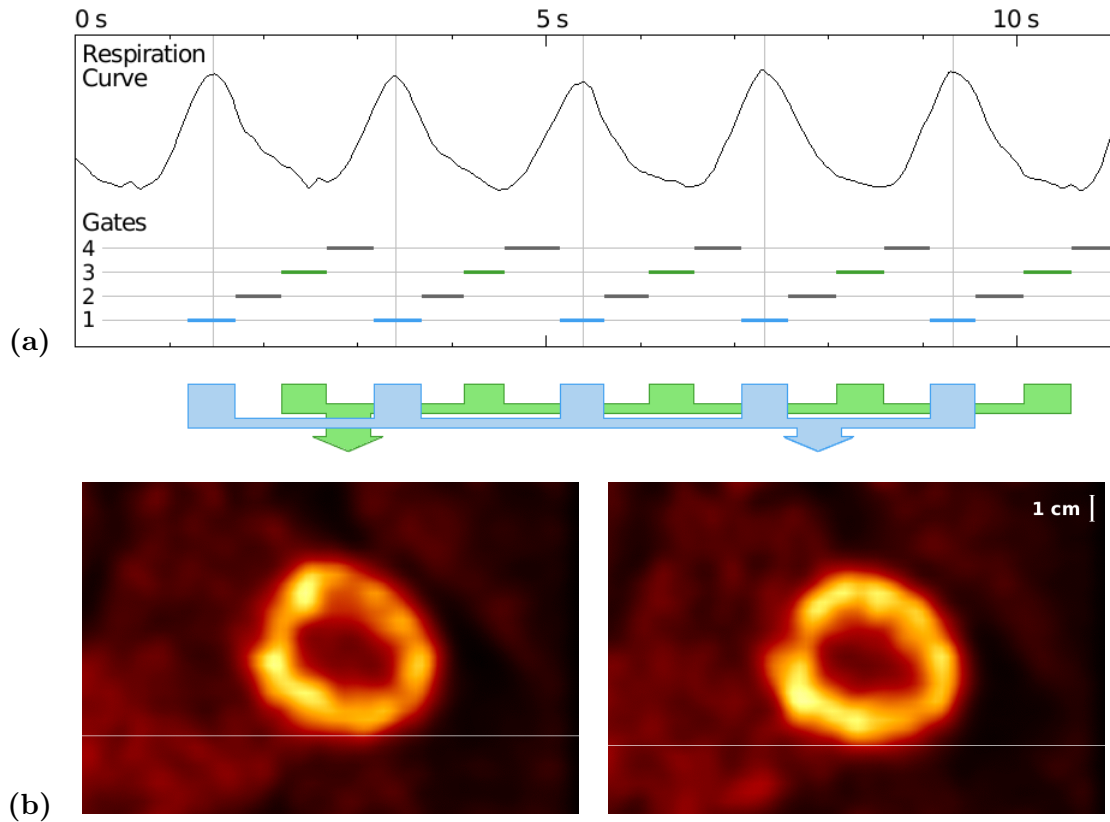


Figure 1.1.: A respiration curve and the corresponding gates (a) can be used to obtain gated PET images of the heart (b).

and the results presented in chapter 3. Finally this introduction gives an overview of other work related to respiratory motion estimation and respiratory gating.

1.1. Overview of Medical Imaging with PET

As already mentioned, a wide range of medical imaging techniques is in use nowadays. Classical *X-ray* or *ultrasound* methods are used to acquire 2D images in different modalities. *Computed tomography (CT)* uses a series of X-ray images to reconstruct 3D volumetric image data. Those data can then be visualized as a 3D model, or as arbitrary projections or cross-sections of the volumetric data.

While not all carry the term “computed tomography” in their names, other medical imaging techniques use very similar methods to reconstruct 3D volumetric data. However, techniques like *magnetic resonance imaging (MRI)*, *single photon emission computed tomography (SPECT)* and *positron emission tomography (PET)* use differ-




Figure 1.2.: Outside view of a commercial PET scanner, the Siemens Biograph Sensation 16.

ent physical effects to acquire the raw data. While CT measures the attenuation of X-rays sent through a patient, PET measures gamma rays emitted from within the patient's body itself.

Actually the gamma rays measured by PET are emitted by radioactive isotopes injected to the patient. Such radionuclides can be built into many compounds appearing in the organism, e.g. glucose, ammonia or water. Substances *labeled* with a radionuclide are called *tracers*. PET can visualize the processes the tracers are involved in, and it even allows for quantification of those processes. This is in contrast to CT, which is specialized on delivering structural information about the patient's anatomy.

Visualization of functional data, also referred to as functional imaging, is essential in many branches of medicine. Oncologists use PET to detect and stage tumors, often using fluorodeoxyglucose with a radioactive fluor isotope as tracer. This makes up the largest part of PET scans performed nowadays. Cardiologists are interested in myocardial perfusion, which can be measured with ammonia tracers. Neurologists use radioactive oxygen isotopes to measure blood-flow in the brain, or special tracers in the context of neurotransmitters.


1.1.1. Physical Basics of PET

There are plenty of books and papers covering both technical and clinical aspects of PET. (E.g. [8], [22], [11].) Nevertheless, some issues of the physics of PET are detailed in the following, due to their implications on respiratory gating.

The fundamental principle underlying PET is a radioactive decay, more precisely the beta plus decay emitting a positron. (Hence the first part of the term PET: **P**ositron



1. Introduction

Emission.) The decay process is summarized by the equation:



Here, the p denotes a proton, whereas the n denotes a neutron. The e^+ is the positron, and is sometimes written as β^+ for historical reasons. (When beta radiation was discovered, it was not known that it was related to the electron. It was just given a sequential letter from the Greek alphabet.) Finally there is the neutrino ν_e , a very lightweight particle hardly interacting with other matter, and therefore of little interest in this context.

Note that β^+ decay does not occur in isolation, since the results of the decay have more combined energy than the original proton. However, β^+ decay can occur inside some radioactive nuclei. If the nucleus is more stable after the decay than it was before, the difference of the binding energies can power the decay process. Assuming that X is a nucleus with mass number a and atomic number z , then β^+ decay can be written:



The second line illustrates this on the example of ${}^{18}\text{F}$, a radionuclide commonly used in PET. Both of those equations are essentially the same as the previous one, only that the original proton and the resulting neutron are part of a nucleus. The proton stays within the nucleus while being transformed to a neutron, so the overall number of nucleons remains constant. On the other hand the atomic number (and name) of the atom, its binding energy, its charge and its chemical properties change, since there is one proton less.

There are several isotopes inclined to β^+ decay, but the ones whose stable counterparts are found numerously in the living organism are most interesting for PET. However, those radionuclides first have to be produced in a rather elaborate process, usually involving a cyclotron. This particle accelerator is used to collide protons, deuterium-cores, or alpha-particles with stable isotopes to create the radio-active nuclides. The bottom of figure 1.3b shows such reactions. (In fact all of the radionuclides presented here can also be created from other isotopes using just protons, so the used cyclotron can be tailored to proton beams.) Figure 1.3a shows the most common radionuclides used for PET in the context of the table of isotopes. This also shows some isotopes they can be created from, and their decay products. Figure 1.3b lists some important properties of the radionuclides. The shorter the half-life of an isotope, the more rapidly the activity decreases. So the radionuclides need to be applied very quickly, which implies that they must be created nearby the PET scanner to reduce transportation time.

However, the chemical process of labeling a tracer with radionuclides can also be time-consuming. For instance, when ${}^{18}\text{F}$ is used in place of hydrogen in fluorodeoxyglucose (${}^{18}\text{F}$ -FDG). This is possible due to the similar chemical properties of hydrogen

and fluorine, but due to the special characteristics of FDG it can be used to visualize certain metabolic processes. Then there is carbon (^{11}C) which can be used to visualize a large variety of functions, since it is quite common in organic compounds. Furthermore, oxygen with ^{15}O and radio-active ammonia ($^{13}\text{NH}_3$) are comparatively simple chemicals which are mainly used to measure the perfusion of the patient's organs.




Figure 1.3.: Radionuclides commonly used in PET: placement in the table of isotopes (a) and physical properties and creation processes (b) (Sources:[19], [11]).

Even though the radionuclides presented above are quite different, the results of the β^+ decay are similar. While the neutron stays in the nucleus, both the positron and the neutrino are emitted from it. The positron annihilates as soon as it meets an electron, as it is its antiparticle:

$$e^+ + e^- \rightarrow \gamma + \gamma.$$

This results in two photons having some interesting properties for medical imaging. Firstly, the complete energy of the two original particles, including their mass, is distributed evenly between the photons. This means each photon has an energy equivalent to the rest mass of an electron, roughly 511 keV. The kinetic energy of the particles before annihilation averagely adds to this only a few eV. Such photons are classified as gamma rays – hence they are denoted as γ – and have the advantage of passing through organic tissue relatively undisturbed. An even more important property of the resulting γ rays is, that they travel in almost exactly opposite directions. Hence, if both of them are detected, it is clear that they originated somewhere on a straight line between the detection points.

A PET scanner seeks to measure the distribution of a radionuclide in the patient's organism, and does this by detecting the γ rays. Unfortunately, there are some physical effects limiting the spatial resolution of PET. Firstly, the paths of the emitted γ rays do not form an angle of exactly 180° , but an angular deviation of averagely $\phi = 0.5^\circ$

does occur (see [11]). When detected by two sensors roughly $d = 1$ m apart, the worst case error is about 2.2 mm (calculated as $\frac{1}{2}d \cdot \tan \frac{\phi}{2}$ which is linear in d). Secondly, there is a small discrepancy between the location of the decay and the location of the annihilation. Though there are plenty of electrons around in normal matter, the emitted positron does move a short distance before it finds an electron with similar thermal energy to annihilate with. Until then, it is statistically scattered and slowed down by coulomb interactions. An average range of the positron can be determined, but it depends on the original kinetic energy of the positron and on the surrounding material. The average kinetic energy of an emitted positron is specific to the radionuclide involved. Figure 1.3b shows average values for the ranges of positrons emitted from different radionuclides. Water was assumed as surrounding medium here, so the given distances should correspond very well with those in organic tissues. The positron range error sums up with the angular deviation error (both in the region of 2 mm), hence the resolution of clinical PET is physically limited to about 4 mm. Note that those statistic effects only affect resolution, while precision of location measurements may well be higher, if enough γ rays are detected.

1.1.2. PET Scanner Technology

Despite the deviation effects, a distribution of decayed radionuclides can be measured by interaction of the emitted γ rays with an appropriate detector. When a 2D slice through a patient's body is to be acquired in clinical PET, the patient needs to be surrounded by γ ray detectors in the plane of the slice. In most commercial scanners this is achieved by a circular configuration of small detector crystals surrounding the patient², as sketched in figure 1.4a. When two γ rays are detected at two different points of the detector ring at the same time, it is assumed that they originate from the decay of a single radionuclide. Such a double-detection is called *coincidence* in the context of PET, while the straight line between the detection points is called *line of response*. Since the emitted γ rays move along the same straight line (but in opposite directions), the decay process causing a coincidence must have taken place somewhere on the line of response. The distribution of radionuclides in the detector plane can be reconstructed from the geometry of all lines of response. (Reconstruction is explained in section 1.1.3.) Due to the physical deviation effects mentioned earlier, the actual decay process did not take place exactly on the line of response, but merely close to it. The angular deviation grows proportionally with the diameter d of the detector ring, so there is a trade-off between scanner size and image resolution. While PET scanners with huge patient gantries are desirable for whole body scans and flexible patient positioning, smaller dedicated head scanners can achieve higher resolutions, let alone special scanners for small animal research.

In order to obtain volumetric data of the tracer distribution in a patient, two approaches are conceivable. Firstly, the patient bed can be moved along the scanner axis

²Other scanner geometries, e.g. a polygonal setup of medium sized detector plates or two large rotating detector plates opposing each other exist, but are not common for dedicated PET scanners.

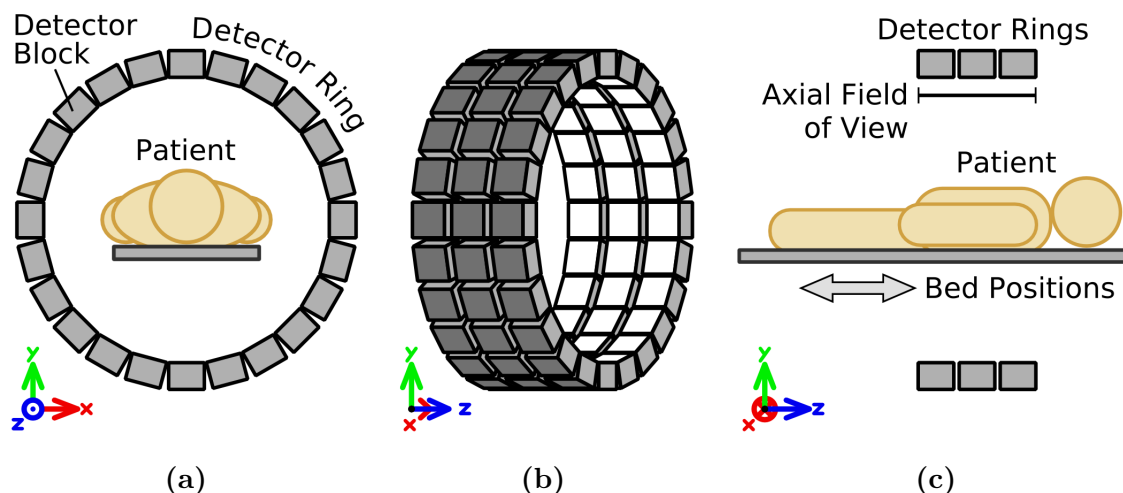


Figure 1.4.: Geometry of a common dedicated PET scanner for medical diagnostics: front view (a), diagonal projection (b), and lateral view (c).

in small steps, while separate 2D slices are acquired at each position. Secondly, the detector ring can be extended by appending additional rings along its axis to form a detector cylinder. (See figure 1.4b.) Neither of these solutions is practical for clinical scanners, since the first one implies enormous acquisition times, while the second one is extremely expensive - the extreme of completely surrounding the patient by detector crystals might be inconvenient anyway. Therefore most commercially available scanners implement a combination of the two approaches. They use a cylindrical configuration of detector rings to acquire a section of the patient volumetrically. If the region of interest is larger than the axial extent of the detectors, the image data have to be acquired in several phases and the patient bed has to be moved accordingly in between. (See figure 1.4c.) Current PET scanners are able to scan an axial region of about 15 – 20 cm per bed position.

Though most available PET scanners use multiple detector rings, they don't all work the same way. Some scanners use each of the rings separately, acquiring one 2D slice through the patient's body per ring. Only coincidences having both detection points on the same detector ring are used for image reconstruction, while lines of response between different rings are discarded. To avoid such inter-ring-coincidences, γ ray absorbing septa are placed between the detector rings. Those septa reduce the detection of γ rays not originating from the plane of the respective detector ring. Scanners using detector rings separately are said to operate in *2D mode*. In contrast, PET scanners operating in *3D mode* accept detections between any two detector rings, and use them for image reconstruction. 1.7. While 3D mode scanners may seem superior at a first glance, since they don't suppress as many γ rays as 2D mode scanners,

1. Introduction

the latter have advantages, too. They face less problems related to detector crystal dead-times and degrading scatter of γ rays. This is illustrated in figure 1.7 further below. In reality, the two concepts are sometimes mixed, but pure 3D mode PET scanners are becoming increasingly mature and popular.

No matter what the exact geometry of the detector ring is, the single detectors are not trivial themselves. This is due to the comparatively high energy of the γ rays used in PET (energy: 511 keV, frequency: $5.4 \cdot 10^{21}$ Hz, wave-length: $2.4 \cdot 10^{-12}$ m) compared e.g. to medical X-rays (energy: 80 keV, frequency: $2 \cdot 10^{19}$ Hz, wavelength: $1.5 \cdot 10^{-11}$ m). Photons of such energies are usually detected with scintillators, which are according to [25] “a device or substance that absorbs high energy (ionizing) electromagnetic radiation then, in response, fluoresces photons at a characteristic [...] longer wavelength”. The scintillators commonly used in PET are crystals, whose electrons interact with γ rays by *Compton scattering*. This interaction changes the heading of the photon by a random angle θ , while increasing the wave-length of the photon by $\Delta\lambda$, where

$$\Delta\lambda = \frac{h}{m_e c} (1 - \cos \theta).$$

Here, h is Planck’s constant, m_e is the mass of the electron and c is the speed of light. While $\Delta\lambda$ is independent of the original energy of the photon, the energy transferred from the photon to the electron is not. For a given energy E_1 of the photon before Compton scatter it is:

$$\Delta E = E_1 - E_2 = E_1 - \frac{E_1}{1 + E_1 \frac{1 - \cos \theta}{m_e c^2}}.$$

Since the γ rays used in the context of PET have energies equivalent to the mass of an electron ($E_1 \approx m_e c^2$), this can be simplified to:

$$\Delta E \approx E_1 - \frac{E_1}{2 - \cos \theta}.$$

Hence the resulting energies are up to 2/3 of the energy of the original γ ray, that is the electron receives 0 – 340 keV. This is often enough to strip the electron of its atom, i.e. to ionize it – hence γ rays are classified as ionizing radiation. In fact a 511 keV γ ray has enough energy for multiple ionization, and can lift several electrons from the valence band to the conduction band. When returning to a lower potential, such stimulated electrons emit lower energy photons, a process commonly known as fluorescence. Those photons, usually close to the spectrum of visible light, can then be detected by photomultiplier tubes [9], which convert them to an electronic signal.

For scintillator crystals used in PET several physical properties are important. Firstly, the *attenuation length*³ for γ rays should be as short as possible, which means

³The attenuation length is the distance into the scintillator when the probability has dropped to $1/e$ that a particle has not been absorbed, where e is Euler’s number.

an increased probability of Compton scatter. This is often the case for crystals with a high density, involving crystals with a high atomic number (having many electrons surrounding them). Lower attenuation lengths increase the sensitivity of a detector while decreasing the necessary thickness. Secondly, the light output should be rather high. It is defined as the number of emitted photons per absorbed γ ray energy. This is affected by several factors, e.g. by ionized electrons returning to a lower potential by non-fluorescent effects, or by the translucency of the scintillator crystal for the photons it emits. By choosing light output properties carefully, the energy resolution of the detector can be adjusted. Thirdly, there is detector dead-time. This is the period of time after a γ ray detection, during which no other γ ray can be detected. It is influenced by the duration of the fluorescent effect, i.e. by the decay time of the ionized electrons. Dead time is important, because it limits the number of γ ray coincidences that can be detected per time interval, and therefore needs to be minimized.

Name	Composition	Attenuation length for 511keV γ rays	Photon output per absorbed Energy	Fluorescence decay time
BGO	$\text{Bi}_4\text{Ge}_3\text{O}_{12}$	10.4mm	9/keV	300ns
LSO	$\text{Lu}_2\text{SiO}_5:\text{Ce}$	11.4mm	30/keV	40ns
GSO	$\text{Gd}_2\text{SiO}_5:\text{Ce}$	14.1mm	8/keV	60ns
NaI	NaI:Tl	29.1mm	41/keV	230ns

Table 1.1.: Properties of some scintillator materials commonly used for PET (Source: [11]).

Though the above properties are the most important for PET scintillators, other criteria are crucial, too. For instance low weight, sufficient stability, low sensitivity to light and other environmental influences, high durability, good compatibility with photomultiplier tubes, and reasonable costs. Some scintillator materials fulfilling these requirements, and therefore commonly used in commercial PET scanners are shown in table 1.1. In the past, the most common PET scintillator was BGO, which has a rather short attenuation length and is easy to handle. Today LSO is becoming more and more popular, due to its low decay time and high light output. However, LSO has the peculiarity of being slightly radioactive and self-emitting γ rays. This effect adds some noise to the measurements.

As a summary, figure 1.5 shows a sequence of some important physical processes occurring in the context of a PET examination. This involves the necessary steps for creating radionuclides and radioactive tracers. Then there are the processes going on in the patient's organism, and the detection of the resulting γ rays by the scanner. Naturally the PET scanner has to perform rather sophisticated signal processing steps to record all γ ray coincidences properly. Those coincidence data can then be used to reconstruct volumetric image data of the tracer distribution. Figure 1.5 gives a preview of the steps necessary to calculate the image, which are detailed next.

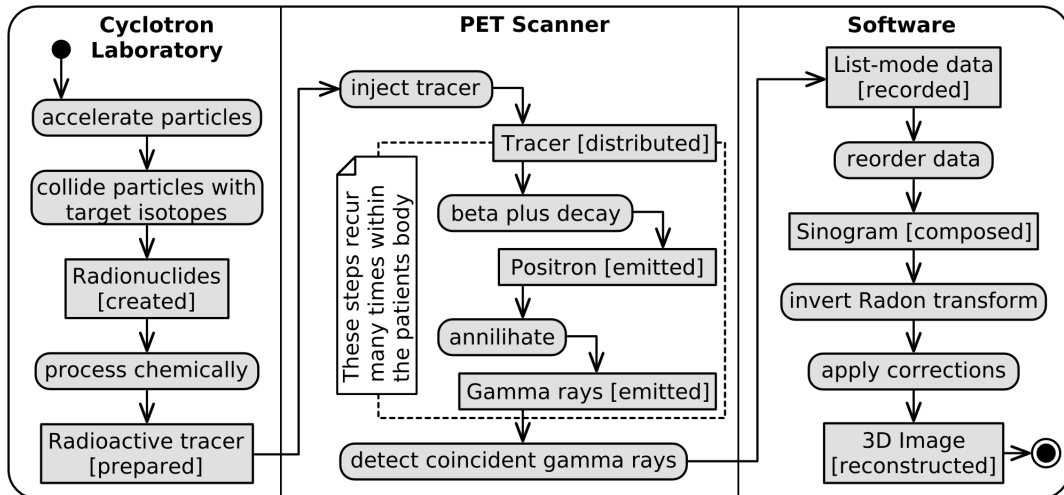


Figure 1.5.: The flow of physical and logical processes essential to PET.

1.1.3. Image Reconstruction

Modern PET scanners perform first processing steps on the detected coincidences on-the-fly, and do not store the exact detection times. However, the raw data resulting from a PET scan can be thought of as a timed list of coincidences – often referred to as *list-mode data*. Together with knowledge of the location of each γ ray detector used in the scanner, a volumetric image of tracer distribution can be computed from the list-mode data. The process is called image reconstruction, and is summarized in this section. Here, only the 2D case of image reconstruction is covered, as it appears in single-plane and 2D mode PET scanners. Image reconstruction for 3D mode is slightly more complicated, but often it is simply divided into several 2D reconstruction steps in a process called rebinning. For more details on image reconstruction issues, see the corresponding literature, e.g. [12] or [21].

Each detected coincidence indicates tracer activity on the line of response (LOR) between the respective pair of detector crystals. In a detector ring with n crystals, there are up to $n^2/2$ lines of response to be considered.⁴ It makes sense to group together all LORs parallel to each other. The LORs within one group can then be sorted by their location, while the groups can be sorted by their orientation. This is depicted in figure 1.6a, where s denotes the location of an LOR (given as distance from the scanner's axis) and α its orientation (given as rotation angle). The tracer activities associated to all LORs can be interpreted as a function of s and α , and visualized as a gray-scale image. An example of such an image, commonly called a

⁴That includes LORs between very proximate or even identical detectors. Such LORs can usually be discarded, since they do not intersect the region of interest. If d is the minimum number of detectors spanned by an LOR, their overall number is actually $n(n - 2d - 1)/2$.

sinogram, is shown in figure 1.6b. Note that α only ranges from 0 to 180 degrees, since the LORs are not directed. Also note that sinograms are common to other tomographic imaging modalities, too. For instance in CT, where the lines are not parallel, but the s corresponds to the camera's field of view and the α to it's rotation. This similarity of tomographic imaging techniques has had a positive effect on available reconstruction software.




Figure 1.6.: Lines of response classified by location and orientation (a), and the corresponding coincidence data drawn as a sinogram (b).

Since the representation of PET data as a sinogram is not too vivid, the next step is to transform it to a 2D image in euclidean space. In the continuous case, let f be a function of tracer distribution along the x and y -axes, then the sinogram R can be obtained as:

$$R(s, \alpha) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \delta(s - x \cos \alpha - y \sin \alpha) dx dy.$$

This integral transform from f to R is called Radon transform [24], and is based on Dirac's delta [23]. The delta distribution $\delta(z)$ has the peculiar properties that it is 0 for any $z \neq 0$ but $\int_{-\infty}^{\infty} \delta(z) dz = 1$. This can be visualized as a function having an infinitely sharp peak around zero. Demonstratively, in the context of computed tomography, the Dirac delta samples the measured distribution f along each LOR. This analytical representation of the tomographic projection can be used for calculations, most importantly by obtaining an inverse of the Radon transform:

$$f(x, y) = \frac{1}{2\pi^2} \int_0^\pi \int_{-\infty}^{\infty} \frac{\partial}{\partial s} R(s, \alpha) \frac{ds d\alpha}{x \cos \alpha + y \sin \alpha - s}.$$

It can also be shown, that the Radon transform is invertible for all cases relevant in computed tomography. Hence the above inversion of the Radon transform could be used to convert from a sinogram to a 2D image in the continuous case. However, the actual PET sinogram is a discrete data set, and there are more suitable approaches to reconstruct images from it. Most popularly, an algorithm named *filtered back-projection* (see [21]) is used to perform the image reconstruction quite efficiently. Many commercially available PET scanners therefore use filtered back-projection to reconstruct 2D images. Recently, algorithms based on statistical iterative reconstruction have become popular, too. The basic idea here is, to start with a rather poor estimation of the desired image, and iteratively improve it. In each iteration the image is forward-transformed to a sinogram, and compared with the measured PET data. The differences are then repeatedly used to improve the estimate, until the desired image quality has been reached. An advantage of statistical iterative reconstruction methods is that the forward-transformation can account for some known physical and technical deficiencies of the PET scanner.

The reconstruction of 3D volumetric data from a PET scanner, can be pictured as stacking several 2D images over each other. In that way, a complete data-set of the tracer distribution in the patient's body can be obtained. Arbitrary slices through the 3D data can then be computed easily by linear transformations. That allows physicians to visualize exactly the portion of the patient's body they're interested in.

To make things more complicated, there are several degrading physical effects, that need to be corrected during image reconstruction. First of all, the single detectors' sensitivities can vary, and need to be calibrated from time to time. The resulting adjustment factors for each LOR need to be considered when the sinogram is created. Second, despite the high energy of the γ rays, they do interact with the surrounding matter. They are both scattered and attenuated within the patient's body to a certain degree, depending on the location of the emission and the surrounding tissue. However, attenuation can be measured as described in the next section, and the resulting attenuation map needs to be applied during image reconstruction. The measured activity for each LOR is increased according to the respective attenuation factor. Scatter can be detected to a certain degree, too, and compensated by statistical filtering. Despite the large number of steps involved in image reconstruction, currently available computing systems can perform it so quickly, that images are available on the operators workstation right after a PET scan. Though many parameters of reconstruction are adjustable, most of the processing is done transparently for the user.

The image reconstruction process is also where respiratory gating can be incorporated. The list-mode data can be divided into distinct gates, according to data obtained from a respiration sensor. This is usually not supported by the integrated software packages and user-interfaces provided for commercial scanners. Therefore chapter 2.1 exemplifies a manual image reconstruction method, for the particular PET scanner system used for this research work.

1.1.4. Limitations of Clinical PET Scanners

The physical principles of PET are rather straightforward, but as indicated previously there are several degrading effects making PET instrumentation rather complicated. The resolution limit imposed by positron range and angular deviation, is impaired even more by the spatial resolution of the detector crystals. However, there are other problems related to the detected γ rays, too. Even though 511 keV radiation is quite permeant, there is quite a chance that it is Compton scattered or attenuated *before* reaching the detector ring. Since both emitted γ rays need to be detected in coincidence, there are different artifacts to be considered. Figure 1.7 shows some possible constellations in the context of 2D-mode and 3D-mode scanners.

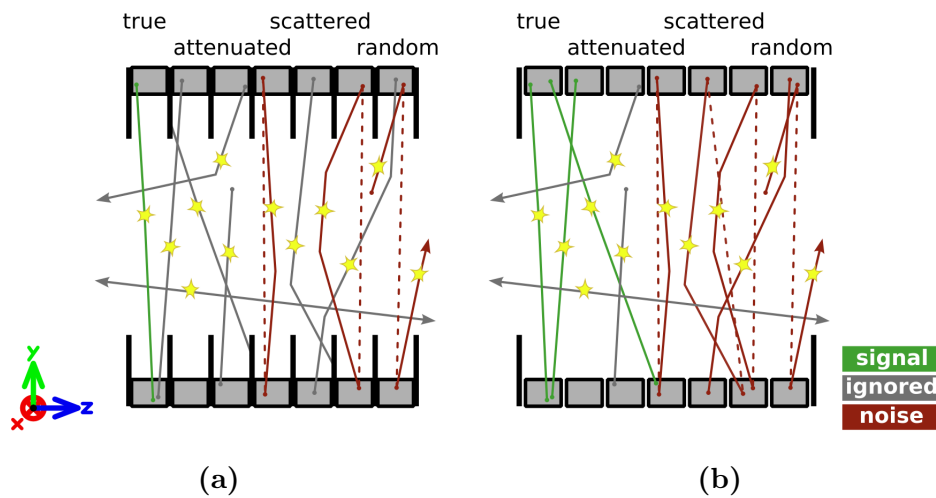


Figure 1.7.: Examples of measurement artifacts occurring in PET. While a large portion of scatter is blocked in 2D-mode (a) it is detected in 3D-mode (b).

If one of a pair of γ rays is attenuated completely or scattered past the detector rings, no coincidence will be registered. This neither improves nor corrupts the measurement. Otherwise, if one of the γ rays is scattered, but both reach the detector, a coincidence is registered, though the rays do not originate from the corresponding line of response. Such a measurement adds noise to the PET image, which is hard to compensate. If the γ ray loses enough energy when scattered, the scatter can be recognized by the detector, but often the energy resolution of the scintillator crystals is not sufficient. Note that a lot of scatter is avoided in 2D-mode, since it is blocked by the septa or the two γ rays hit different detector rings. Only in-plane scatter influences measurement in 2D mode, but a huge part of true coincidences is blocked, too. According to [22], 10% – 20% of the γ rays detected in 2D mode are scattered, while this fraction can be as high as 30% – 60% in 3D mode.

Both scatter and attenuation are highly dependent on the body region scanned and

on the size of the patient. In fact, γ rays originating from the center of the body are less likely to reach the PET detectors unhindered, compared to those emitted close to the surface. The probability that both emitted γ rays are not attenuated over a given distance d is $e^{-\mu d}$, where μ is called the attenuation coefficient.⁵ In the human body, μ varies rather heterogeneously: it is high for bones and low for other tissues, especially the lungs. Luckily, an effective attenuation coefficient μ_{LOR} can be determined for each line of response, resulting in a so called attenuation map. The amount of coincidences A_m measured along a LOR can then be adjusted according to the corresponding attenuation coefficient:

$$A = A_m \cdot e^{\mu_{LOR} d}.$$

The attenuation map can be obtained experimentally by sending a known amount of 511 keV γ rays through the patient along each line of response. The intensity of γ rays measured in the opposing detector can then be used to obtain the respective μ_{LOR} . Naturally, this may not be done while a significant amount of radioactive tracer is present in the patient's body. Many commercial PET scanners are equipped with a small and precise source of γ rays, which are rotated around the patient for measuring the attenuation map. This part of a PET scan is called *transmission scan*, while the actual measurement of tracer distribution in the patient's body is called *emission scan*.

The recent appearance of combined PET and CT scanners [20] allows for a slightly different approach. CT is a transmissive modality, which measures the attenuation of x-rays sent through the patient. Hence, a CT image (respectively its sinogram) practically is an attenuation map. Unfortunately the attenuation coefficient μ is different for 511 keV γ rays and for x-rays, and does not even scale uniformly for all materials. However, an attenuation map for PET can still be derived from a CT image by different approximation methods. In the most simple case, two different correction factors are applied for bone-like tissue on the one hand and for any other tissue on the other hand. In commercial scanners more sophisticated variations of this method are implemented, and most of them yield a decent attenuation map. Since current CT scanners are fast compared to PET scanners, using CT for attenuation correction is quite convenient. The combination of PET and CT in the same gantry has other advantages, too. It facilitates the creation of well-registered fused images containing information from both modalities.⁶

In fact, time consumption is one of the major limitations of PET. One reason for this is the temporal resolution of the detectors and the signal processing electronics. Though the temporal difference of two coincident γ rays arriving at the detector ring lies below a nano-second, the electronics need to establish a slightly larger time window

⁵The attenuation coefficient μ is just the inverse of the attenuation length mentioned in section 1.1.2 in the context of scintillators.

⁶Most recently, combined PET-MR scanners have been announced, too. Though magnetic resonance tomography is not based on electromagnetic radiation, the feasibility of MR images for PET attenuation correction is being investigated. [27]

to detect coincidence events. However, if emissions of two or more positron decays reach the detector crystals during the same time window, they cannot be classified correctly. The rightmost part of figure 1.7 shows such a random event, where two unrelated γ rays are detected at the same time, while the other two are scattered away. When γ rays are falsely detected as coincidence, and noise is added to the sinogram. In other random cases, three or more γ rays can arrive at the same time. Usually they all have to be discarded by the scanner, since the corresponding lines of response cannot be determined unambiguously. In addition to the problems between different detectors, each detector crystal has an intrinsic dead time. Due to this even longer time period, caused by slow fluorescence decay (look back at table 1.1 for data), a scintillator cannot detect two γ rays arriving shortly after each other. This leads to some more coincidences ignored by the scanner.

The higher the decay activity within the PET scanner, the more severe all the timing related artifacts. On the other hand, to obtain good image data, as many coincidences as possible need to be measured. Due to the noise added by scatter and random events, only a limited number of true coincidences contributes to the signal quality of the image. As a rule of thumb, about $5 \cdot 10^6$ true coincidences are needed to acquire a 10 cm cross-section of a human body in reasonable quality. The over-all amount of detected coincidences including randoms and scatter is almost one order of magnitude higher. The total amount of positron emissions in the scan subject is another order of magnitude higher, due to attenuation and especially γ rays missing the detector crystals. Considering that the count-rate for coincidences is best around 10^6 decays per second for current scanners, PET acquisition times of several minutes are inevitable. Note, that the duration of a PET scan cannot be simply decreased by increasing the activity of the radio-active tracer. If it gets too high, the result can even be worse, due to an over-proportional increase of random events and scintillator dead time.⁷

Due to the long time required for PET scans, several biological artifacts can degrade image quality. Whereas CT scans can be performed within a few seconds during breath-hold, a patient will necessarily breathe during a long PET acquisition. This considerably affects image quality of the thorax region, since the contained tracer activity is moving around with the tissue. In effect, the activity distribution measured by PET will be smeared over all phases of the recurring respiration cycle. Though this motion blur is only in the region of one centimeter, it can complicate the localization of tumors and the quantification of cardiac perfusion. Hence, it is advisable to correct respiratory motion as far as possible. A approaches for that will be presented and discussed in the rest of this report.

⁷There are also medical and regulatory limits to the amount of radio-activity that can be applied to a PET patient. As for today, those limits do not need to be exploited, since the PET detectors don't work efficiently with that much radiation in the first place.

1.2. Objectives of this IDP

To correct the negative effects of respiratory motion in medical imaging, the respiration has to be measured first. The primary goal of this IDP was the evaluation of different sensors suitable for respiratory motion estimation. The actual sensors used are described in more detail in the next chapter. The list comprises a stress measuring belt (Anzai, chapter 2.2), an air-flow sensor (PMM, chapter 2.3), a thermometer (Biovet, chapter 2.4), and a marker-based 3D camera system (ART, chapter 2.5). System components of three of them are depicted in figure 1.8. The PMM spirometer sensor was a prototype, we could only use for a short time, so no photograph can be presented here.




Figure 1.8.: Some of the evaluated Sensors, as presented by their manufacturers: the Anzai belt (a), the Biovet system (b), and the ART camera (c).

Chapter 2 focuses on the use of the sensors, rather than on internal details and sensor principles. An important part of this IDP consisted of getting the sensors to work, and access their data for further processing. The purpose of the measured respiration data was their use for respiratory gating. Hence it was not intended to derive quantitative physiological values like lung volume or air-flow from the sensors. They were merely used, to partition the progression of respiration into gates representing different poses of the thorax. It was not deemed important, which physiological mechanisms were related to thorax motion, or how they related to the measured sensor data. Hence, the data of all respiratory sensors could be treated as a function over time, called respiration curve or respiratory signal below.

While some of the respiratory sensors delivered such respiration curves right away, data from other sensors had to be post-processed. Since our PET research group had already experimented with the Anzai belt before this IDP started, it could be used as a reference system. Software to examine, edit and process the data was already available. (It was written in IDL, the *Interactive Data Language*, a matrix-oriented interpreter language intended for scientific computing.) As for the other sensors, the

software had to be enhanced to handle their data, too. This was most challenging for the ART camera system, which is not a dedicated respiration sensor. Rather sophisticated processing software had to be written here, to convert the multiple motion measurements delivered by the camera to a respiratory signal. A large work portion of this IDP consists of the design and implementation of this software, with all the special cases that had to be considered.

Once respiration curves from all the sensors were available, their use for respiratory gating was evaluated. First checks were performed by visual comparison of plotted respiration curves. Then systematic statistical evaluations of both the respiration curves and the gates derived from them were performed. The details about the comparisons and their results are found in chapter 3.1. That chapter also evaluates gated PET images reconstructed with data from some of the sensors.

The image evaluation focuses on gating with the Anzai belt, and the ART camera. Anzai is used because it was already tested beforehand, and we had more experience with this system than with any other. While it was estimated in advance, that the ART Camera system would be quite cumbersome to use, we focused on evaluating this system, too. This is due to the fact, that it delivers much more data, than any of the other systems. An optional goal of the IDP was to check the feasibility of more general movement correction based on the ART data. It should be investigated whether respiration artifacts could be corrected a more sophisticated way, and whether other non-periodic body motion could be corrected, too.

1.3. Respiratory Motion Correction in Medical Imaging

In pulmonology, the analysis of respiration by air-flow and lung-volume measurements is quite common. The respective devices, summarized as spirometers, are not commonly used for respiratory gating yet. Some spirometers are simply too bulky and offer an unnecessary range of sophisticated diagnostic functions. Yet there are some, which might well be suitable for respiratory gating in the context of PET. Indeed, an air-flow based spirometer was evaluated in this IDP. Nevertheless, some other respiration sensors have been suggested for respiratory gating, too. Experience can also be drawn from existing cardiac gating systems. Some related technologies already in clinical use are presented in the next section. Subsequently, some recent scientific developments directly related to respiratory gating in PET are highlighted.

1.3.1. Gating Systems in Clinical Use

Though respiratory gating is not a clinical standard in the context of PET yet, several related techniques are in use already. For instance, cardiac gating is applied successfully in PET and other medical imaging modalities. Cardiac gating is similar to respiratory gating, but the degrading effect which is corrected here is the motion of the heart. Cardiac gating relies on *ECG* as sensor input, which is already integrated with many commercial PET scanners. Though the heart-rate is normally higher than

the frequency of respiration, some aspects of cardiac motion are easier to handle. The motion is less continuous, and is performed in short pulses during which the heart contracts (*systole*). By contrast, the heart is expanded for a comparatively long time (*diastole*), during which PET data can be acquired undisturbed.

ECG based cardiac gating is applied likewise in other medical imaging techniques, e.g. MR. There it is sometimes combined with respiratory gating, which is already commercially available. An extrinsic measurement of respiration is not needed in this context, since it can be derived from the raw MR data itself. Such intrinsic respiratory gating is possible, since MR can acquire enough data during short periods of time to distinguish inspiration and expiration.⁸ Due to the low temporal resolution, such intrinsic methods are not easily portable to PET. However, some research on this is currently being done at the PET group, but that is not subject of this report.

Commercial respiratory gating systems are also available for external beam radiotherapy. This is a treatment for various types of tumors, which is based on the irradiation of malignant tissue with a high dose of γ rays. The radiation needs to be focused on the tumor as precisely as possible, which can be impaired by respiratory motion. Therefore, respiratory gating is sometimes used for tumors in the thorax region. The radiation beam is focused on the location of the tumor for instance during exhalation, and is deactivated for other phases of the respiratory cycle. Respiratory gating in radiotherapy is not common yet, but a few related systems are commercially available. In fact, the Anzai belt (see chapter 2.2) is primarily intended for respiratory gating in radiotherapy.

1.3.2. Current Research

Radiotherapy was also the background of one of the first research projects evaluating different sensors for respiratory gating systematically. In [13] H. Kubo and B. Hill were comparing a temperature sensor, a strain gauge (similar to the Anzai belt used here) and a spirometer. Checking, if their promising results could be repeated with current sensors was one of the objectives of this IDP.

While the above work did not include a 3D camera system, this has already been evaluated variously. Nehmeh et al. used a camera-based gating system in the context of PET imaging [16]. Their focus was on the improved quantification of FDG uptake in lung lesions. In [4] and [3] Beach et al. explored the use of such a camera system for respiratory gating and motion detection for SPECT imaging. A related research group explored the feasibility of separate signals for respiration and motion correction. In [6] they suggest a sophisticated algorithm which splits a respiration curve into large-scale body motion and respiration. Body motion could then be corrected by applying rigid transformations to the raw data, while gating takes care of respiratory motion.

⁸Actually respiration is not measured directly from the MR image, which is to be gated. Instead, additional MR images are acquired in parallel by time slicing. This results in a so called navigator, a sequence of low resolution images of e.g. the diaphragm, from which respiration can be measured.

There are also efforts to supplement respiratory gating, by applying additional transformations to the gated images. While ordinary gating results in gated images only containing a subset of available data, it has been suggested to recombine these images for enhanced image quality. In [14] Livieratos et al. use linear transformations based on radio-active markers placed on the patient's thorax to register the gated images. Though precise measurements with a 3D camera could be used in a similar approach, no such effort was made in scope of this IDP. Instead, image based registration of gated PET image sets is currently being evaluated in the PET research group at Klinikum Rechts der Isar. Though this is not detailed in this report, it should be mentioned, that there were very promising trials with a new deformable registration algorithm [28].

1. Introduction

2. Methods for Testing Respiration Sensors

Before the applicability of a sensor for respiratory gating can be considered, it has to fulfill several criteria. The four respiration sensors used in this study were the Anzai belt (section 2.2), the PMM spirometer (section 2.3), the BioVet temperature sensor (section 2.4), and the ART camera system (section 2.5). Being quite different, all the sensors were first checked against the following quality requirements:

Usability A respiration sensor needs to be easy to operate. The additional time needed to set it up for a PET scan should be short. Also, the necessary devices should not take up too much place in the examination room, or interfere with other medical equipment.

Invasiveness A respiration sensor may not be too invasive. Due to the long duration of PET scans it is important that patient comfort is assured. Otherwise, increased patient motion might occur. Since there are enough alternatives, sensors potentially harmful to the patient can be ruled out in the first place.

Reliability A respiration sensor must be highly reliable. On the one hand this means, that the respective sensor principle can track respiration under varying circumstances. Neither differing patients nor external influences should prevent it to work. The acquisition of suitable respiratory signals should be reproducible in any situation. On the other hand it means that sensor hardware and the related software work reliably. Unpredictable system failures or limited availability cannot be tolerated in the context of a PET scan.

Costs A respiration sensor should not be excessively expensive. However, considering the costs of PET and of medical equipment in general, the costs of respiration sensors are secondary. In this report, little can be said about the costs, since they are highly negotiable.

In the end, the most important requirement for the respiration sensors is to deliver data suitable for respiratory gating of PET images. For this work a gradual approach of evaluation was followed. First, each sensor was tested for its general suitability to deliver respiration curves. This was done by the author and his supervisors for each sensor separately. These first tests included some short sample acquisitions with one of us playing test subjects. Those early tests were meant to show, if the respective

sensors could principally fulfill the requirements above, and if the resulting respiration curves might be used for respiratory gating.

Subsequently some longer data acquisitions were performed, with several persons (both male and female) of our PET research group volunteering. These series of tests focused on the comparison of different respiration sensors. Unfortunately it was not possible to acquire the same respiration sequence with all of the sensors under consideration simultaneously. This was due to organizational restrictions, but it would have been difficult to combine the BioVet temperature sensor and the PMM spirometer anyway. Instead, the Anzai belt was used as a reference system, and all of the other sensors were compared to it. This was viable due to its dedication to and approval for respiratory gating, and due to previous experiences of the PET research group. The Anzai belt was technically suited for combination with all the other sensors, too.

Most of the comparative measurements were performed over two or three minutes of regular breathing, plus some additional time with simulated breathing artifacts. This time span seems short compared to the 11 minutes we used for gated PET acquisitions (see below). However, it was sufficient to span about a dozen respiration cycles for direct comparison. This approach was favored, in order to acquire more comparison sequences with different volunteers. While it was unrealistic to simulate the exact circumstances of a lengthy PET scan in this early studies, it was desirable to challenge the sensors with as many different patients as possible.

While it was expected that all sensors show regular respiratory motion decently, focus was laid on their handling of respiration artifacts. For this, short phases of speaking, coughing, heavy breathing, or breath-hold were performed by the volunteers. It was then analyzed, how those phases were reflected in the respective respiration curves. Since short breath-holds at full inspiration were recorded nicely by all sensors, they were used as landmarks for signal synchronization. This was necessary since we did not bother to synchronize all of the sensors electronically. Instead the volunteers were asked to perform a breath-hold at the beginning and at the end of each acquisition, and the respiratory signals were synchronized manually according to those landmarks.

This was done by different software tools written in IDL [17] and offering graphical user interfaces for respiratory signal comparison. Though some of the sensors brought their own software to analyze the respiration curves, a separate program was needed to compare data from all the different sensors. The manufacturer-provided programs were only used for data recording and for live-monitoring where applicable. All other signal evaluation was performed with custom software only working on recorded data, and not able to display data from any of the sensors on the fly. Conveniently it was able to re-sample the different signals automatically, thus compensating different sampling rates. One variant of the program (`respcurves.pro`) was developed by A. Martinez-Möller and R. Bundschuh before this IDP started. Some corrections and feature-additions were then performed by the author. Another variant of the IDL program, supporting a module for handling ART camera data (as described in sections 2.5.4 and 2.5.5), was developed as part of this IDP. This `Signal_Analysis_GUI.pro` program also features additional support for comparing respiratory signals and the

resulting gates quantitatively, as detailed in chapter 3.1.2.

After sufficient experience was gained with the different sensors, the next step was actually using them for reconstruction of gated PET images. Here the use of volunteers seemed inadequate, due to radiation exposure and the costs of a PET acquisition. Instead we cooperated with a study conducted on cardiac patients and measuring their myocardial perfusion. This study was well suited for evaluation of respiratory gating, since only PET images of the thorax region were acquired. Before the diverse respiratory sensors are presented, some details of this study, the PET scanner used, and respiratory gating are described in the next section.

2.1. Use of Respiration Sensors in the Context of PET

The patient study was conducted on a combined PET and CT scanner, the Biograph Sensation 16 PET/CT available from Siemens Medical Systems (Erlangen, Germany, <http://www.medical.siemens.com>). It combines an ECAT Accel PET scanner and a Somatom Emotion 16 CT scanner in one gantry, and offers integrated data processing for the two. As for the PET part, it uses LSO scintillator crystals arranged in 8×8 detector blocks. Each crystal has a surface of $6.45 \times 6.45 \text{ mm}^2$ facing towards the scanner axis, and a radial depth of 2.5 mm. There are 3 rings of detector blocks, with 48 blocks per ring, totaling in 9216 detector crystals. The detector ring diameter of 82.4 cm results in an effective radial field of view of 58.5 cm, while the axial field of view is 16.2 cm per bed position. The manufacturer designates the spatial resolution as about 6 mm in the center of the scanner, slightly decreasing towards the outside. The PET scanner always operates in 3D mode, which is possible due to the rather short coincidence timing resolution of 3 ns.

The cardiac study protocol involved ^{13}N -labeled ammonia, a tracer used to monitor blood perfusion. PET acquisitions were performed two times per patient, first under rest conditions, and then forty minutes later under stress. One minute before latter, stress was induced pharmacologically by a six minutes infusion of *Adenosine* at 0.6 mg/min per kilogram of patient weight. For each scan, the patient received 300 to 500 MBq of tracer activity, one minute after the start of the respective PET acquisition. The first 11 minutes, of each scan were performed as list-mode acquisition, making raw coincidence data available for manual processing. This could be used to reconstruct gated images later. Each scan continued for another 10 minutes of standard PET acquisition, from which static PET images for medical diagnosis were reconstructed by the integrated PET software.

Besides the PET scans, several CT scans of the patient were acquired. In the beginning of the procedure, a 2D topogram of the patient was obtained and the region of interest for successive scans was framed. Then a low-dose CT scan was acquired, to be used for the attenuation map correcting measured tracer intensities in the PET scans. This CT scan was acquired during patient breath-hold at medium inspiration, and was usually suitable for attenuation correction of both gated and non-gated images.

In rare cases the low-dose CT had to be repeated after the PET scans, if the first one turned out to be poorly aligned with the static PET images. After all those PET related scans, various forms of CT angiography (*CTA*) were performed. Those procedures shall not be detailed here, since they were not related to respiratory gating.

For the patient study, we focused on just two respiration sensors: the Anzai belt, and the ART camera. ART camera data were acquired over the complete duration of the PET examination, usually around $1\frac{1}{2}$ hours altogether. This allowed for the evaluation of body motion over the whole examination time. We expected some unpredictable patient motion, due to the long time the patients had to lie still and due to their rather poor health. In contrast, the Anzai belt was only used during list-mode PET scans. Actually the respiration curve was recorded for a slightly longer interval, and the electronic synchronization feature described below was used to trim it.

Synchronization of the ART camera data was performed manually. For that the approximate time between the start of ART acquisition and the list-mode PET scan was measured. The derived respiratory signal was then shifted accordingly, and aligned with the one from Anzai as thoroughly as possible. Significant respiration peaks and occasional breathing artifacts helped to verify that signal alignment was not off by one respiration cycle.

To obtain gated images from the list-mode data and the respiratory signals, several reconstruction steps had to be performed manually. These were performed by command-line programs from the `e7recon` suite, provided by Siemens Medical Systems. Figure 2.1 gives an overview of the various reconstruction activities necessary, and some intermediate products. The required programs, respectively the file formats used, are given in brackets. Besides the `List-Mode Data` and the `Respiratory Signal` the `Low-Dose CT` for attenuation correction, and a `Normalization File` for calibrating scintillator sensitivity are needed as input. (The `Normalization File` is generated regularly by scanning a phantom containing well defined decay activity.) Respiratory gating is introduced quite early in the reconstruction process. When the `List-Mode Data` are gathered to a `Sinogram`, certain intervals can be omitted according to a `Gates File`. The `Gates File` consists of an array of gate-numbers, one for each time-frame of `List-Mode Data`, and only data from gates with the desired number are adopted to the `Sinogram`. In our case, the `Gates File` first needs to be transformed to a `Listmode Header`, before the `Sinogram` is created. Afterwards, the gated `Sinogram` can be processed as usual. First some corrections are applied, using the `Normalization File` and the `CT Attenuation Map`. The resulting `Corrected Sinogram` is then reconstructed to a volumetric `PET Image`.

The resulting `PET Image` and many intermediate files are in ECAT format. This is a proprietary binary format for volumetric data in medical imaging, similar to the DICOM standard [2]. In contrast, the `Respiratory Signal` and the `Gates File` use a text-based format. The creation of the `Gates File` is a reconstruction step not covered by the `e7recon` suite. However, several gating algorithms defining the gates according to a given `Respiratory Signal` have been implemented at our PET group prior to this IDP. Axel Martinez-Möller wrote these algorithms in IDL (`RespGatingGUI.pro`),

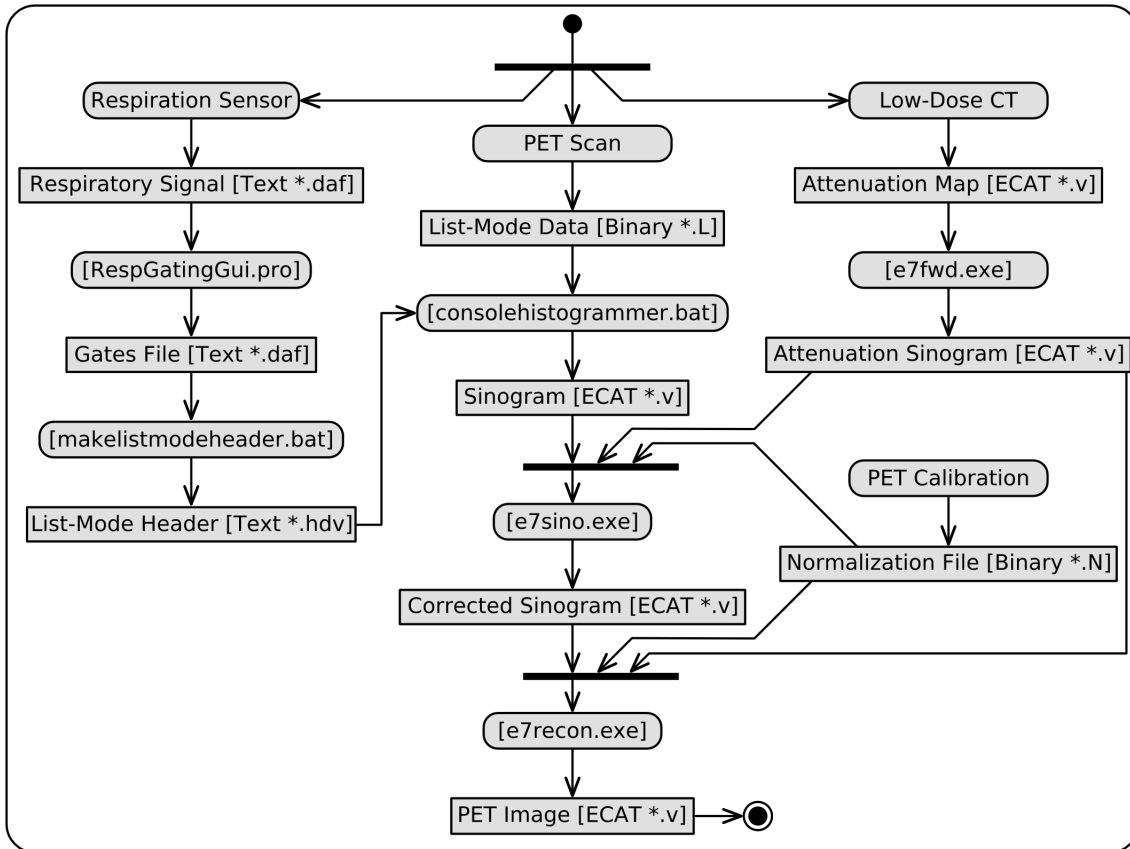


Figure 2.1.: The processing steps performed to reconstruct gated PET images. (A detailed view of the right part of figure 1.5.)

and was evaluating them in parallel to this IDP. While some of the algorithms are more sophisticated, he suggested to focus on a rather simple one for comparison of the respiration sensors. That algorithm splits the respiratory signal into single respiration cycles, by determining all moments of peak inspiration. Then it regularly divides the resulting intervals according to the desired number of gates. For most of the comparative study of respiration sensors, this rather robust algorithm was used with four gates.

To compensate for all image degrading motion in the thorax area, respiratory gating has to be combined with cardiac gating. The **Gates File** for such dual gated reconstructions is obtained by a combination of a respiratory and a cardiac **Gates File**. Such a file then contains gate numbers for each combination of a respiratory gate and a cardiac gate. However, for the evaluation of respiration sensors, no cardiac gating was applied. It seemed more reasonable to focus on the gating signals obtained from

the sensors presented below, without introducing additional complexity.

2.2. Anzai Belt

The *AZ-733V* by Anzai Medical (Tōkyo, Japan, <http://www.anzai-med.co.jp>) is a commercially available respiratory gating system. It is a dedicated respiration sensor for use in radiation therapy and medical imaging. According to the manufacturer, the *AZ-733V* is FDA approved and CE mark certified for those clinical uses. This makes the system rather mature and controllable, compared to some of the other respiration sensors tested. Therefore it was used as a reference system in this IDP.

There are two kinds of sensors available as part of the *AZ-733V* system, a laser and a belt. We used the belt version, hence this respiratory motion sensor is referred as *Anzai belt* throughout the text. The sensor part consists of an elastic belt, which can be fastened around the patient's abdomen or thorax by hook and loop tape, as shown in figure 2.2. During inspiration, the circumference of those regions increases and the belt is stretched. The increased stress results into increased pressure to the inside of the belt. This pressure is then measured by a small load-cell – an electronic force sensor – placed in a pocket inside the belt. Figure 1.8a in the introduction shows both the belt and the load-cell. There, the photograph of the load-cell is magnified, but in reality its diameter is about half the width of the belt.



Figure 2.2.: The Anzai belt strapped around a patient's body.

The complete *AZ-733V* setup consists of several interconnected devices. The load-cell placed in the belt is connected to a sensor port which in turn is connected to a wave deck. The wave deck can be connected to a standard PC by USB, and the whole system is delivered with a notebook computer running the corresponding Anzai software. The windows program `az733v.exe` offers a graphical user interface for calibrating, monitoring, processing and recording data from the Anzai belt. Manual calibration is necessary, since the signal measured by the load-cell has no quantitative meaning.

It is linked to thorax-expansion, but there are simply too many unknowns, e.g. the anatomy of the patient and how tight the belt is fastened. The calibration controls can magnify and shift the signal according to user needs. Besides the respiratory signal, always in a range from -25 to 125 , some additional curves are displayed. One curve optionally shows gates computed from the respiration curve on-the-fly. Another one shows an external signal, attached to the serial port of the wave deck.

To synchronize the recorded respiratory signal with the list-mode PET acquisition, a custom built cable was used. It connects two pins of a parallel port on the PET server¹ to two pins of the serial port of the wave deck. A custom extension of the list-mode acquisition software generates a signal on those two pins, for the duration of a list-mode acquisition. This signal is evaluated by the wave deck, and the Anzai software displays and records it as the external signal. Its value is 1 normally, but it switches to 0, while a signal is received on the serial port.

For patient studies the belt was placed on the upper abdomen, right where the costal arch begins. Experience showed, this is a reasonable setup for a large variety of patients. Of the three alternate belt lengths delivered with the AZ-733V, the medium one was suitable for almost all subjects encountered. The system is also delivered with two alternate load-cells, varying in sensitivity. We always used the one designated “low” (curiously being the more sensitive one). For PET scans, the sensor port remained on the patient bed, while the long cable allowed to put the wave board and the notebook in the control room. There the respiratory signal was calibrated to a range of about 25 (expiration) to 75 (inspiration) for normal breathing. That did not tap the full available range, but the signal did not have to be re-calibrated, if the patient’s breathing pattern changed a little. Altogether the setup procedure of the Anzai belt was quite easy to handle, and did not impair patient comfort at all. An additional benefit of the system is that all components are reusable, including the belt.

The recording of the Anzai respiration curve was started manually, shortly before the PET acquisition in list-mode started. After the list-mode PET scan finished, recording with Anzai was manually ended. The Anzai software stored the recorded signals to a file, using a custom text format called DAF to represent data. DAF files consist of a header holding various information like date, time, comments, and sampling rate – we always used 40 samples per second for this sensor. The subsequent data section of the file contains the signal values, one line for each sample. Various auxiliary signals are stored, too, containing values for e.g. detected peaks, computed gates and external signals. However, for this IDP only the pure respiratory signal was utilized. The external signal received from the PET server was also evaluated, too, but it was merely used to trim the respiratory signal to the duration of list-mode PET acquisition.

The respiratory signal stored in the DAF file is normalized to the same range as the respiration curves displayed in Anzai software. Data exceeding this range by more

¹More precisely the ACS, a machine delivered as part of the Siemens Biograph Sensation 16 PET/CT and being responsible for PET data processing.

than 25 is simply truncated by the software. That's why proper calibration is essential for further evaluation of the sensor data. Also, the respiratory signal is only given in integer precision. But this turned out to be more than sufficient with respect to the calibrated amplitude of about 50. Since the AZ-733V was used as a reference system in this study, the signals obtained from the other respiration sensors were also converted to DAF for further processing. For instance, the `RespGatingGUI.pro` tool accepts DAF files as input format. More importantly, the resulting gates files are in DAF format, too, since this format can be used directly as input for respiratory gating. (Compare to figure 2.1 above.)

2.3. PMM Spirometer

Another respiratory sensor evaluated in this IDP is the PMM spirometer by Siemens Medical Systems (Erlangen, Germany, <http://www.medical.siemens.com>). This device is dedicated for respiratory gating in medical imaging, but unfortunately it was at disposal for only a few days. Also, it was in an early prototype stage, which did not fulfill some basic quality criteria yet. Therefore, reliability and usability of the PMM are not considered here. We only evaluated if the sensor principle works, and if such a sensor might be used for respiratory gating in the future.

The PMM is a spirometer, that is a sensor measuring air flow to and from the lungs. The patient gets an air-tube placed to the nose, as seen in figure 2.3. Part of the respired air flows to the sensor attached to the other end of the tube, and the data it acquires is sent to a PC for further processing. Rather than displaying the raw flow rate, the processing software uses it to calculate the progression of air-volume in the patient's lungs. This can be thought of as integrating over the flow rate signal. It makes the resulting respiration curve quite similar to the one obtained from the Anzai belt. The PMM performs automatic normalization of the signal, if it drifts out of a predefined range.

The early software version of the PMM prototype was able to write respiration curves to text files. The file format used is very simple: each line contains one data sample. It consists of a time-stamp, given in seconds, and a corresponding air-volume, given as integer number. Unlike other spirometers, the values given by PMM do not have any absolute meaning. However, their progression is more or less proportional to the progression of air-volume in the patient's lungs. This is similar to the Anzai belt, and turns out to be sufficient for respiratory gating.

For the short tests conducted with the PMM spirometer, the data were always recorded with the default 50 samples per second. An electronic synchronization with other devices was not performed. We never used the system in the context of an actual PET scan. However, several short respiration sequences were acquired with volunteers. In those acquisitions the PMM spirometer was used together with the Anzai belt, and they were started manually at the same time. To correct for small inaccuracies in synchronization, the obtained respiration curves were later shifted manually. They



Figure 2.3.: The PMM spirometer is plugged to an air-tube placed in the patient's nose.

could be synchronized quite precisely according to intentional breath-holds performed by the volunteers at the beginning or end of each sequence.

Due to the prototype status of the PMM spirometer and the limited availability, our experiences should not be overestimated. Yet it can be said, that the system basically works and its handling seems pretty straightforward. However, the reliability of the prototype was unacceptable, and some of the respiration curves contained serious artifacts, as detailed in chapter 3.1.1. Much of this seemed to be related to software issues, rather than to physical limitations of the sensor. It is likely that those problems will be solved when the final version of the PMM spirometer is released. Nevertheless, it has some disadvantage compared to Anzai: problems may occur for patients needing oxygen supply, and for patients breathing through the mouth. Also, some may find a tube stuck in their nose less agreeable than a belt.

2.4. BioVet Temperature Sensor

Another type of sensor evaluated in this IDP, is a special thermometer placed to the patients nose. It measures the changing temperature of air respired by the patient with a high temporal resolution. Assuming room temperature, the outside air passing the sensor on inspiration is colder than the air from the lungs passing it on expiration. For this study, we used the temperature sensor included in the *BioVet CT1 System* by Spin Systems (Brisbane, Australia, <http://www.spinsystems.com.au>). The BioVet system digitally processes the input of several physiologic sensors, and is intended for use in medical imaging research on small animals. The scope of applications includes both cardiac (ECG triggered) and respiratory gating. Though it is not advisable, to use this system clinically on humans, experiences with the BioVet system should be

transferable to many other temperature sensors.

Note, that there is another respiration sensor provided with the BioVet system. It is a small elastic pad, approximately the size of a coin, which can be taped to the animal's abdomen and measures pressure pneumatically. This is similar to the functional principle of the Anzai belt, and works quite well for mice. Tests of attaching this respiration sensor to a human thorax or abdomen quickly showed, that the sensor is not suited for clinical respiration measurement. While it might have worked, it turned out soon that noise was very high, calibration for humans rather poor, and reliability extremely low compared to the Anzai belt. We figured, that one sensor based on body expansion was enough for the comparative study, and stuck with the Anzai belt. However, pressure data were always acquired alongside temperature data, when using BioVet.

The setup of the BioVet system is quite similar to the one of the Anzai belt. The temperature sensor and other accessories are attached to the BioVet Acquisition Module (BVAM, see figure 1.8), a small box handling the physical sensor inputs. This is connected via a sufficiently long cable to the BioVet Control Module (BVCM), a slightly larger box performing signal processing and offering auxiliary inputs and outputs. The BVCM can be controlled by PC software to which it is attached by USB. The BioVet CT1 software delivered by the manufacturer offers a wide range of control and processing functions. However, in this IDP it was only used to write respiratory signals acquired by the sensors to files.

Once again the output is written to plain text files, using a comma-separated values format. Each line contains one data sample for each of the supported sensors and some user signals. The temperature values are given as integers in units of 0.1 °C. For the respiratory signals acquired in this IDP, the default sampling rate of 1000 Hz was used, even though this seems exaggerated for respiratory gating. As for the PMM, no synchronization with other devices was implemented. Instead the recorded respiration curves were synchronized manually, using a landmark based approach.

Since this sensor was lent to us by another group for a short time, it was only evaluated on two volunteer sequences. In both cases, the ductile end of the temperature sensor was placed close to the subject's nostril, and provisorily taped to this position (see figure 2.4). This solution is naturally a bit cumbersome and comparatively inconvenient for the patient. However, it was sufficient to evaluate the applicability of a thermometer for respiratory gating in principle. The elastic pad of the pressure based respiration sensor was used, too, and was taped to the volunteer's abdomen. Then short respiration sequences were acquired using the Anzai belt and the ART camera as comparative sensors in parallel.

Since the BioVet temperature probe is not a dedicated respiration sensor, the resulting signal is not proportional to the thorax expansion during breathing. Instead, the temperature rises during expiration and drops during inspiration. This makes a landmark based synchronization with other respiration sensors rather hard. As a work-around, we leveraged the fact that the BioVet temperature data and BioVet pressure data were read from the same file, and hence perfectly synchronized. While the pres-



Figure 2.4.: The Biovet temperature probe placed close to the patient's nose.

sure based BioVet signal was not worth further evaluation, each sequence contained at least a few recognizable respiration cycles. Since this signal was based on the same physiological effect as the Anzai belt and the ART camera (namely thorax expansion), it could be used for synchronization. The BioVet temperature curve was then shifted accordingly. When it was additionally negated, this respiration curve corresponded reasonably well with respiratory signals from the other sensors. To calculate an even better respiration curve from the temperature data, an integral of the signal might help, probably alongside some constant correction factors and offsets. However, such sophisticated steps were omitted in this short feasibility study.

While the BioVet CT1 System was relatively easy to use, that is not relevant for the use of temperature measurements for respiratory gating. According to the manufacturer, the clinical use of BioVet on humans is not advisable, so completely different system would be necessary in the end. Such a respiration sensor could be focused on the temperature part, and be tailored to medical imaging. In particular, a more sophisticated method to attach the temperature sensor to a patient's nose would be required.

2.5. ART Stereo Infrared Camera

The last respiration sensor examined is probably the most interesting one. Actually the ART Stereo Infrared Camera offered by Advanced Realtime Tracking GmbH (Weilheim, Germany, <http://ar-tracking.de>) is not a respiration sensor, but an all purpose tracking system. It is a marker based system, that can be used to track 3D locations and motions of a large variety of objects. Its dedicated fields of operation are virtual and augmented reality, measuring, motion capturing, industrial tracking and medical applications. In the scope of this IDP, the capabilities of the ART camera in the context of respiratory gating were examined. The basic idea is, to track

2. Methods for Testing Respiration Sensors

the motion on several points on the patients abdomen or thorax, and reconstruct a respiration curve from this motion.

The feasibility of such an approach has already been proved by Beach et al. and is documented in [4]. The same research group also investigated the use of a stereo infrared tracking system² for measurement of rigid body motion in the context of PET imaging [3]. The use of the ART camera in this IDP was intended both to verify their findings and to explore further possibilities to correct respiratory and irregular body motion in medical imaging. Note that a motion tracking system as flexible as the one provided by ART has the potential to measure unwanted body movements more directly than any other respiration sensor. In the end it is not breath flow, temperature, or body circumference decreasing image quality, but the superposition of variety of patient movements.

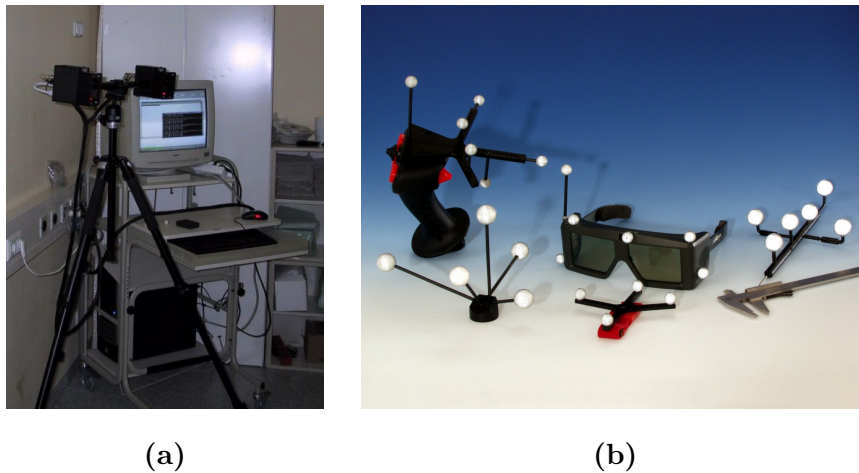


Figure 2.5.: The ART tracking system, including stereo infrared cameras and computer (a), and a repertory of reflexive ART markers mounted to various objects (b).

The 3D tracking systems offered by ART come in different variations. For this IDP, the comparatively compact and mobile *smARTtrack1* model was used. Other infrared tracking systems offered by ART use three or more cameras, which have to be mounted permanently and need to be re-calibrated when moved. The *smARTtrack1* uses only two cameras, which are permanently connected to each other. (This unit of two cameras is depicted in figure 1.8c in the introduction.) Both distance and angle of the two cameras are fixed, so the system can be moved as a whole, without the need for re-calibration.³ The *smARTtrack1* system, simply called ART camera throughout

²They used the Polaris Tracking System from Northern Digital Inc. (Waterloo, Ontario, Canada, <http://www.ndigital.com>).

³Due to small mechanical defects caused by transportation and due to thermal drift, re-calibration

this text, is usually placed on an ordinary tripod otherwise used for photography. The necessary mounts are delivered with the ART camera. The only other component necessary to operate the ART camera is a desktop computer running the accompanying software, *DTrack* [1]. The infrared cameras are connected to the PC by Ethernet, which transfers all the data acquired. However, there is another proprietary connection between the two cameras and the computer, needed for synchronization. The computer needs to be equipped with a special PCI extension card, handling that synchronization connection. Compared to the notebook computers needed for the other respiration sensors, the whole ART camera system is rather bulky. To enhance mobility, all necessary computer equipment was placed on a small trolley. This is depicted in figure 2.5a behind the tripod holding the cameras. Part (b) of the figure shows examples of objects equipped with markers that can be tracked.

That said, the next section will give a brief overview of the tracking principles the ART camera is based on. This is necessary to understand some of the decisions made for our clinical setup of ART in PET examinations, which is detailed in the subsequent section. Afterwards the software for processing raw ART camera data (2.5.4) and the software for extracting respiration curves and other useful data (2.5.5) is presented. The design and implementation of this software probably made up the most time consuming part of this IDP. Yet, the limited scope of the project prevented some desirable properties of the software to be realized. The resulting programs offer all the functionality needed to perform respiratory gating, but some aspects like usability or performance had to be neglected.

2.5.1. Principles of 3D Tracking with ART

The smARTtrack1 camera system provides optical outside-in tracking, meaning that the cameras constitute a fixed spatial reference, while the tracked objects can move around. It is also a marker based system, meaning that objects cannot be tracked unless they are labeled by special markers. ART uses passive spherical markers, being retro-reflective in the spectrum of infrared light. The ART cameras also contain diodes emitting infrared-flashes, which are reflected by all markers in the field of view. The reflected infrared light can then be detected by the ART cameras, whose sensitivity is tailored to the respective wave-length. Normally background noise is comparatively small, so each marker can easily be segmented.

When the location of a marker in the camera image is known, a straight line from the camera to the marker can be reconstructed. Below, such a straight line is referred as optical ray of the marker. The distance between marker and camera cannot be estimated from the optical ray, unless the marker's size is known very precisely. As a matter of fact, ART systems work with markers of different size, and do not depend on this knowledge. Instead, the principles of stereoscopic vision are utilized to determine the distance of a marker. This is depicted in figure 2.6a for three markers M_1 , M_2 ,

is still advisable from time to time.

and M_3 . Both of the smARTtrack1 cameras, deliver an optical ray to the marker. The precise marker location can be determined by the intersection of the two optical rays.

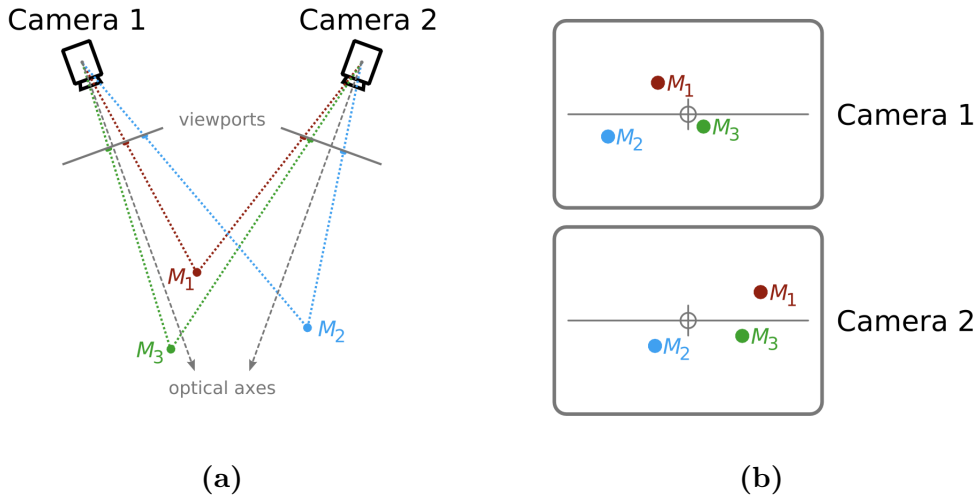


Figure 2.6.: The principle of stereoscopic vision with two cameras.

The markers in 2.6a are colorized for the sake of clarity only. In reality, the ART system cannot distinguish a single marker from another one. That’s why correspondences between markers viewed by the two cameras have to be determined computationally. In the rare case that two or more markers and both of the cameras are coplanar, the correspondence may be ambiguous. In this case, phantom markers can appear, but heuristics to handle such situations exist. Other problems occur, when markers partially overlap each other from a camera’s point of view. As far as the ART system is concerned, such details are handled automatically by the corresponding software, *DTrack*.

Though the *DTrack* software is capable to display raw camera data (as in the foreground of figure 2.7), it is focused on additional processing. It computes exact 3D coordinates for each marker according to the above principles of stereoscopic vision. Each marker location is given as a tuple of x, y, and z coordinates, with respect to the reference coordinate system of the ART setup (“room”). Hence single markers are sometimes called three-degrees-of-freedom-objects (3DOF) in the context of ART data. Each 3DOF object is assigned an ascending ID number when it first appears, and it keeps this ID as long as it is tracked. However, when the marker cannot be tracked temporarily – because it leaves the cameras’ field of view or due to occlusion – it is assigned a new ID after it reappears.⁴ In the background of figure 2.7 another *DTrack* window is shown, which contains a list of 3DOF markers with their respective coordinates on the right side.

⁴Such unique marker IDs seem to be an advantage compared to random switches of markers tracked with *Polaris*, as described in [3].

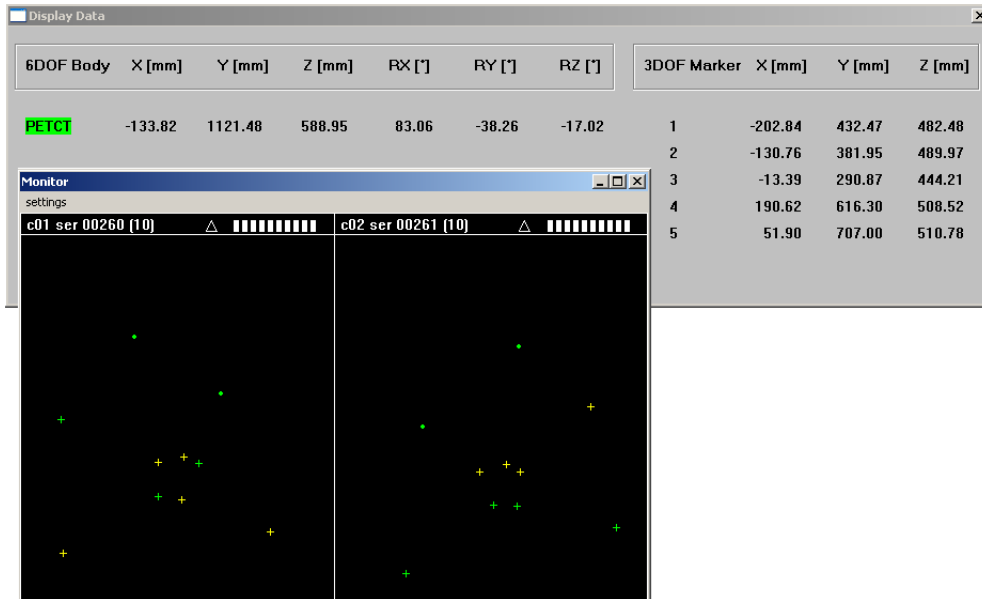


Figure 2.7.: The DTrack software controlling the ART tracking system, with a marker monitoring window and a corresponding list of tracked objects.

Markers are tracked continuously, and their locations are recomputed 60 times per second. The precision of marker coordinates given by an ART system depends on a huge variety of factors, like camera setup, marker quality, and marker location with respect to the camera. An accuracy of 3 mm is guaranteed within the camera's field of view, but in many cases sub-millimeter precision is realistic. For smARTtrack1 motion along the camera's optical axis is generally tracked less accurate, than motion perpendicular to it.

While marker locations can be tracked quite precisely with the ART camera, they don't reveal any information about their orientation. To obtain the orientation of objects, several markers have to be combined. When three or more markers are connected rigidly, both location and orientation of the resulting object can be tracked. For this, the tracking system has to know the constellation of those markers, so it can try to match this template with the tracked locations of the markers. This kind of registration problem can be solved by an iterative closest point approach, as the one described in [26].

In the context of the ART camera, such rigid marker groups are called bodies. Both their location in three dimensions, and their orientation with respect to three axes is tracked, hence they're also called six-degrees-of-freedom-objects (6DOF). The DTrack software handles 6DOF object transparently, meaning that it recognizes them among other markers and calculates their orientation. For this to work, the user has to

define the body by calibrating a constellation of four or more markers, and assigning an ID and name. Unlike independent markers, a body’s ID is persistent, and it can be recognized even if it was temporarily occluded. The window in the background of figure 2.7 shows a body named “PETCT” along with its location offsets and its rotation angles. Note that ART does not display markers as 3DOF objects, when they are part of a known 6DOF object. However, when enough markers comprising a body are temporarily occluded, the body cannot be recognized any more, and the remaining markers are displayed as 3DOF objects.

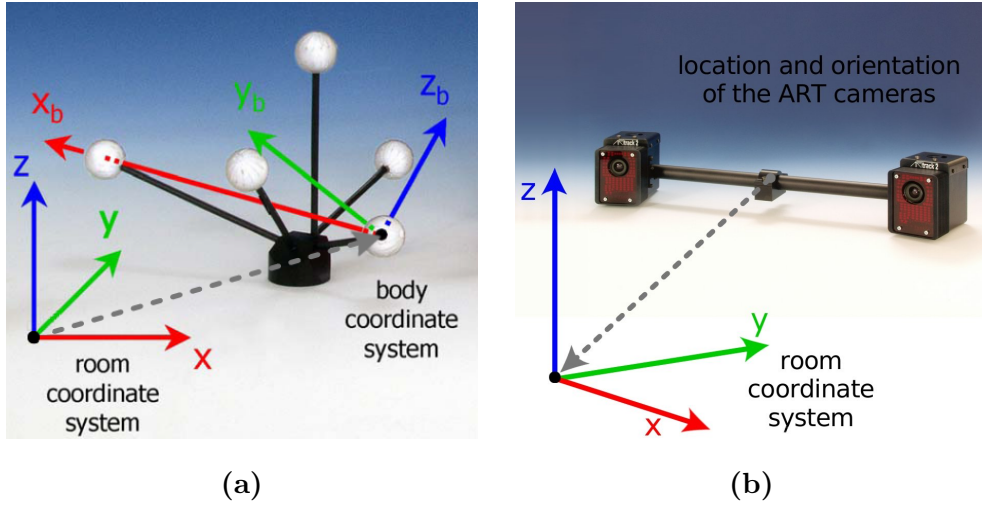


Figure 2.8.: Each 6DOF body defines an own coordinate system (a) while the location of the cameras defines the room coordinate system (b).

It should be mentioned, that the location and orientation of an ART body can be expressed by means of a linear transformation from body coordinates to the room coordinates. Each 6DOF object defines such a body coordinate system⁵, as is visualized in figure 2.8a. Let \vec{b} be the current position vector of a body and \mathbf{B} the current rotation matrix giving its orientation. Then a point \vec{p}_b given in body coordinates can be transformed to the corresponding point p given in room coordinates as follows:

$$\vec{p} = \mathbf{B}\vec{p}_b + \vec{b}.$$

The so called room coordinate system is the coordinate system in which all ART tracking data are given. Even this coordinate system can be calibrated by the user, but the calibration process is slightly more elaborate. Once the room coordinate system is calibrated, it is fixed with respect to the location of the smARTtrack1 cameras. This

⁵ART offers different strategies how the body coordinate system is defined. For details see appendix A.2 of the ARTtrack & DTrack Manual [1].

relationship is illustrated in figure 2.8b. Note that both room and body coordinate systems are always orthonormal, right-handed, and labeled in millimeters.

2.5.2. Clinical Setup for Tracking Respiratory Motion

Based on the principles of stereoscopic vision, and the capabilities of the ART camera, respiratory motion was estimated in this IDP. For this purpose, markers had to be attached to the patient. The first decision was whether to use 6DOF bodies or 3DOF markers for tracking. While 6DOF bodies provide more information about the object they are attached to, they are rather bulky, usually having a diameter of at least 5 cm. It did not seem feasible to affix more than two or three such bodies on the patients. Even for a single 6DOF body it is not trivial to link it appropriately to a deformable surface such as the patient's body. Due to its leverage the 6DOF body might jiggle too much, leading to noisy tracking data. Besides, wearing huge marker mount, like the ones depicted in figure 2.5, might seem quite awkward to the patient. Though the use of such a rigid marker block lying on the patient's abdomen has been shown by Nehmeh et al. in [16], we decided to focus on the use of multiple independent markers.

For the patient study the patients were labeled with three markers on the patient's upper abdomen (one medial and two slightly lateral), and two additional ones on their upper thorax, slightly below the clavicles. Preliminary tests showed, that a single marker placed almost anywhere on the upper body of the patient is sufficient to obtain a respiration curve. However, the use of additional markers improved reliability, and made the tracking more robust against both patients with varying breathing patterns and temporary marker occlusion. Also, a larger number of independent markers distributed over the patient's body, partially makes up for the lack of orientation information. By evaluating motion from all the markers together, translation and rotation of the thorax can be estimated.




Figure 2.9.: ART markers alongside some accessories (a) were mounted to patients (b) in a preliminary study.

The use of independent markers made labeling comparatively easy. To assure good marker visibility small plastic mounts were crafted, as shown in figure 2.9a. Such mounts can be attached to patients in different ways, e.g. applied on a plaster, placed on a tight shirt, or via elastic belts, as proposed in [4]. Since we already used the Anzai belt in the comparative study, we simply attached three of the markers to the belt.⁶ The markers were screwed to the marker mounts, which in turn were taped down to the belt. The additional two markers were fixed to ECG electrodes, which were then stuck directly on the patients skin, along with the one used for ECG monitoring. (See figure 2.9b.)

Proper positioning of the ART camera was another challenge, due to the narrow patient port of the PET/CT scanner. While [16] and [4] suggest a camera location at the patient's feet, the Siemens Biograph Sensation 16 PET/CT promoted camera placement at the bed-head. That's because the other end of the scanner is occupied by the elaborate patient bed, and also the PET part is build into the front end of the scanner's case. As a result, the patient's thorax is more visible from the bed head, during a PET scan of the region. Hence the ART camera was placed on the scanner axis, approximately $1\frac{1}{2}$ meters in front of it. It was then raised to about 1.8 meters, and the two cameras of the smARTtrack system were aimed approximately to the expected location of the markers during PET acquisition, as shown in figure 2.10. Though the patient port of the Biograph Sensation 16 is comparatively broad, this was about the only camera location avoiding marker occlusion by both the scanner's case and the patient itself. Luckily this camera position corresponds well with the distance recommendations given for the ART system. The bulky computer trolley cabled to the ART cameras was placed to the side of the scanner, out of the patient's view.

Since the ART camera was shared with other projects and had to be moved around in the hospital, we were not able to install it permanently in the PET/CT room. This meant, that the camera had to be build up and adjusted to an appropriate position each time a gated PET scan was made. This setup procedure is far more time consuming than for any other of the tested sensors. Additionally the system has to be re-calibrated occasionally and camera sensitivity needs to be adjusted from time to time. The data processing described further below is more complex, too, since some parameters have to be adjusted manually to obtain good respiration curves. While some of these caveats could be avoided, if the ART camera were fixed permanently to the wall, the use of an infrared 3D camera system only seems worth the effort, if the additional data can actually be harvested to improve PET image quality. However the results described in chapter 3.2 indicate that this might be the case indeed.

From the patient's point of view, the use of the ART camera only causes little discomfort. That is because the markers placed on the patient may not be covered by a blanket or clothing. The rest of the ART system works contact-less. Though it does emit electromagnetic radiation, it is definitely not ionizing, since it's in the spectrum

⁶This approach is independent from the rest of the Anzai system, since it also works when the load-cell sensor is not placed in the belt.

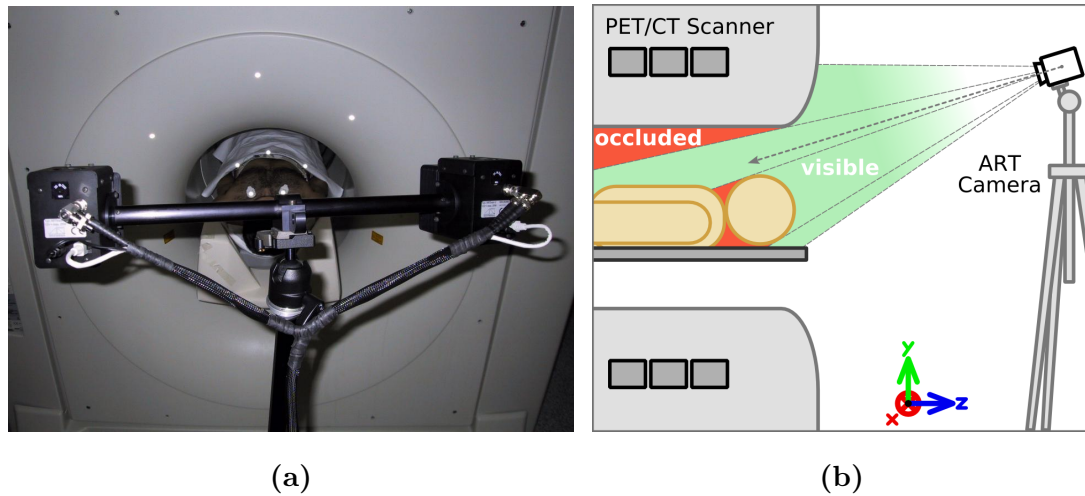


Figure 2.10.: The setup of the ART camera system with respect to the PET/CT scanner (a) and a sketch of the camera's field of view (b).

of infrared light. The power output of the infrared flashes is up to 30 W according to the manufacturer, so thermal effects can be neglected, too.

2.5.3. Coordinate System Calibration

When ART camera systems are mounted permanently, the room coordinate system can be calibrated to match some real-world features of interest. In our context, a correspondence between ART coordinates and the ones used in PET images was desirable. But since the room coordinate system is bound to the camera location (which could never be reconstructed precisely enough) it would be different each time. Recalibrating the room coordinate system takes several minutes, and requires a special room calibration body. This 6DOF object would have to be aligned with the coordinate system used by the PET/CT scanner, which is not explicitly known. Due to those issues, an arbitrary room coordinate system was used, not having any special relation to the scanner.

To obtain calibrated marker locations nevertheless, additional calculations were necessary. The transformations were based on a 6DOF body, permanently attached to the PET-CT scanner's case. To minimize mechanical modifications of the scanner, this reference body was made up of flat ART markers stuck around the patient port, as shown in figure 2.11a. Flat markers were permissible here, because they stay almost perpendicular to the camera's line of sight all the time. Since the reference body is mechanically linked to the scanner, an invariant transformation from the body's coordinate system to the one defined by the PET/CT scanner exists. On the other hand, tracked marker positions (\vec{p}) can easily be transformed to the coordinate system

2. Methods for Testing Respiration Sensors

defined by the tracked reference body (\mathbf{B}). This transformation is denoted \mathbf{B} in figure 2.11b and can be calculated by reforming the corresponding equation given in section 2.5.1:

$$\vec{p}_b = \mathbf{B}^{-1}(\vec{p} - \vec{b}).$$

The resulting marker positions (\vec{p}_b) in reference body coordinates can then be converted to PET/CT coordinates by applying a constant transformation. To obtain the corresponding transformation matrix (\mathbf{C}), calibration measurements have to be performed. Normally this is done by measuring a set of points in both coordinate systems, and applying a registration algorithm like the one described by Horn in [10]. In [7] Bruyant et al. demonstrated how this is done using multi-modal markers visible by both PET and the stereo infrared cameras. This was achieved by placing a drop of a radioactive tracer in the well of the marker spheres before screwing them to the mounts.

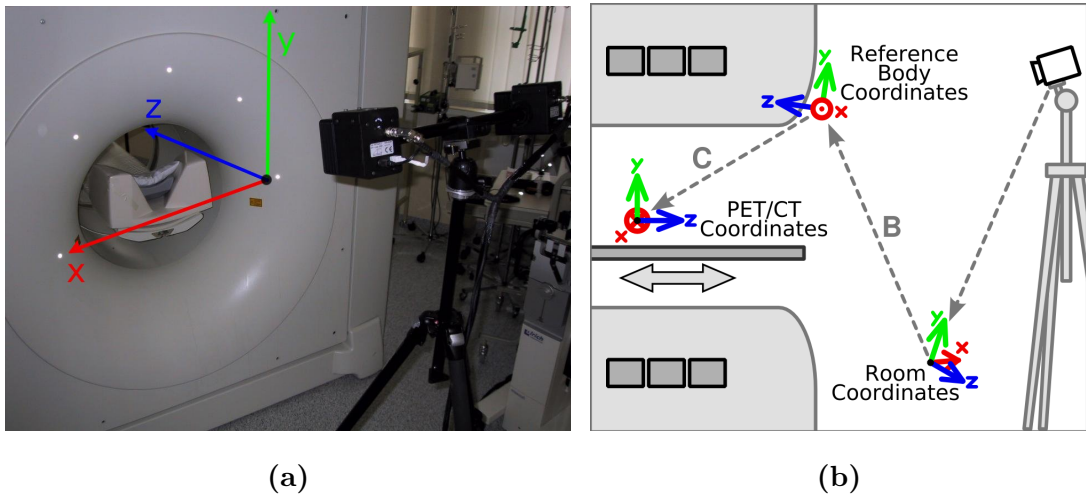


Figure 2.11.: A reference body affixed to the PET/CT scanner (a) is the basis of the coordinate system calibration applied (b).

For this feasibility study, we relied on a even more primitive calibration. For the measurements of patient motion and respiration we were primarily interested in calibrating orientation, while absolute translation information were negligible. To obtain calibration measurements, we utilized the fact, that the patient bed could be moved exactly along the scanner axis, that is along the z-axis of the PET/CT coordinate system. The bed could also be lowered and raised along the y-axis of that coordinate system. We placed an ART marker on the bed and moved it along the axis, while tracking it with the ART camera. We also wrote down the corresponding bed positions, which were displayed in millimeters by the scanner electronics. A list of such correspondences can be used to register the two coordinate systems, e.g. by using the

absolute orientation algorithm by Horn.

However, in our case an even simpler solution for the rotation matrix \mathbf{C} was available. When normalizing the movements along the PET/CT axes, two very simple position vectors \vec{p}' and \vec{q}' are obtained. These can be related to the corresponding normalized motion vectors \vec{p}_b and \vec{q}_b from the ART measurements:

$$\vec{p}_b = \begin{pmatrix} p_{bx} \\ p_{by} \\ p_{bz} \end{pmatrix} = \mathbf{C}\vec{p}' = \mathbf{C} \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \vec{q}_b = \begin{pmatrix} q_{bx} \\ q_{by} \\ q_{bz} \end{pmatrix} = \mathbf{C}\vec{q}' = \mathbf{C} \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

Due to the many zeros appearing in the vectors, and due to normalization, six components of \mathbf{C} are explicitly given by the above equations:

$$\begin{aligned} c_{12} &= p_{bx}, & c_{13} &= q_{bx}, \\ c_{22} &= p_{by}, & c_{23} &= q_{by}, \\ c_{32} &= p_{bz}, & c_{33} &= q_{bz}. \end{aligned}$$

To obtain the missing components of the calibration matrix \mathbf{C} , additional constraints are utilized. Since the sought-after \mathbf{C} transforms between two orthonormal coordinate systems, it must be a pure rotation matrix (without scaling or skewing). Hence \mathbf{C} must also be an orthogonal matrix, that is, all its column vectors must be orthogonal to each other and must have the length 1. This results in the following equations constraining the missing column vector \vec{c}_1 :

$$\begin{aligned} \vec{c}_1 \circ \vec{c}_2 &= c_{11}c_{12} + c_{21}c_{22} + c_{31}c_{32} = 0, \\ \vec{c}_1 \circ \vec{c}_3 &= c_{11}c_{13} + c_{21}c_{23} + c_{31}c_{33} = 0, \\ |\vec{c}_1| &= \sqrt{c_{11}^2 + c_{21}^2 + c_{31}^2} = 1. \end{aligned}$$

These three equations can easily be solved arithmetically, but the closed form solution is rather bulky and is omitted for the sake of brevity. It should be mentioned, that the equations yield two possible solutions for \vec{c}_1 , differing only in sign. The correct solution for the calibration matrix \mathbf{C} is chosen, so that its determinant of \mathbf{C} is positive.⁷ This means, \mathbf{C} is a right-handed matrix, which makes sense since it transforms between two right-handed coordinate systems.

The calibration matrix \mathbf{C} describes a rotation from PET/CT coordinates to the coordinates of the reference body. By applying the inverse of \mathbf{C} to tracked marker positions in reference body coordinates, they are transformed to a coordinate system aligned with the PET/CT scanner. This transformation is denoted \mathbf{C} in figure 2.11b. Altogether a position vector \vec{p} given in the raw ART marker data can be transformed

⁷Actually the determinant should be exactly +1, if the measurement error is omitted.

to a position vector \vec{p}' in PET/CT coordinates according to this equation:

$$\vec{p}' = \mathbf{C}^{-1}\mathbf{B}^{-1}(\vec{p} - \vec{b}).$$

Note that the axes of the resulting coordinate system are aligned with the PET/CT scanner, but the origin is arbitrary. This is sufficient to quantify relative movements of a patient lying on the PET/CT bed and to classify them. For instance, it can be told whether a movement occurred in cranial-caudal direction or in anterior-posterior direction. However, the coordinate system does not reveal any absolute position information, so it is not known where exactly an ART marker would appear in a PET image. It should also be mentioned, that the patient bed remains unmoved during the acquisition of one PET position. That's why it is not necessary to calibrate with respect to the location of the bed. Otherwise a reference body placed to the bed would be required.

2.5.4. Processing of Raw ART Camera Data

Before any ART data can be evaluated at all, they have to be recorded using the DTrack software. While DTrack makes tracking data instantly available via the network, no real-time processing was performed for this feasibility study. Instead, the data were written to text files by an off-the-shelf tracking server. Those ART data files were then processed by the IDL software package `ARTCamera` which was written in scope of this IDP. This section explains the processing steps implemented by the `ARTCamera` software. They can be called on the IDL command-line, but they are also integrated in the IDL application `SignalAnalysisGUI`. For software design details and usage hints related to those two software packages, please refer to appendix A, respectively B.

ART data files consist of sequential tracking data, organized by samples. The sampling rate of ART is limited to 60 Hz but lower rates can also be emitted by DTrack. For ART files representing respiratory motion, sampling times as low as 20 Hz turned out to be sufficient. The values in ART files are whitespace-separated and are tagged with a short string, which is given in braces in the following. Each sample in an ART file consists of a header line containing a serial number (`frame`), a timestamp (`ts`), and some information on the numbers of objects tracked. After the header, information for each 6DOF body is given in an own line. Then come 3DOF markers one per line, and finally the sample ends with an empty line as separator. Each 6DOF entry contains, the ID (`bod`) specified in in DTrack during body calibration. Then comes a 3D position vector (`loc`) representing the body's current location in the ART room coordinate system. Finally, the orientation of the body is given twice, first represented by three rotation angles (yaw, pitch, roll), second represented by a 3×3 rotation matrix (`rot`). Entries for 3DOF bodies are similar, but only an ID (`mar`) and a position vector (`loc`) is given.

Figure 2.12 summarizes the structure of ART files, explained above. The first step in processing ART tracking data is parsing the file using regular expressions to extract




Figure 2.12.: A diagram depicting the structure of ART text files – they contain samples consisting of markers and bodies.

this data structure. The result is an array of samples, each containing data of all tracked objects visible at the respective moment. This number can vary due to marker occlusion, so the data has to be consolidated and interpolated where necessary. The first task is to extract a uniform array containing all the marker positions. The array shall contain three dimensions, the first one representing the consecutive samples, the second one representing the distinct markers, and the third one representing the three location coordinates. Such an array is used by all further processing steps, and is called marker array below.

When extracting the marker array, a problem is that markers can get occluded and they have a different ID when reappearing. The marker array is filled up with NaNs for all samples not containing the respective marker. A newly appearing marker ID is first checked for correspondence with all currently occluded markers. If such a marker exists and has a position close enough to the new one, they are considered to be the same marker. Hence, the location of the new marker is assigned to the existing one. This behavior can be affected by adjusting a distance threshold, which estimate the maximum distance a marker moves during occlusion. New marker IDs appearing further away from any occluded markers are actually considered new, and get their own column in the marker array. For the patient study distance thresholds of 1 cm to 3 cm were chosen as necessary by the user. Such big tolerances are required, since a marker might get occluded e.g. during inspiration and become visible during expiration, hence changing its location quite a lot. In particular, the markers placed on the abdomen could get out of the ART cameras view for quite a time, when the bed was moved out to apply drugs to the patient. (Luckily that never happens during the actual PET scan.)

Due to measurement artifacts and unnecessary assignments of new marker IDs, the resulting marker array often contains far more markers than attached to the patient. Experience showed, that an hour of monitoring five patient markers yielded about 10 to 15 columns in the marker array. One particular effect contributing to this number is that the markers comprising the reference body can temporarily appear as single 3DOF markers. This happens when the reference body is partially occluded (e.g. by a technician passing through the cameras field of view) and the DTrack software cannot recognize it as 6DOF body. The markers actually used for patient labeling can cause

problems, too, e.g. when long occlusions occur due to poor marker placement. Even short marker occlusion can lead to new columns in the marker array, if the distance threshold is chosen too small.

Such effects lead to a marker array containing lots of missing data, marked by NaNs. Before further processing, markers with little meaningful data should be discarded. The `ARTCamera` software offers a filter discarding markers exceeding a certain number of missing samples. The filter value can be specified by the user in terms of seconds. This filtering method proved very useful to retain meaningful marker data. Alternatively the user could specify the desired markers manually according to their location.

Even when the marker array is reduced to the markers required for further processing, it still contains missing data due to temporal marker occlusion. Therefore the `ARTCamera` software interpolates the location data, for each interval a marker is missing. Linear interpolation of the boundary values in 3D is used where possible. But for occlusions at the beginning or the end of the sequence, boundary values are simply repeated. Interpolation leads to straight lines in a marker's motion where tracking data was missing. It turns out that the negative effect of such false data on respiratory motion estimation is negligible, if other markers contain useful data for the same period of time.

By now, the marker array contains all the location data required. However, they are given with respect to the ART room coordinate system, which has no particular meaning. So the locations need to be transferred to the PET-CT coordinate system, as explained in section 2.5.3. In order to do so, the reference body has to be extracted from the ART data structure. This is easy, since ART bodies do have a unique ID, in contrast to markers. From each sample, position vector and rotation matrix are obtained, and combined into one 4×4 homogenous matrix. These homogenous matrices representing the reference body are collected in an uniform array. Such a body array has three dimensions, the first one representing the consecutive samples, while the second and third one contain the rows and columns of the homogenous matrix. Body data missing due to occlusion are interpolated by closest neighbor sampling. In the setup used for the patient study, neither the ART camera nor the reference body are moved during the acquisition. Hence the homogenous matrix representing the reference body should be almost the same for all samples. However, the body array containing independent values for each sample, makes the setup tolerant against small accidental camera movements.

The transformations contained in the body array can now be applied to the marker array. Since IDL is an array oriented language, it is easy to iterate over all samples, invert the transformation matrix, and multiply all the markers to it. The resulting marker array is in reference body coordinates, but in order to transform it to PET/CT coordinates, another transformation has to be applied. This is given by a constant calibration matrix \mathbf{C} , which was obtained by calibration measurements, as described in section 2.5.3.

Figure 2.13 visualizes the steps necessary to obtain consolidated, interpolated and calibrated tracking data. The resulting marker array contains data quantifying respi-

ratory motion. To use this data for respiratory gating, further processing is required, as indicated in the figure. The challenge here is to digest a one-dimensional respiration curve from all the independent marker movements, as detailed in the following.

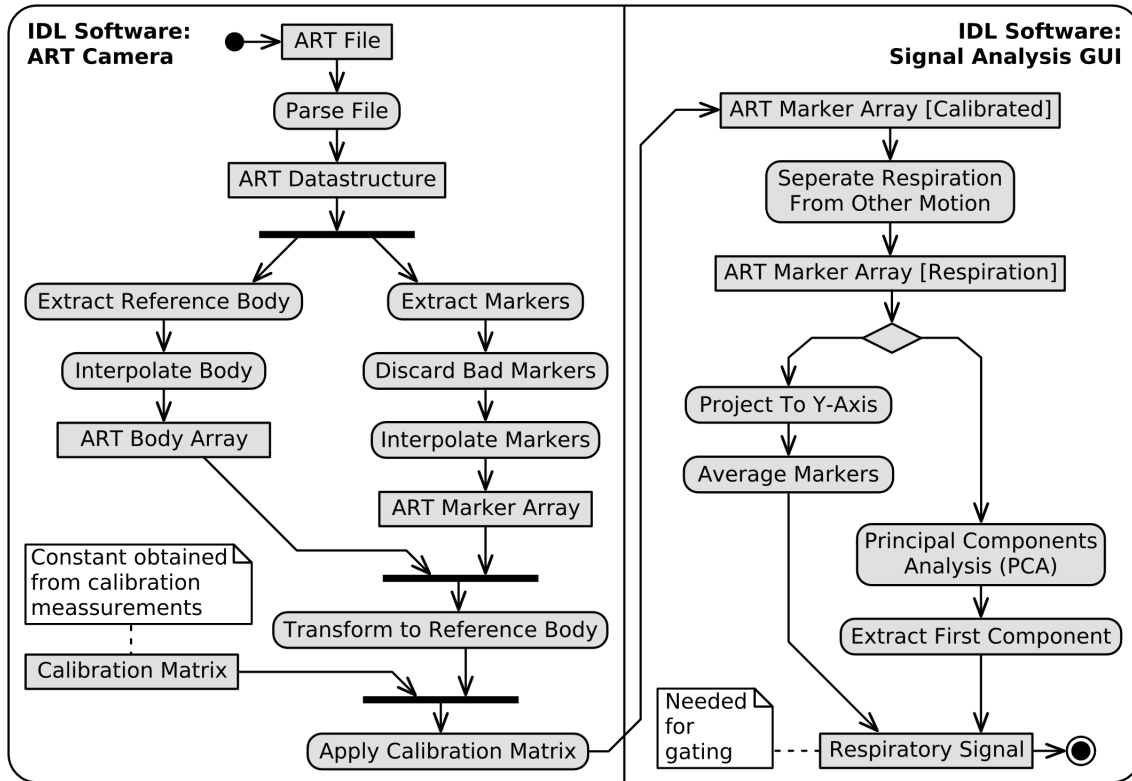


Figure 2.13.: The processing steps necessary to consolidate ART tracking data, and to extract a respiration curve.

2.5.5. Extraction of Respiration Curves

The marker array obtained by the previous processing steps contains a huge diversity of information for each sample. When n markers are tracked, there are $3n$ degrees of freedom involved. In contrast, a respiratory signal used for gating only consists of one value plotted against time. This section describes how all the available marker data are filtered and digested to obtain a respiration curve. As a by-product, information about large-scale non-periodic body motion is obtained, which can be considered independent of respiration. Such body motion information cannot be used for respiratory gating, but it is worthwhile to analyze it more closely. (See chapter 3.2.)

Actually, the separation of respiratory motion from other body motion is the first

step to perform on the marker array. Since it turned out to be easier to filter for data not related to respiratory gating, large-scale body motion is obtained first. Subtracting the resulting data from the marker array yields the respiration related portions of the marker motions. In [6] a sophisticated method for separating body motion is presented. Unfortunately this method could not be applied in scope of this IDP, since it has not been published in time. Instead, a rather simple approach using a low-pass frequency filter was used. This was implemented by setting up a digital filter for frequencies lower than 0.1 Hz and convoluting it with the marker array. This filter was applied for all spatial components of all markers separately.

While this filter method is not perfect, it is good enough to separate respiratory motion and other body motion for further processing. Note, that for ART acquisitions over the whole duration of a PET/CT the body motion also contains movements of the patient bed. This movement can be quite substantial (up to one meter along the scanner axis) and it is important to filter it, when a respiratory signal over the whole ART sequence is desired. In this IDP two independent methods to digest respiratory signals from a marker array were implemented. The comparisons in chapter 3.1.2 show that the results of both methods are virtually equivalent. However, both of the methods have their assets and drawbacks, which are stated after the following explanations.

The first, more trivial, approach is to project the 3D marker locations to one axis. This results in a one-dimensional motion for each marker, which is quantified by a scalar number. Any of these numbers could be used as respiratory signal, but by computing the average over all markers a more robust signal is obtained. For the sake of simplicity, we projected the marker locations measured in the patient study to the y -axis of the PET/CT scanner. That corresponds to an anterior-posterior movement of the patient's abdomen, and makes respiration well visible. Projecting to the z -axis of the scanner (cranial-caudal direction) would also yield viable results. Actually, projection to any linear combination of those two axes shows significant respiration. Only the scanner's x -axis is not suitable to detect respiration, since the symmetrically placed markers compensate each other's motion. Respiration curves obtained by projection to an axis have the following implications:

- + The resulting data represent the respiratory motion along the axis quantitatively. That's useful for comparative patient studies, and might even have applications in improved image correction techniques. However, temporarily occluded markers can mess up the quantitative measurement.
- Coordinate system calibration is strictly required in order to specify the correct axis.
- Any information contained in motion orthogonal to the axis gets lost. In rare occasions, the chosen axes may not work well with the patient's physiology.

An alternative method to digest a respiration curve utilizes a *principal components analysis* (PCA, see [18]). This mathematical method is used to consolidate multi-dimensional data, by transforming it to a new basis sorted by the variance of the

data-set. That is, the first component of the basis is chosen such, that the data-set has the highest variance in that dimension. When applying PCA to the marker array, each spatial coordinate of all the markers is considered an independent dimension. Then the first component of the resulting PCA is the desired respiration curve. This works, because the only source of significant variance in the marker location is the respiratory motion. The use of this method to digest respiratory signals has several pros and cons:

- + Does not require any coordinate system calibration at all.
- + Combines the most meaningful data from all dimensions. This is true independently of any physiological peculiarities of the patient.
- Only works, if large-scale body motion and bed-movements are removed from the marker array. Otherwise those processes having a huge variance would affect the direction of the first principal component.
- The range of the resulting respiration curve is chosen arbitrarily, and does not have any quantitative meaning. The curve can even be upside down, meaning that inspiration is represented by smaller values than expiration. This has to be corrected manually for respiratory gating.

Despite all the differences, both methods were applicable in scope of our clinical setup. However, the PCA method may be useful in situations, where the tedious process of coordinate system calibration is omitted. In that case it is the more reliable way to obtain a respiration curve. On the other hand, coordinate system calibration is essential, if any more general motion correction based on ART tracking data is considered. As far as only the respiration curves are concerned, both digesting methods are almost equivalent. This is shown in the next chapter, where the ART respiration curves are compared to those obtained from the other sensors.

3. Evaluation of Resulting Sensor Data

The ultimate goal of using the respiration sensors presented here is to improve the quality of PET images by respiratory gating. As shown in [16] the improvement can be quantified, e.g. by segmenting a tumor in the PET image, and calculating its volume or the amount tracer activity contained. Due to reduced motion blur, gated images show the tumor much more precisely, and therefore lead to reduced tumor volumes. The difference can be significant for small tumors, but in the context of our cardiac study it is far harder to quantify. There is a difference in respiratory gated PET images of the myocardium, as shown in section 3.1.3, but it is not as significant as it would be for smaller structures.

That's why this evaluation of respiration sensors is focused on the comparison of the different sensors among each other. The goal was to find out, if all of the sensors deliver good respiratory signals. Also, particular strengths and weaknesses of the signals should be pointed out. The results can be used to pick the appropriate sensor for future patient studies. Due to the special capabilities of the ART camera, there is an own section discussing its potential for more general movement correction at the end of this chapter.

3.1. Comparison of the Sensors

The comparison of the respiration curves was conducted using three IDL programs `RespCurves`, `RespGatingGUI`, and `SignalAnalysisGUI`. (The latter was implemented as part of this IDP, and is described in appendix B.) The synchronization of the signals was performed manually with those tools, as is detailed in the descriptions of the respective sensors in chapter 2.

Due to the organizational issues mentioned earlier, not all of the sensors could be tested together. In the following, some interesting portions of respiration curves are compared visually for two sensors at a time. The Anzai belt is used as the reference system here, and is involved in each of the comparisons. Among other things the analysis of respiratory signals shows different artifacts, caused by non-regular respiration. Such artifacts can propagate into the gates file, and reduce the quality of gated PET images. After these qualitative considerations, a statistical analysis of the similarity of several respiration curves is given.

3.1.1. Qualitative Comparison

Regular breathing was detected similarly by all evaluated sensors. In the respiration curves it manifests as a zigzag curve with slightly rounded peaks. Often the ascending parts of the signal are a bit steeper and shorter than the descending ones. This corresponds to the fact, that the inspiration phase is shorter than the expiration phase for most people. The left part of figure 3.1 shows such a sequence of regular respiration cycles acquired with the Anzai belt and the PMM spirometer. Since none of those respiratory signals have a quantitative meaning, the respiration curves were rescaled and shifted, to be more illustrative. Note that the slightly varying frequency of the respiration cycle matches well in both curves and all peaks are aligned perfectly.




Figure 3.1.: Two corresponding respiratory signals, acquired with the Anzai belt, respectively with the PMM spirometer.

The rest of the sequence shown in figure 3.1 reflects different kinds of irregular breathing. First, the volunteer performed a small snatchy movement, which resulted in a slight distortion of its respiration. Note that the PMM spirometer software, which performs automatic calibration continuously, got confused by that and produced some artifacts. Later, the characteristic visualization of a breath-hold at full inspiration appears in the respiration curve. Though slightly differing, breath-holds were recognizable in data acquired with all the sensors. They were hence used for manually checking the correct synchronization of respiratory signals acquired simultaneously. Note, that the Anzai belt shows a lower frequency signal during breath-hold, which can be attributed to the volunteer’s heart-beat. While heart-beat cannot normally be detected in the respiration curve, it becomes visible when the Anzai belt is under high pressure, and no overlaying respiration is present. The breath-hold is also seen in the PMM respiration curve, but once again it caused some calibration artifacts.

Unfortunately, the PMM spirometer was even less reliable for some patients, and the automatic calibration feature messed up the signal completely. While failures with some test subjects can be blamed to its prototype status, irregular breathing always

resulted in PMM artifacts worse than for Anzai. Anomalous breathing situations as the one visualized in figure 3.2a could also occur during a PET scan, especially if the patient’s health is impaired. If such artifacts are only temporal, the effect on respiratory gating is negligible, but a good respiration sensor should not get too confused. Another problem encountered with the PMM spirometer prototype is that its temporal resolution jitters sometimes. This makes it impossible to synchronize it precisely with other sensors, let alone with a PET list-mode acquisition. As shown in figure 3.2b, the time seems to progress slower in the PMM respiration curve, as compared to Anzai. Unfortunately this behavior is not reproducible exactly, and no constant correction factor for the sampling rate could be determined.

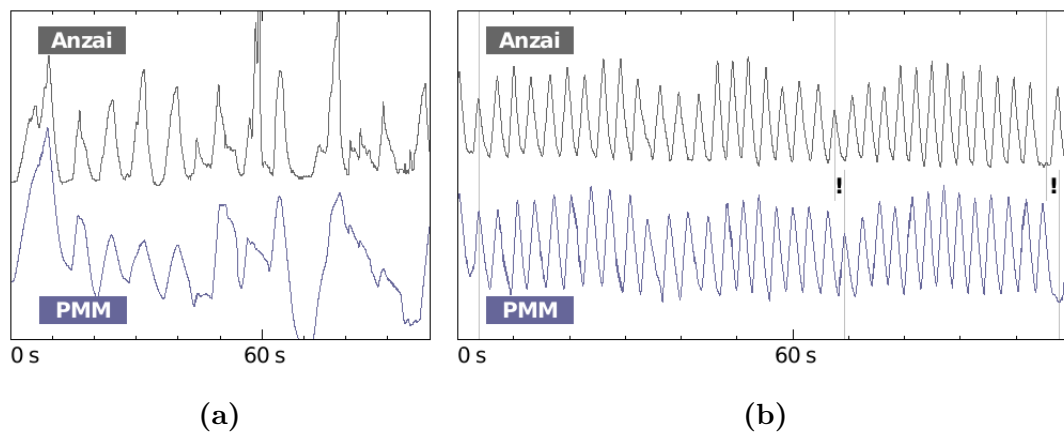


Figure 3.2.: Respiration curves obtained with the PMM spirometer prototype show several artifacts, mainly due to poor calibration (a) and timing problems (b).

Next, respiratory signals acquired with the BioVet system are evaluated, though it has not been tested too thoroughly. However, figure 3.3 shows that the temperature sensor can monitor respiration nicely. The figure also includes data obtained from the BioVet pressure sensor pad. It shows good signal peaks in the beginning, but the low signal-to-noise ratio renders it useless at the end of the sequence. We were not able to affix the pad well enough, to get useful results over a longer period of time.

The fact that both of the BioVet signals are synchronized implicitly was used to check proper synchronization of the BioVet temperature sensor with the Anzai belt. In figure 3.3 all of the respiration curves correspond quite well, only the sign of the BioVet temperature sensor data needs to be inverted. In some of the other sequences acquired with BioVet, the temperature signal was also lagging behind slightly. To get results comparable to the Anzai belt, those respiratory signals had to be shifted manually. In the scope of our short experiments, we were not able to quantify this effect properly. But it seems, that it might be linked to physiological discrepancies of

3. Evaluation of Resulting Sensor Data

different patients.

The occasional lag of the temperature signal implies, that gates calculated from the Anzai belt and the BioVet temperature sensor divide the respiration cycle differently. However, the negative effect on gated images should be small, if respiration is regular enough. In the end, only the duration of the respiratory cycle is important for the gating algorithm used in this IDP. And that can be measured reliably with the BioVet temperature sensor.




Figure 3.3.: An Anzai respiration curve compared to the BioVet pressure and temperature sensors.

Since both the Anzai belt and the ART camera measure respiration by the expansion of thorax and abdomen, the resulting respiration curves correspond best. As the statistical comparisons in the next section show, the signals are very similar for both measurements on volunteers, and the patient study involving actual PET scans. The comparison also showed, that it does not matter whether projection or PCA is used to extract the respiration curves from ART data. Respiration curves obtained by projection respectively PCA all showed a correlation (see next section) of more than 98%, usually around 99.5%. So the use of PCA was chosen arbitrarily for this section.

In some cases, especially during pharmaceutically induced stress in scope of the PET scan, the waveforms in the two respiration curves can look rather different. Also, the respiratory signal provided by Anzai, sometimes has a baseline varying at a very low frequency. Such effects are caused by the actual breathing process, and can hardly be suppressed by the patient. Respiration curves calculated from ART camera data tend to hide such effects, since a frequency filter is applied during the extraction process. All this may cause the two respiration curves to look slightly different, as for instance the ones in figure 3.4a. However, the effect on respiratory gating is often negligible. Figure 3.4b visualizes the gates computed from the signals above, and in fact they're almost identical.

On the other hand, some artifacts may lead to rather different gates, even if the




Figure 3.4.: Slightly irregular respiration curves acquired during a stress PET scan with Anzai and ART (a), and the resulting gates (b).

respiration curves bear resemblance to each other. As for the curves depicted in figure 3.5a, the Anzai belt regularly shows an additional peak at full expiration. This turned out to be due to poor placement of the belt, and the particular situation could have been avoided by more careful preparations and monitoring. However, it shows, that even relatively small changes of the wave-form can confuse the gating algorithm quite a bit. The resulting gates shown in figure 3.5b indicate a big difference of the respiratory signals. This is caused by the additional peaks in the Anzai signal, which are sometimes falsely identified as full inspiration.

In this case, an operator could just reduce the sensitivity for detecting peaks in the gating algorithm, and both curves would result in identical gates. However, it makes sense to use a sensitivity working in most of the cases without need for manual adjustment. In other cases, actual physiological effects may cause respiration curves from different sensors to yield completely different gates. In that case, it is hard to tell, which one represents the physiology of respiration better. Larger patient studies including gated PET images, as the ones presented in section 3.1.3, are needed to evaluate such special cases.

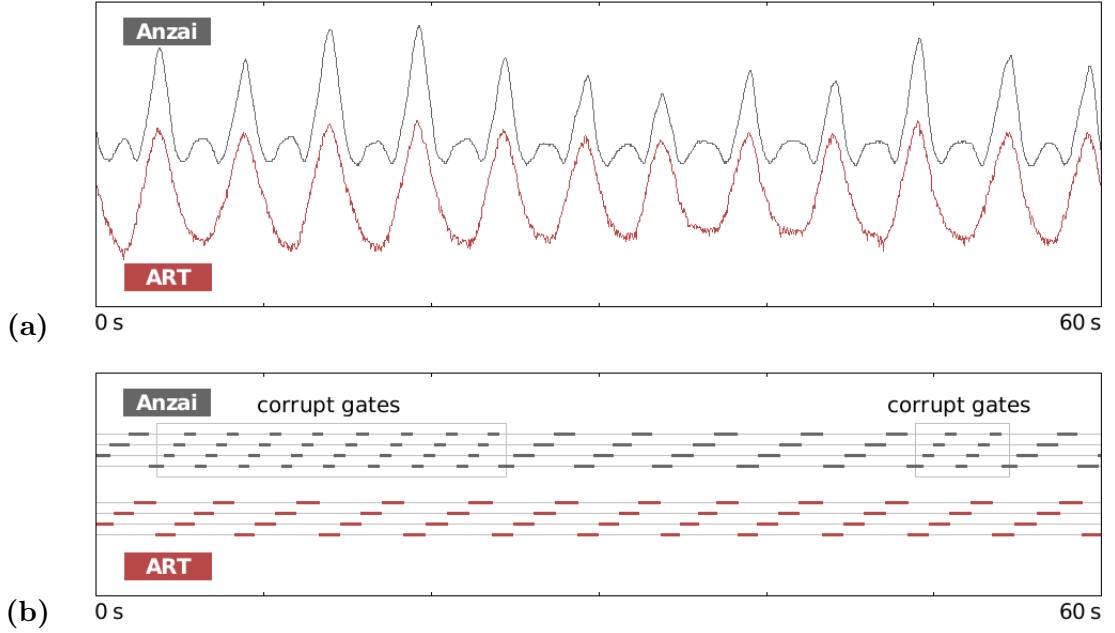


Figure 3.5.: Respiration curves from a sub-optimally placed Anzai belt and the ART camera (a), and the resulting gates (b).

3.1.2. Quantitative Comparison

The above comparison of respiratory signals reveals some similarity in almost any case. However, to extract meaningful results from several acquisitions with volunteers and patients, the similarity of respiration curves had to be quantified somehow. Hence, comparison functions were built into the `SignalAnalysisGUI` software detailed in appendix B. After a short description of the comparison techniques used, the respiration sensors are compared quantitatively in this section.

An obvious approach to compare two signals, is to calculate their correlation coefficient. This statistical method is based on the covariance of the two data sets, but it also involves normalization. Hence the correlation coefficient is robust to different signal amplitudes. For the comparison of two respiration curves of the same duration, they were first re-sampled to the lower sampling rate, if necessary. Then the following formula was applied to obtain the correlation coefficient $c_{R,S}$:

$$c_{R,S} = \frac{\text{cov}(R, S)}{\sigma_R \sigma_S} = \frac{E((R - \mu_R)(S - \mu_S))}{\sigma_R \sigma_S}.$$

Here R and S are distributions representing the signals to be compared. E and μ give the expected value of a distribution, while σ gives its standard deviation. The

covariance function (cov) is just a different notation for the corresponding term on the right-hand side of the equation.

The correlation coefficient gives a good estimate of the similarity of respiratory signals. However, the last section showed, that there is no general relationship between the similarity of the curves, and the correspondence of the resulting gates. In particular, certain differences of the waveform can result in a comparatively bad correlation, while the computed gates are identical. Therefore a comparison in the context of respiratory gating should be performed on the gates, rather than on the original respiratory signals.

The trivial approach to compare the gates data by correlation, too, does not seem sensible at the second glance. Mismatches between the lowest and the highest gate number would contribute to the result overly. This effect has no correspondence in the real-world application of the gates. Hence, another method for comparing gates was conceived for use in this study, lapidary dubbed *gated correspondence*. This method simply calculates the fraction of samples having the same gate number assigned in both of the gated signals. If G and H gated signals of the length n , the gated correspondence $c_{G,H}$ is computed as:

$$c_{G,H} = \frac{1}{n} \sum_{i=1}^n \begin{cases} 1 & \text{if } G_i = H_i \\ 0 & \text{otherwise} \end{cases} .$$

The higher the gated correspondence value, the more list-mode data are assigned to the same gate. Hence the gated correspondence suggested here is linked to the similarity of gated PET images based on the respiratory sensors compared. It is considered the preferred comparison method in the following. Note, that values of gated correspondence and correlation cannot be compared directly, though both have absolute values in the range from 0 to 1.

Figure 3.6a shows values for both correlation and gated correspondence of the Anzai belt compared to the ART camera. The underlying respiration curves were acquired in scope of the ongoing patient study, and each sequence monitors respiration during an 11 minutes list-mode scan. In the diagram, each pair of bars represents the comparison values – correlation and gated correspondence – resulting from one respiration sequence. Note that each patient’s respiration was monitored twice per PET scan, once in rest and once under stress. The first three comparisons in the diagram, are based on respiration curves of three patients in rest. The next three comparisons originate from the same three patients (in the same order) monitored during pharmaceutically induced stress.

Of course three patients are not sufficient to extract reliable information. Unfortunately list-mode PET scans of the thorax are rather rare at our department of nuclear medicine. However, the patient study will be continued over the next few months, which will also deliver more reliable data on respiratory gating. In this IDP, respiration of volunteers was acquired with the Anzai belt and the ART camera, too. However, those sequences cannot be compared to the patient study directly, since they were much shorter and did not include stress conditions. Generally speaking, the re-

3. Evaluation of Resulting Sensor Data

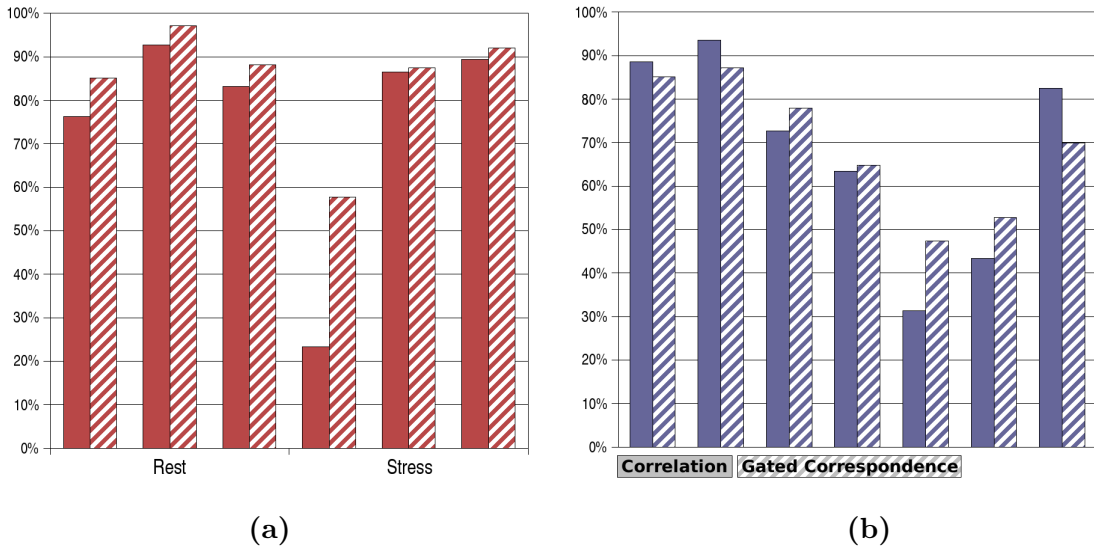


Figure 3.6.: Statistical comparison of respiratory signals acquired with the Anzai belt to those of the ART camera (a) and the PMM spirometer (b).

spective comparison results were in the upper region of the results shown in figure 3.6a.

In spite of the few measurements underlying the comparison, some tendencies can be derived from the results. For instance, the fourth column shows, that even poor correlations of different respiratory signals can lead to passably good correspondence of the resulting gates. This can be explained by effects described in the previous section. Another tendency seen in the data, is that respiration is harder to measure and yields less comparable curves, when the patient is under stress. This seems natural, since the stress leads to faster, and often more irregular breathing. However, those effects are not stringent, as the data of the third patient show.

The PMM spirometer could not be used in the patient study, due to organizational problems. Hence, comparison could only be based on acquisition with volunteers. Figure 3.6b shows results of the PMM spirometer compared to the Anzai belt for short sequences of three minutes. The data for the first two volunteers show, that the PMM spirometer can produce good respiration curves, corresponding to the ones obtained with the Anzai belt. The reduced correspondence for the rest of the volunteers can be blamed partially on some timing problems encountered with the PMM spirometer, as shown in figure 3.2b. This was obviously a technical problem of the PMM spirometer, and will most likely be overcome when it is released as a commercial product.

The BioVet temperature sensor was not used in the patient study, either. In fact, we were only able to acquire the respiration of two volunteers with it. The comparison results for both of them were quite similar to each other, and it is not overly illustrative

to show them in an own diagram. Instead, figure 3.7 shows the mean comparison results acquired for each pair of sensors. These results should be regarded with utmost cautions, since the data underlying the three columns are rather different.




Figure 3.7.: An overview of the average correspondence of respiratory signals acquired with the different sensors.

For instance the first two bars represent the correlation and gated correspondence of the Anzai belt compared to the ART camera are based on the data from the patient study presented above. Though some volunteer acquisitions were added to compute the mean values presented here, they incorporate real-life respiration monitoring of actual PET patients. In contrast, the two bars in the middle are based on only two short volunteer measurements, obtained under laboratory conditions. Also, the respiratory signals of the Anzai belt and the BioVet temperature sensor compared here, were synchronized manually in order to fit best. However, the occasional lag of the BioVet temperature sensor might impair the comparison results substantially, and could have negative effects on respiratory gating. Finally the last column in figure 3.7 shows the mean comparison results of Anzai belt and PMM spirometer. As noticed before, the comparatively low correspondence can be blamed on the PMM's timing problems.

Although a larger number of test acquisitions would provide more accurate estimations, the values in figure 3.7 make clear, that the signals of all respiration sensors evaluated in this IDP are somehow similar. None of the sensors was a complete failure, and basically all of them might be used for the reconstruction of gated PET images. That is also true for the PMM, if the occasional timing problems are solved.

3.1.3. Comparison of Gated Images

Reconstruction of respiration gated PET images was only done with the Anzai belt and the ART camera. The main focus in this feasibility study was, to find out if respiratory gating had any effect on the reconstructed PET images at all. In the long run, one of the benefits expected from respiratory gating is, that the tracer distribution in a given region can be quantified more precisely, due to reduced motion blur. However, such quantitative measurements were not performed in scope of this IDP.




Figure 3.8.: Images reconstructed from different gates of the same PET scan, the first on expiration (a), the second on inspiration (b).

Instead, figure 3.8 shows the visual effect of respiratory gating. The two gated PET images depict the same coronal slice through the thorax region of a patient at rest.¹ The first image only contains PET data acquired at full expiration, while the second one contains PET data from full inspiration. The myocardium surrounding the left ventricle of the heart has a high perfusion, and can therefore be seen as region of high tracer activity. In the pseudo-colors used for the images, it is displayed as a bright orange to yellow ring.

When looking at the two gated images closely, it can be seen that the heart is slightly displaced between the two. In particular, the lower part of the bright ring (representing the inferior wall of the heart) is shifted down in the second image. This reflects the physiological motion of the heart during the respiration cycle. It is moving upwards on expiration and downwards on inspiration. Using the gated PET images, a rough estimate of the magnitude of this motion can be given. For the patient depicted in figure 3.8 it is about 4 millimeters. This measurement is not overly precise, due to the limited spatial resolution of PET and the limited temporal resolution of the respiratory

¹For cardiac diagnosis, PET images of the hearth are usually rotated, so that the axis of the left ventricle is exactly horizontal. In contrast, the slices shown here are aligned with the PET coordinate system.

gates.

While the images shown in figure 3.8 were gated using the ART camera, the use of the Anzai belt yields quite similar results. For instance, the respiratory signal obtained with the Anzai belt was used to reconstruct the gated PET images shown in the introduction (figure 1.1b) from the same PET scan. Though the correlation of the two respiratory signals was only 76%, the resulting sets of gated images are quite similar. Only small regions vary in intensity, due to a slightly differing noise distribution. But the overall motion of the heart is clearly visible, regardless of the respiration sensor used.

3.2. Motion Estimation Beyond Respiratory Gating

The reconstruction of respiration gated PET images was the only motion correction technique applied in this IDP. However, the motion data acquired with the ART camera revealed far more information on the patients' respiration. The ART camera was the only respiration sensor in this study, which could actually quantify magnitude of the motion of the patient's abdomen and thorax. Table 3.1 shows the amount of respiratory motion measured along the PET/CT coordinate axes for all three patients. No data are given for respiratory motion components along the x -axis (left-to-right direction), since this component of respiratory motion is not significant. In fact, it was below 1 mm for all the patient's observed.

	x -axis	y -axis	z -axis	absolute
Patient 1	N/A	5.0 mm	2.5 mm	5.6 mm
Patient 2	N/A	5.5 mm	2.5 mm	6.0 mm
Patient 3	N/A	10.0 mm	4.5 mm	11.0 mm
Average	N/A	6.9 mm	3.2 mm	7.5 mm

Table 3.1.: The maximum magnitude of respiratory motion of three different patients.

The values in table 3.1 were obtained manually, by searching the respiration curves for a section with high amplitude. Then the marker showing the highest motion amplitude was selected, which was the marker placed in the middle of the Anzai belt in all examined cases. For this marker the difference between its location at full inspiration and at full expiration was calculated in each dimension. While these values represent the motion for the axis-aligned components, the absolute value was determined by simply calculating the length of the resulting motion vector.

The above values quantifying respiratory motion are just maximum values. For the first two patients, visual inspection of the respiration curves showed, that the magnitude of respiratory motion lies close to this maximum during the whole PET scan. In contrast, the magnitude of respiratory motion varied a lot for the third patient. Here,

phases of extremely shallow breathing periodically changed for phases with extremely deep inspirations. This resulted in the high maximum respiration magnitude for this patient. Note, that such a breathing pattern is hardly suited for respiratory gating, which is based on the assumption of similar respiration cycles repeating periodically.

Another degrading effect in PET images, that cannot be corrected by respiratory gating, is the presence of non-periodic body motion. Such motion can be caused by small twitches due to patient discomfort or by muscle relaxation and motion of internal organs. As mentioned before, a main advantage of the ART camera is, that it can estimate such body motion. However, only the motion of the patient's surface can be measured by placing markers on it. It is hard to extrapolate the motion of the PET scan's region of interest, since that requires detailed knowledge of the patient's anatomy.

Moreover it is expected, that measured body motion can slightly differ among the markers placed on the patients thorax and abdomen. This region of the body is moving non-rigidly, and cannot be described by rigid or affine transformations perfectly. Nevertheless, to get an estimate of the body motion occurring during a PET scan, it is useful to assume simple translation of all the markers. In this IDP, the centroid of the markers was computed for each sample, and the resulting signal was filtered to remove periodic respiratory motion. This is similar to the filtering step described in chapter 2.5.5, but in this case additional smoothing was applied, to remove respiratory motion even more thoroughly.

Figure 3.9a shows the average magnitude of body motion measured in the patient study, along each PET/CT coordinate system axis. Three values are given, first for the motion during the PET scan under rest condition, second for the motion under stress (both 11 minutes). The third value reflects the overall motion during the complete PET/CT examination, including the rest and the stress scan.² Note, that the motion of certain markers may be much higher, but since only the centroid of marker motion is given here, different markers compensate each other.

The amount of body motion measured was quite similar for all three PET patients examined. So was the progression of patient motion: in most cases it consisted of slow drifts over a few dozen minutes. However, all three of the patients showed a comparatively abrupt body motion lasting only one or two minutes. After some investigation we found a temporal coincidence between this event and the injection of Adenosine. This is the drug artificially causing stress, which was started to be applied about one minute before the stress scan. Figure 3.9b shows motion curves for the three PET patients at the time of the Adenosine injection. The depicted interval reflects the beginning of the six minutes of Adenosine injection, where motion was visible most clearly.

Note, that figure 3.9b shows the body motion along the z-axis of the PET/CT

²Actually, body motion could only be measured while the patient bed was not moving, since bed motion was not calibrated. In the PET protocol used, the bed was only moved for CT scans, and stood still during the majority of time. Therefore the overall body motion given here covers about 60 minutes per patient.

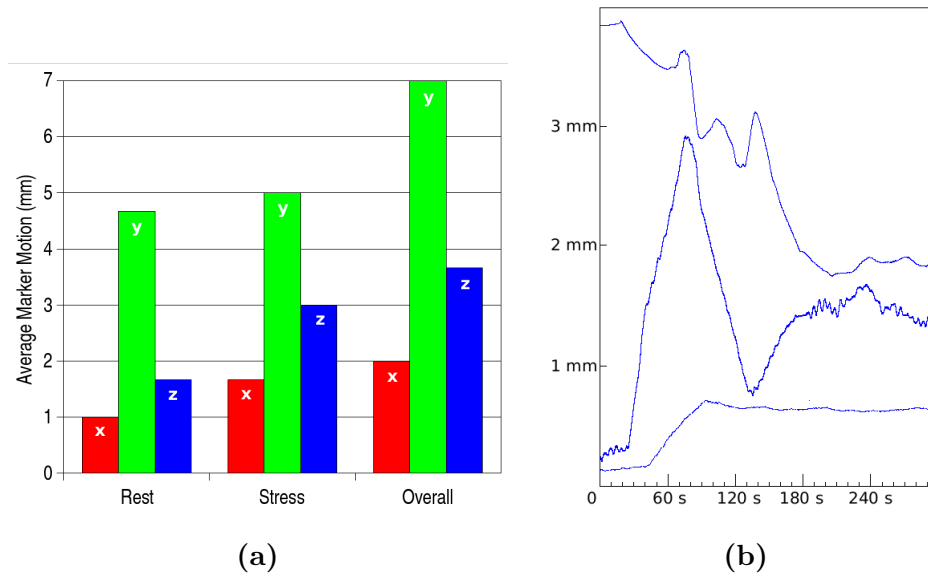


Figure 3.9.: The average body motion of PET patients (a), and the progression of abrupt body motion upon Adenosine injection.

(cranial-caudal direction). The magnitude of the motion is rather small in these curves, since they display the average motion of all markers. Yet, for all three patients, the motion which occurred during this short period of time was about half as much as the motion measured during the overall PET scan. The respiratory motion after Adenosine injection also reflected the effect of the drug. Respiration tended to become heavier, and also more irregular. This led to worse comparison results of the respiration sensors during stress PET scans, as has been shown in figure 3.6.

3. *Evaluation of Resulting Sensor Data*

4. Conclusion

While the benefit of sophisticated patient motion estimation with the ART camera is yet to be proved, plain respiratory gating does make sense in PET imaging of the thorax region. This will be even more true in the near future, when hot-spot studies revealing small structures using specialized tracers become more common. This IDP helped to gain experience with some respiration sensors and gated reconstruction techniques on the cusp of clinical use. A summary of those experiences is given in the next section, while the subsequent one suggests some future improvements.

4.1. Outcome

The central statement, resulting from the comparison of the four respiration sensors, is that all the different sensors are basically suitable for respiratory gating. Due to technical and organizational issues, not all of them were actually used for reconstruction of gated PET images. Nevertheless, all obtained respiration curves showed patient respiration clearly. The fact, that the sensors are exchangeable with respect to respiratory gating, is underlined by the quantitative data given in chapter 3.1.2.

This makes clear, that the sensor principles used as a basis for respiratory gating can be quite different. Whether they're based on pressure, temperature, air-flow or motion tracking, they all can yield suitable respiration curves. Since none of the respiratory signals has principal advantages over the others (as far as respiratory gating is concerned), practical considerations can be focused. That is, whenever respiratory gating is deployed, the choice of the respiration sensor should be based on "soft" factors like usability, invasiveness, reliability and costs.

As seen in chapter 2, not all of the actual sensor implementations are fulfilling these quality requirements. Figure 4.1 classifies the respiration sensors according to our practical experiences. The relative values given here are based on the sensors as they were used in scope of this IDP. However, there is a lot of room for improvement for some of the sensors, no matter on which principle they are based on.

In fact, of the compared sensors, only the Anzai belt is ready for clinical use today. The PMM spirometer prototype and the BioVet temperature sensor are impaired by their low reliability. As far as the BioVet temperature sensor is concerned, this might be overcome by additional data processing steps, to handle possible signal lag correctly. In this IDP however, the developed processing software was dedicated to the ART camera, due to the promising possibilities for more general motion estimation. Though not a dedicated respiration sensor, the ART camera was able to monitor patient respiration

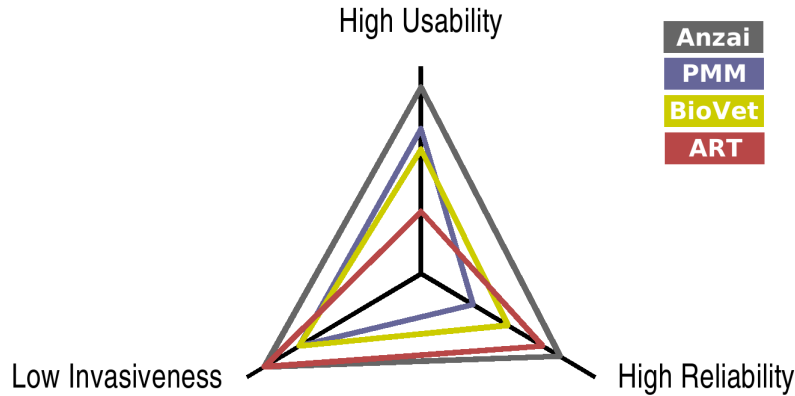


Figure 4.1.: Subjective evaluation of the respiration sensors according to some important quality requirements.

quite reliably. It’s main deficiency is the low usability, due to rather tedious setup and calibration procedures.

In the end it can be said, that sophisticated sensors like the ART camera only make sense, if the additional data are actually used for enhanced image reconstruction. As used in this IDP, the ART camera is far too cumbersome for respiratory gating alone. Nevertheless, the ART camera is a promising tool for further research on patient motion and improvement of PET images. Some ideas that may be worth realization are sketched in the following.

4.2. Future Work

Most importantly, the respiratory motion measurements and gated image reconstructions have to be repeated on additional PET patients. Unfortunately, the cardiac patient study involving list-mode PET scans, started off with a rather moderate patient throughput. In fact, during the first three months of the study, only three patients were actually examined. However, the situation has improved recently, and more patient data are available by now. Those could not be evaluated for this report anymore, but will be examined by the PET research group.

Acquisitions of additional patients can be used to examine large-scale non-periodic patient motion during long PET scans more thoroughly. Though the motion examined for the three study patients was similar in magnitude and characteristics (e.g. the reaction to Adenosine injection), it may be way different for others. Only additional measurements can show, if non-periodic body motion has a significant impact on PET image quality in the first place. The motion magnitudes measured so far are in the region of PET scanned resolution, and therefore do not contribute to image blur overly.

Then again, the values for body motion given in chapter 3.2 only give a rough estimate of the actual body motion. They merely quantify the translation of the centroid of all tracked markers. The variance in translation of the different markers needs to be examined, too. It is likely, that single markers show more significant body motion. Generally, the area of abdomen and thorax can exhibit highly non-rigid motion.

The degree of non-rigid motion involved should be estimated by applying a rigid point based registration method between marker locations measured at distinct times. If the error in the resulting rigid transformation is big, the patient motion has to be considered non-rigid. However, if it turns out that a large portion of body motion occurring during PET scans can be described by rigid transformations, the motion might be corrected during image reconstruction. This can be done directly on list-mode data, as shown by Bloomfield et al. [5] in the context of neurological PET, but it requires detailed knowledge of the geometry of the PET scanner's scintillator crystals. Alternatively the motion can be segmented into several parts with only little motion, and separate images can be reconstructed. Those can then be combined to one image, after appropriate transformation is applied.

If body motion of PET patients turns out to be highly non-rigid, corrections will be hard without sufficient knowledge of the patient's anatomy. This is because the ART camera can only track motion on the surface, and does not provide enough information to interpolate motion inside the body. In any case, a big part of non-rigid motion in the thorax area is related to respiratory and cardiac motion. Therefore future methods correcting large-scale non-periodic patient motion should always be combined with appropriate gating techniques. If patient studies show the need for ART camera based motion correction, it is likely to be combined with both cardiac and respiratory gating.

While a lot of research has yet to be done in the field of general motion estimation and correction, there are many pending improvements for respiratory gating, too. In our department, precise electronic synchronization to the list-mode PET scan was only implemented for the Anzai belt. For the ART camera it could be done similarly, but that would require modifications of the recording software for tracking data. Alternatively, precise time-stamp information might be included into the data recorded. Since all compared respiration sensors are PC based, clock synchronization could be performed by network protocols like NTP [15]. If that is not feasible, a radio controlled clock could be attached.

For the ART camera, software support for live monitoring of the patients' respiration would be a useful feature. While the IDL software created in this IDP is not suitable for that, a reimplementaion of the respective processing steps e.g. in C++ might be extended towards such a feature. Improved interpolation features would also be nice to have.

In general, any respiration sensor to be used clinically, needs improved usability. However, this cannot successfully happen as part of a research project. Instead, it's the manufacturers' turn now, to integrate facilities for respiratory gating with their PET scanners. Only then, larger patient studies can show the actual value of respiratory

gating for medical diagnostic. In the end, the physicians working with PET have to decide, if the benefits of respiratory gating outbalance the additional effort put into respiratory motion estimation.

4.3. Acknowledgments

This IDP was initiated by Prof. Dr. N. Navab, and was kindly advised by Prof. Dr. P. Böni. The technical respectively physical aspects of the work were supervised by Axel Martinez-Möller and Ralph Bundschuh. Further acknowledgments go to Dr. Maria-Jose Martinez, Dr. Stephan Nekolla, Dr. Sibylle Ziegler, Marc Huisman, Virginia Spanoudaki and Irene Torres for sharing their knowledge on PET and supporting the work in various ways. Isabelle Stangier made a large part of this work possible in the first place, by providing access to the ART camera, and giving advice on its usage, which is also true for Martin Bauer and Daniel Pustka.

Special thanks go to those of the above persons, who volunteered for respiratory motion acquisitions with the tested sensors. The same for all patients taking part in the ongoing cardiac study, which allows for evaluation of respiratory gating, and for all the physicians involved. Last but not least, we all thank the members of medical and technical staff who support the efforts of the PET research group.

A. ART Camera Software

This appendix gives a brief overview of the software for processing ART camera data, which was written as part of this IDP. The `ARTCamera` software is a library, which can be used from the IDL command line, or by other software packages. It can perform all the processing steps that have been shown in figure 2.13, but the functionality is split up into several classes. Figure A.1 depicts the structure of those classes:

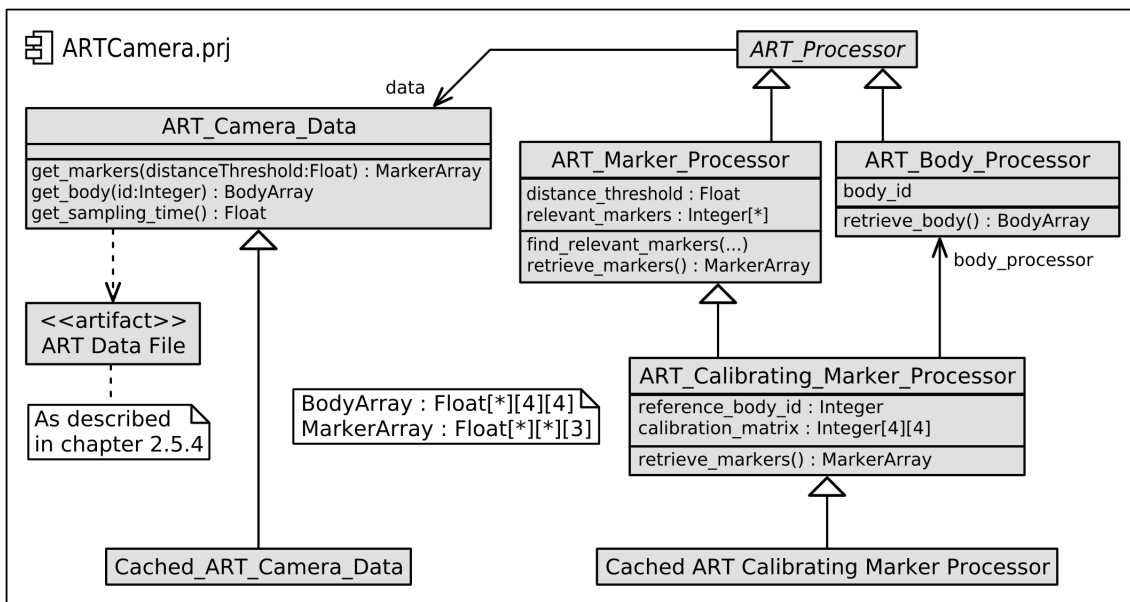


Figure A.1.: A class diagram of the ARTCamera software.

The class `ART_Camera_Data` is responsible for reading ART files and internally stores them in a data structure similar to the one that has been shown in figure 2.12. This class is also responsible for extracting marker arrays and body arrays from the raw ART camera data. If appropriate, the markers are reassigned to other IDs to handle occlusion properly. The underlying semantics are explained in chapter 2.5.4.

A hierarchy of `ART_Processors` is responsible for further processing of the ART data. Each of them can be set up with a number of parameters, and then performs processing operations accordingly. `ART_Marker_Processor` and `ART_Body_Processor` are responsible for fetching marker arrays respectively body arrays from an `ART_Camera_Data` object, and perform interpolation of missing data. `ART_Marker_Processor` can also

filter for relevant markers. `ART_Calibrating_Marker_Processor` additionally applies coordinate system transformations to the tracked markers. Therefore it delegates to a `ART_Body_Processor` retrieving the location of the reference body for each sample. It then applies the corresponding transformation to all markers, and finally applies another transformation given as constant matrix. (See chapter 2.5.3.)

Since some of the processing steps are rather slowish, `Cached_ART_Camera_Data` and `Cached_ART_Calibrating_Marker_Processor` add caching to their respective super-classes. Whenever they processing a step already performed with exactly the same parameters, they simply retrieve the result from a cache. It is advisable to use these caching enabled classes whenever possible:

```
; example values for constants used below:
DISTANCE_THRESHOLD = 10    ; in mm
REFERENCE_BODY_ID = 1     ; as defined in DTrack software
MIN_SAMPLE_DURATION = 600  ; in s

; retrieve interpolated and calibrated marker data from an ART file:
DATA1 = obj_new('cached_art_camera_data', '/PATH/TO/DATA.art')
PROC1 = obj_new('cached_art_calibrating_marker_processor')
PROC1->set_art_data, DATA1
PROC1->set_distance_threshold, DISTANCE_THRESHOLD
PROC1->set_reference_body_id, REFERENCE_BODY_ID
PROC1->find_relevant_markers, min_sample_duration=MIN_SAMPLE_DURATION
PROC1->load_calibration_matrix, '/PATH/TO/CALIBRATION_MATRIX.sav'
MARKERS1 = PROC1->retrieve_markers()
```

The above sequence can be entered on the IDL command line, to retrieve interpolated and calibrated marker data from an ART file. Users are encouraged to replace the capitalized expressions to suit their needs. The resulting marker array (`MARKERS1`) can then be evaluated on the command-line. For details on each class and function see the documentation comments in the respective IDL files.

B. Signal Analysis Software

The `SignalAnalysis` software provides a graphical user interface for the `ARTCamera` software described above. Additionally it includes a set of filters, which can be used on marker arrays, e.g. to extract respiratory signals. The `SignalAnalysis` software also contains tools to compare those respiratory signals to each other. Figure B.1 gives an overview of the classes involved:

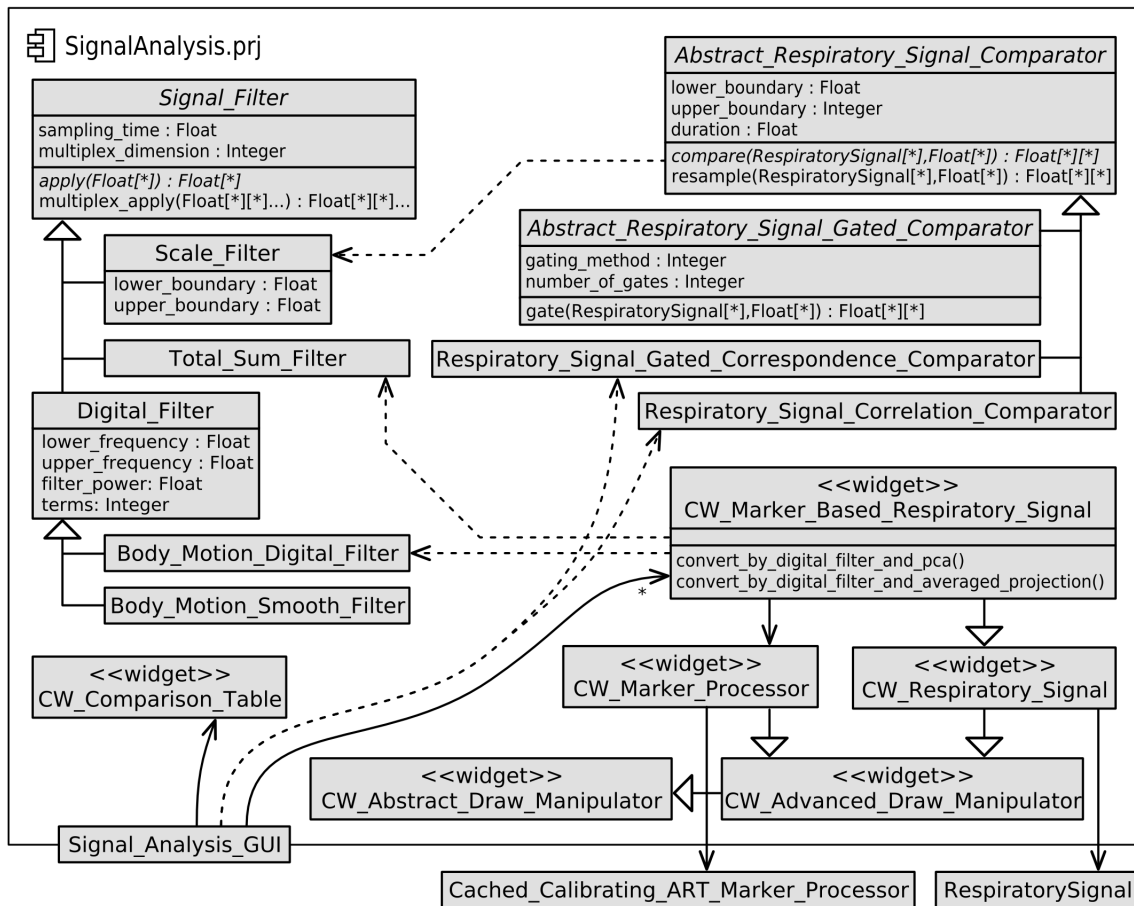


Figure B.1.: A class diagram of the Signal Analysis GUI, and the included signal filters and comparators.

The upper left part of figure B.1 contains a hierarchy of `Signal_Filters`. Each subclass implements a filter function working on a one-dimensional array. The superclass `Signal_Filter` contains the template method `multiplex_apply`, which can apply such a filter along a given dimension of a multi-dimensional array. The concrete filters offer functionality like scaling or frequency filtering. Note that there is no filter dedicated to respiratory motion. That can easily be extracted by filtering for non-periodic body motion, and subtracting it from the original signal.

Another hierarchy, shown in the upper right part of figure B.1, is dedicated to the comparison of different respiratory signals. The classes on top of the hierarchy offer utility methods for resampling and normalizing the signals before comparison. They also define a comparison method which takes an array of respiratory signals to be compared and an array of offsets into these signals. Implementing subclasses process the respiratory signals starting at the given offsets and quantify their similarity by a scalar comparison value. They return a (symmetric) matrix containing a comparison value for each pair of given respiratory signals. Currently, subclasses for the two comparison methods presented in chapter 3.1.2 are implemented.

The lower part of figure B.1 depicts the widgets used by the `Signal_Analysis_GUI` software. They are all derived from some helper classes, but the most important widgets are `CW_Marker_Processor` and `CW_Respiratory_Signal`. They display ART camera data respectively respiration curves, and offer means to manipulate them (mainly scaling and shifting their graphical representation). A screenshot of the `Signal_Analysis_GUI` showing a `CW_Marker_Processor` widget and a `CW_Respiratory_Signal` widget is shown in the screenshot in figure B.2. Note, that the widgets shown here are actually `CW_Marker_Based_Respiratory_Signals`, which combine the functionality of the two others by using tabs. This class also implements methods to extract respiration curves from marker arrays, as described in chapter 2.5.5.

The complete `Signal_Analysis_GUI` consists of several `CW_Marker_Based_Respiratory_Signal` widgets. In the bottom (not visible in the screenshot) it has controls to compare all the respiration curves. This is done by calling one of the respiratory signal comparators, and the resulting comparison matrix is then shown in a new `CW_Comparison_Table` window.



Figure B.2.: A screenshot of the Signal Analysis Software, with a Marker Processor widget in the top and a Respiratory Signal widget in the bottom.

Bibliography

- [1] ADVANCED REALTIME TRACKING GMBH, *ARTtrack & DTrack*, version v1.22 ed., November 2004. <http://ar-trcking.de>.
- [2] N. E. M. ASSOCIATION, *Digital Imaging and Communications in Medicine (DICOM)*, 2004. <http://medical.nema.org/>.
- [3] R. BEACH, H. C. GIFFORD, S. SHAZEEB, P. P. BRUYANT, B. FENG, M. A. GENNERT, S. NADELLA, and M. A. KING, *Stereo-Infrared Tracking to Monitor and Characterize Rigid-Body Motion and Respiration During Cardiac SPECT Imaging: Progress Towards Robust Clinical Utilization*, in Nuclear Science Symposium Conference Record, IEEE, 2005.
- [4] R. D. BEACH, H. PRETORIUS, G. BOENING, P. P. BRUYANT, B. FENG, R. R. FULTON, M. A. GENNERT, and M. A. KING, *Feasibility of Stereo-Infrared Tracking to Monitor Patient Motion During Cardiac SPECT Imaging*, IEEE Transactions on Nuclear Science, 51(5), October 2004, pp. 2693–2698.
- [5] 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, 48(8), April 2003, pp. 959–978.
- [6] P. BRUYANT, R. BEACH, B. FENG, and M. KING, *Compensating Combined Body Translation and Respiratory Motion in Cardiac List-Mode Spect*, Journal of Nuclear Medicine, 46 (supplement 2), May 2005, p. 476.
- [7] P. P. BRUYANT, S. NADELLA, M. A. GENNERT, and M. A. KING, *Quality Control of the Stereo Calibration of a Visual Tracking System (VTS) for Patient Motion Detection in SPECT*, in Nuclear Science Symposium Conference Record, IEEE, 2005.
- [8] P. S. CONTI and D. K. CHAM, *PET-CT. A Case Based Approach*, Springer, December 2004.
- [9] M. W. DAVIDSON, M. ABRAMOWITZ, and O. AMERICA, *Photomultiplier Tubes*, April 2006. <http://microscopy.fsu.edu/primer/flash/photomultiplier/>.
- [10] B. K. P. HORN, H. M. HILDEN, and S. NEGAHDARIPOUR, *Closed-form solution of absolute orientation using orthonormal matrices*, Journal of the Optical Society of America, 5, July 1988, p. 1127.

- [11] J. L. HUMM, A. ROSENFELD, and A. DEL GUERRA, *From PET detectors to PET scanners*, European Journal of Nuclear Medicine and Molecular Imaging, 30(11), November 2003, pp. 1574–1587.
- [12] A. C. KAK and M. SLANEY, *Principles of Computerized Tomographic Imaging*, Society of Industrial and Applied Mathematics, 2001.
<http://www.slaney.org/pct/>.
- [13] H. D. KUBO and B. C. HILL, *Respiration gated radiotherapy treatment: a technical study*, Physics in Medicine and Biology, 41(1), January 1996, pp. 83–91.
- [14] L. LIVIERATOS, L. STEGGER, P. M. BLOOMFIELD, K. SCHAFERS, D. L. BAILEY, and P. G. CAMICI, *Rigid-body transformation of list-mode projection data for respiratory motion correction in cardiac PET*, Physics in Medicine and Biology, 50(14), July 2005, pp. 3313–3322.
- [15] D. L. MILLS, *Network Time Protocol (NTP)*. RFC, September 1985.
<http://www.faqs.org/rfcs/rfc958.html>.
- [16] S. A. NEHMEH, Y. E. ERDI, C. C. LING, K. E. ROSENZWEIG, H. SCHODER, S. M. LARSON, H. A. MACAPINLAC, O. D. SQUIRE, and J. L. HUMM, *Effect of respiratory gating on quantifying PET images of lung cancer*, Journal of nuclear medicine, 43(7), July 2002, pp. 876–881.
- [17] RESEARCH SYSTEMS, INC., *IDL*, 2005. <http://www.rsinc.com/idl/>.
- [18] L. I. SMITH, *A tutorial on Principal Components Analysis*, February 2002.
http://www.cs.otago.ac.nz/cosc453/student_tutorials/principal_components.pdf.
- [19] A. A. SONZOGNI, J. K. TULI, and D. F. WINCHELL, *NuDat 2.2*.
<http://www.nndc.bnl.gov/nudat2/>.
- [20] D. W. TOWNSEND, J. P. J. CARNEY, J. T. YAP, and N. C. HALL, *PET/CT Today and Tomorrow*, The Journal of Nuclear Medicine, 45(1 (Suppl)), January 2004, pp. 4S–14S.
- [21] D. W. TOWNSEND and M. DEFRISE, *Image-reconstruction methods in positron tomography*, CERN, Geneva, June 1993.
- [22] T. G. TURKINGTON, *Introduction to PET Instrumentation*, Journal of Nuclear Medicine Technology, 2002.
- [23] E. W. WEISSTEIN, *Delta Function*, 1999-2006. From MathWorld –
<http://mathworld.wolfram.com/DeltaFunction.html>.
- [24] E. W. WEISSTEIN, *Radon Transform*, 1999-2006. From MathWorld –
<http://mathworld.wolfram.com/RadonTransform.html>.

- [25] WIKIPEDIA, *Scintillator*. from Wikipedia, the free encyclopedia, April 2006.
<http://en.wikipedia.org/wiki/Scintillator>.
- [26] Z. ZHANG, *Iterative Point Matching For Registration of Free-Form Curves*, Tech. Rep. 1658, Institut National de Recherche en Informatique et en Automatique, March 1992.
- [27] D. ZIKIC, *MR Based Attenuation Maps for PET Measurements*. Inter-Disciplinary Project, Technische Universität München, 2005.
- [28] D. ZIKIC, *Deformable Registration Methods*, Master's thesis, Technische Universität München, 2006.