Fortgeschrittenenpraktikum (FP-95)

Medical Image Analysis

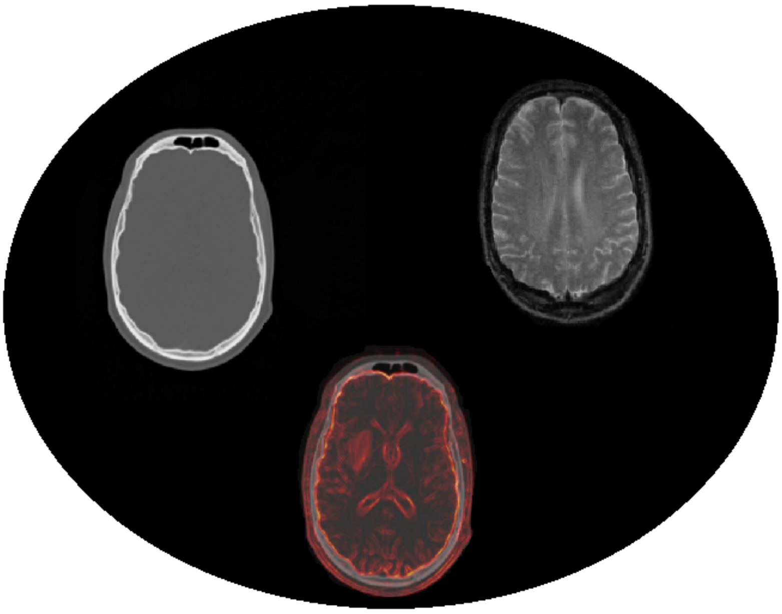

Version 1.6
(June 27, 2025)

Research Group
Computer Assisted Clinical Medicine
Mannheim Institute for Intelligent Systems in Medicine
Theodor-Kutzer-Ufer 1-3, 68167 Mannheim, Haus 3, Ebene 4

The instructions and experiment are continuously revised. To improve the experiment, please share your comments and remarks with the supervisors! We look forward to your feedback!

Contents

List of Abbreviations

B_0 Static Magnetic Field B_0

CT Computed Tomography

DT Dwell time

FID Free Induction Decay

MR Magnetic Resonance

MRI Magnetic Resonance Imaging

NMR Nuclear Magnetic Resonance

PD Proton Density

RF Radio Frequency

T_1 Spin-Lattice Relaxation Time T_1

T_2 Spin-Spin Relaxation Time T_2

TE Echo time

TR Repetition time

Introduction

Medical image registration is an essential component in the planning of medical interventions or the diagnosis of diseases. The aim of image registration is the geometric linking of image information from different images. Patient movements, different internal coordinate systems of the imaging systems, different recording perspectives or distortions in the image recordings can, for example, lead to different geometric representations of the recorded object. In these cases, it is necessary to determine the transformation between the image data so that corresponding image information can be assigned to each other. In particular, the linking of information from different imaging modalities such as CT and MRI plays a major role here.

In this practical course, the basics of image registration will be developed and tested using various examples. Due to the scope of this topic, the experiment is limited to a few aspects of registration, on the basis of which the concepts can be applied and discussed. The experiment is carried out in two full-time days. A total of 16 hours are set aside for completion, which means that the participants should prepare well in advance in order to complete the tasks within this time frame.

An important note before we start: As the experiment is still in the test phase, we are grateful for any feedback regarding problems with the individual parts of the experiment and the comprehensibility of these instructions. If you have any questions, the tutors will of course be happy to help you, but the experiment should essentially be carried out independently. The support provided by the tutors cannot and should not replace familiarization with the topic.

Basics

This section contains an introduction to the basics of medical imaging and image registration. Information on implementations can be found in the experiment execution. To prepare for the experiment, it is also worth going through the list of questions at the end of the instructions.

1 Image Modalities

Let's first look at the basics of medical imaging. Two key imaging techniques in everyday clinical practice are MRI and computed tomography (CT). Since different imaging techniques can be used to highlight different information, the choice of imaging system depends on the question being asked. To understand this, let's first look at the physics of MRI.

1.1 MRI imaging

1.1.1 Larmor Freqeuncy

The nuclei of hydrogen atoms, the protons and their spin properties, are crucial for magnetic resonance imaging. As a rotating mass, it has angular momentum and as a charge carrier it has a magnetic moment. It therefore has the properties of both a spinning top and a magnet. As with a spinning top, the proton tries to maintain the position of its axis of rotation and, like a magnet, can be influenced by other magnetic fields. As with magnets, a voltage can therefore also be induced in a receiving coil by movements of the proton. This means that both the magnetic field vector, i.e. the position of the proton's axis of rotation, and a change in its orientation, which induces a measurable voltage in a receiving coil, can be observed.

In an external magnetic field B_0, a magnet (as in a compass, for example) aligns itself along the magnetic field. However, as the proton wants to maintain the position of its axis of rotation due to its angular momentum, it performs a precession movement in an external magnetic field B_0. This has a characteristic frequency, the Larmor frequency

$$f_0 = \gamma_0 \cdot B_0, \tag{1}$$

which is proportional to the applied magnetic field. Here f_0 is the Larmor frequency in MHz, γ_0 is the gyromagnetic ratio, a material-specific constant and B_0

is the strength of the magnetic field in T. In an external magnetic field in the z-direction, the spins are gradually aligned. With this alignment, both parallel and anti-parallel alignment is possible, whereby the parallel alignment is energetically preferred the greater the applied magnetic field. The fact that more spins are aligned parallel than antiparallel results in a small measurable longitudinal magnetization M_Z. If you now want to add energy to this system, this can be done with an electromagnetic wave that has the same frequency as the precession of the spins in the magnetic field, i.e. the Larmor frequency. This excitation causes the axes of rotation of the spins to tilt out of the z-direction again. With a precise high-frequency pulse with a specific power and duration, the axes of rotation can be tilted by a specific angle, e.g. 90° or 180°. With a 90° pulse, the magnetization M tilts from the z direction to the x-y plane and rotates there at the Larmor frequency, inducing an AC voltage with this frequency in the receiving coils. This voltage is the MR signal on which MRI imaging is based.

1.1.2 Relaxation and Bloch equations

Let us now look at how the spins fall back from the excited states to the stable initial state over time and how the transverse magnetization in the x-y plane slowly decays. This occurs through the spin-lattice interaction and the spin-spin interaction. These processes are also known as T_1 and T_2 relaxation and characterized by the corresponding tissue-dependent time constant T_1 and T_2. These processes are described by the Bloch equations:

$$\begin{aligned} \frac{d}{dt}M_x &= -\frac{M_x}{T_2} \\ \frac{d}{dt}M_y &= -\frac{M_y}{T_2} \\ \frac{d}{dt}M_z &= -\frac{M_z - M_0}{T_1} \end{aligned} \tag{2}$$

where M_0 denotes the equilibrium magnetization and M_i with $i = x, y, z$ the magnetization in the corresponding direction. The solution to these differential equations after a 90° are:

$$\begin{aligned} M_x &= M_0 \cdot e^{-\frac{t}{T_2}} \\ M_y &= M_0 \cdot e^{-\frac{t}{T_2}} \\ M_z &= M_0 \cdot \left(1 - e^{-\frac{t}{T_1}}\right) \end{aligned} \tag{3}$$

The relaxation mechanisms are described in the following:
By releasing energy to the environment, the spins realign themselves along the magnetic field, whereby the transverse magnetization decreases and the longitudinal magnetization M_Z increases. This process is therefore referred to as spin-lattice or longitudinal relaxation. The characteristic Spin-Lattice Relaxation Time T_1 (T_1) for this process depends on the magnetic field strength and the material-dependent internal movement of the molecules. For tissue, it is in the order of half a second to several seconds when a magnetic field of 1.5 T is applied.

However, in addition to the release of energy to the environment, a second effect also destroys the transverse magnetization. Immediately after excitation, all spins precess synchronously, i.e. have the same phase. Energy exchange between spins, caused by rapidly changing local magnetic field changes due to neighboring spins, causes the phase coherence to decay. As a result, the magnetic vectors partially cancel each other out instead of adding up and the transverse magnetization decreases. This is referred to as the spin-spin interaction. It has the time constant Spin-Spin Relaxation Time T_2 (T_2).
In addition to the influence of other spins, inhomogeneities in the external magnetic field also lead to dephasing of the spins. These inhomogeneities are caused by the device itself as well as by the object under investigation in the magnetic field, e.g. a human body. This additional dephasing of the spins has the time constant T_2^* , which is usually shorter than the spin-spin relaxation time T_2. The T_2^* effect in the measurements can be suppressed by special sequences.
The spin-lattice interaction and the spin-spin interaction occur simultaneously and independently of each other. However, since the T_2 time is much shorter than the T_1 time, the measurable transverse magnetization disappears long before the longitudinal magnetization has built up again.

1.1.3 Image Weighting

Three tissue parameters determine the intensity with which the tissue is displayed in an MRI image. Firstly, the proton density ρ, i.e. the number of excitable spins per unit volume. Then the T_1 time, i.e. the time until the spins have realigned themselves along the magnetic field after an excitation, and finally the T_2 time, i.e. the time until the signal decays after an excitation due to the spins running out of phase. Depending on which property is emphasized in an image, we speak of proton-weighted, T_1-weighted or T_2-weighted MR images. Examples of this are shown in Figure 1.

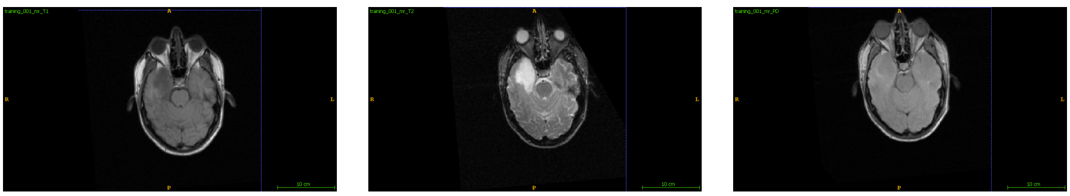

(a) (b) (c)

Figure 1: View of different modalities: (a) MR-T_1, (b) MR-T_2, (c) MR-PD.

How this works is now illustrated on the example of a spin-echo sequence which will be used in this experiment to get differently weighted images. The signal can be described by:

$$S \propto \rho \left(1 - e^{-\frac{TR}{T_1}}\right) e^{-\frac{TE}{T_2}} \tag{4}$$

where Repetition time (TR) and Echo time (TE) are the measurement parameters that can be set on the scanner and that are described in more detail in the following:

Several excitations of the same layer are necessary for spatial resolution. The time between these excitations is called TR. The longer you wait between two excitations, the more spins realign themselves along the magnetic field and can contribute to the signal again during the next excitation. A short TR means that, depending on the tissue, the spins can relax to different degrees before a new excitation occurs. As a result, different tissues deliver different intensities depending on their T_1. However, if TR is chosen to be very long, the spins in all tissues have time to relax again, which means that the image contrast is no longer dependent on the T_1 times of the tissues. Another time that determines the subsequent image contrast is the echo time (TE). This is the time between the excitation of the spins and their measurement. Switching the gradient coils on and off during spatial encoding creates inhomogeneities in the magnetic field, which amplifies the T_2 and T_2^* effects. With a very short TE compared to the T_2 times of the examined tissue ($TE \ll T_2$), the T_2 and T_2^* effects are barely measurable.
In concrete terms, this now means that for T_1 wheighted images TR needs to be chosen smaller than T_1 ($TR < T_1$) to obtain changes due to different T_1 values of different tissues. On the other hand, T_2 should not influence the measuremet, which means that acte should be chosen much smaller than T_2 of the tissue ($TE \ll T_2$). THe explanation in the previous paragraph was a qualitative description, but you can get the to the same setting by looking at equation 4. To obtain T_1 contrast, you want to have changes in the term $e^{-\frac{TR}{T_1}}$, while $e^{-\frac{TE}{T_2}}$ stays constant. This can be achieved by having $TR < T_1$ and therefore changes in the T_1 dependent term, while the T_2 dependent term stays approximately 1 when choosing $TE \ll T_2$.
With the same chain of reasoning you can derive that for T_2 weights imaging $TE < T_2$ to see differences in the T_2 dependent term, while $TR >> T_1$ ensures that the T_1 dependent term is approximately constant. The relationship how to choose TE and TR for differently weightes images is summarized in table 1.

	TR	TE
T_1-weighted	short	short
T_2-weighted	long	long
PD-weighted	long	short

Table 1: Image weighting through the choice of TR and TE.

1.1.4 Spatial Encoding

Now that we have discussed the generation of MR signals and the possibilities for different weighting of the imaging, we will take a look at spatial encoding in the last step. In order to be able to record a three-dimensional image, coordinates must be assigned to the MR signals. It must therefore be clear where in the body they originated from. To resolve the z-coordinate, an additional magnetic coil is used, which gives the magnetic field a gradient along the z-direction. As the Larmor frequency depends on the strength of the external magnetic field, the Larmor frequencies of the spins now differ along the z-direction. This means that only one layer in the body is excited with one frequency at a time. A strong gradient can

therefore be used to resolve fine layers, whereas a weak gradient can only resolve coarse layers.
For the resolution in the y-direction, the phase encoding, an additional gradient is applied in the y-direction. As a result, the Larmor frequency in the upper area of the MRI is slightly higher than in the lower area, which means that the spins in the upper area orbit faster than those in the lower area. This results in a phase shift of the spins. If the phase gradient is switched off again shortly afterwards, all spins move at the same speed again, but can be assigned to their y-position based on their phase.
With an additional gradient in the x-direction, the spins precess faster in the area with a stronger magnetic field than in the area with a weaker magnetic field. As a result, the MR signal no longer has just one frequency, but a frequency spectrum in which the height of the frequency provides information about the x-position. The frequencies can be analyzed using a Fourier transformation. However, one measurement is not sufficient to analyze the phase information; several measurements with different phase encodings must be carried out. In this way, the phase information can be decoded from the results using a further Fourier transformation. The time between these measurements is TR as already discussed.

1.2 CT-Imaging

Let us now turn our attention to another imaging system, the computer tomograph (CT). A CT consists of an X-ray tube, apertures, a table and a detector. A computing unit is also required to calculate an image from the detector data, as well as operating and display devices. The apertures ensure that the radiation generated in the X-ray tube passes through the body to be examined in a controlled manner. After leaving the body, it then hits a detector. The attenuation of the radiation in the tissue changes the energy that reaches the detector. By rotating the radiation source and detector around the patient, a large number of X-ray images are obtained from different perspectives. The contrasts in the subsequent images depend on the absorption capacity of the materials. This is measured on the Hounsfield scale, which indicates the relative difference in absorption of a material compared to water. Air has a value of -1000 HU, bones are in the order of 500-1500HU. A radon transformation is used to reconstruct the image data from the detector data, which only contain the sum of all absorptions.

2 Image registration

In order to be able to use image information from different modalities simultaneously, the data must be transferred to a common coordinate system. This is done as part of image registration. A data set is used as a reference and then a transformation is searched for that optimally overlays the second data set with the first. The reference can also be used to remove artifacts such as distortions from a data set. Since a combination of several imaging modalities is usually used for diagnostics, image registration is of great importance in medical image processing frameworks. This section covers the basics of image registration.

Let's first look at an image and a transformation that describes how the image should be transformed. The transformation can be used to calculate directly for each pixel where it is mapped to in the transformed image. Since an image is discretized due to its pixels, an interpolation method is also required to determine the intensity values at the pixel positions in the transformed image. This type of problem is also known as the forward problem. In contrast, in the registration problem, it is not the image and the transformation that are known, but the original image, referred to below as the reference image, and a transformed image, referred to below as the moving image. This image information is used to determine the transformation that best converts the moving image into the reference image. This type of problem is also known as an inverse problem and can generally only be solved using an optimization method. This has the general form:

$$\bar{\boldsymbol{u}} = \arg\min_{\boldsymbol{u}} (C(\boldsymbol{u})), \tag{5}$$

where u are the parameters of the image transformation to be optimized, $C(\boldsymbol{u})$ is the cost function that defines the similarity of the images and $\bar{u}$ are the parameters that describe the optimal transformation between the images.
The image registration algorithm thus consists of a metric that defines the similarity of two image datasets, a transformation that determines how the moving image may be transformed to increase the similarity between the moving image and the reference image, an interpolation method and an optimization algorithm that optimizes the metric value. In the following, we focus on similarity measures and transformations of medical image registration and their possible applications.

2.1 Similarity measures

Similarity measures can initially be divided into intensity-based and landmark-based measures. For intensity-based measures, the intensity values in the images are used to determine the similarity. For landmark-based approaches, significant positions in the images to be registered are first extracted and used to determine the similarity. In this practical course, we limit ourselves to intensity-based methods for which no landmark determination is necessary.
With intensity-based metrics, a distinction is made between metrics that directly compare the intensity values with each other and those that only use statistical information about the intensities of the images. The former can often only be used in the monomodal case, as the assumption that the same information is represented by the same intensities in the images is not guaranteed in the multimodal case.
The Sum of Squared Differences (SSD) can be used for a direct comparison of image intensities. Here, the square differences of the intensities between the reference image and the moving image are compared. In order to obtain a mean difference, the result is often normalized by the number of observed points N, whereby the Mean Squared Difference (MSD) is obtained:

$$C_{MSD} = \frac{1}{N} \sum_{i=1}^{N} \big(R(x_i) - M^{\mathcal{T}}(x_i)\big)^2 \tag{6}$$

$M^{\mathcal{T}}$ describes the moving image that was transformed into the coordinate system of the reference image using the transformation $\mathcal{T}$. x_i are the voxel coordinates of the reference image. In some cases, however, a direct comparison of the image intensities is not possible. One example is the combination of information from different medical imaging such as CT and MRI images. Although the same object is shown in both data sets, not all of the information can be seen in both data sets or has different intensities for the same information. In this case, Mutual Information (MI) has proven itself as a similarity measure. Here, only statistical information about the image intensities is used on the assumption that the intensities in the two images correlate.
Mutual information is defined as

$$S_{\mathrm{MI}}(I_A, I_B) = H(I_A) + H(I_B) - H(I_A, I_B), \tag{7}$$

where $H(I_A)$ and $H(I_B)$ are the respective marginal entropies of the image data I_A and I_B. $H(I_A, I_B)$ is the joint entropy of both image data. The marginal entropy is defined as

$$H(I) = -K \sum_a p_I(a) \cdot \log\left(p_I(a)\right), \tag{8}$$

where I is the image volume under consideration, $p_I(a)$ is the probability of intensity a in image I and K is a positive constant. The constant K depends on the selected base of the logarithm, but does not influence the subsequent optimization, as it only scales the cost function. The probability $p_I(a)$ can be determined by creating a histogram of the image intensities (see Figure 2(a)). To do this, the intensities occurring are entered in a histogram and this is then normalized so that the sum of all entries is 1. In many cases, it makes sense to assign an intensity interval to each entry in the histogram instead of a single intensity value. Accordingly, the joint entropy $H(I_A, I_B)$ is defined as

$$H(I_A, I_B) = -K \sum_{a,b} p_{I_A,I_B}(a, b) \cdot \log\left(p_{I_A,I_B}(a, b)\right), \tag{9}$$

where $p_{I_A,I_B}(a, b)$ is the probability for the intensity pair (a, b) in the image pair (I_A, I_B), which can be determined from a joint histogram of the image volumes. A 2D histogram can be created to calculate the probabilities for the occurrence of intensity pairs $p_{I_A,I_B}(a, b)$ (see Figure 2(b)). Here, the intensity in the images I_A and I_B is determined for each position x_i and the corresponding entry in the histogram is increased by 1. At the end, the 2D histogram is also normalized so that it reflects the probability of a pair occurring. There are even more complicated methods for constructing histograms in the literature, but only the simplest case is explained here for illustrative purposes.

2.2 Transformations

There are two types of transformations that can be considered in image registration. These are, on the one hand, transformations of image intensities and, on the other hand, geometric transformations. In this practical course, we will limit

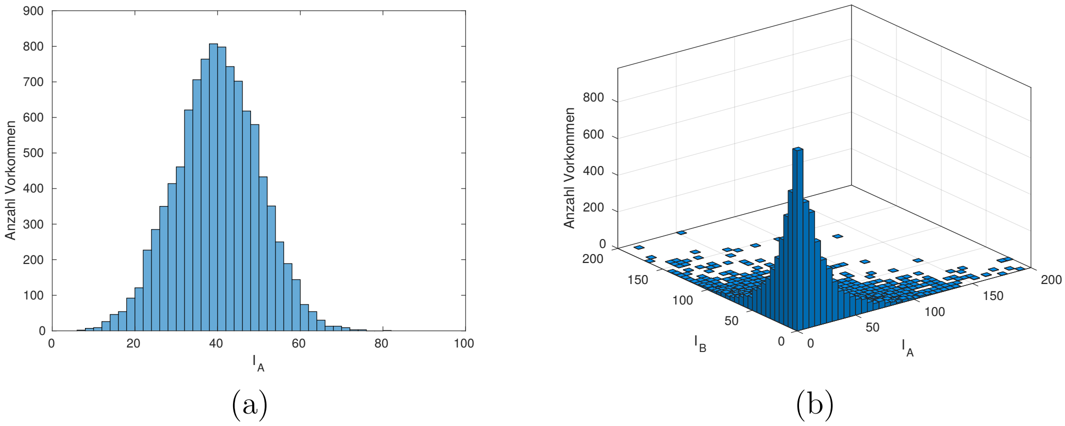

Figure 2: Representation of histograms to determine the probabilities $p_I(a)$ and $p_{I_A,I_B}(a,b)$: (a) 1D, (b) 2D.

ourselves to geometric transformations.
With geometric transformations, the basic distinction is between a global transformation of the image data, i.e. a transformation that is the same for all image points, and local transformations, which prescribe a separate transformation for each point in the image. The simplest model for global transformations is a translation. In this case, all points of an image are shifted by a constant vector. For the transformed coordinates (x', y', z') of a point (x, y, z), the following applies in the case of a translation with the vector (t_x, t_y, t_z) in three dimensions:

$$x' = x + t_x \tag{10}$$
$$y' = y + t_y \tag{11}$$
$$z' = z + t_z \tag{12}$$

If rotation is permitted in addition to translation, there are 6 degrees of freedom for the transformation (3 translation and 3 rotation degrees of freedom). This type of transformation is also referred to as rigid. A key property of the rigid transformation is that the distance between two points in the image does not change as a result of the transformation. If you also want to allow scaling and shearing of the images, the number of parameters to be optimized increases to 12. The individual components of the transformation are shown in Figure 3. This transformation is referred to as affine mapping. If the 3D coordinates are written as 4D vectors, these parameters can be represented in a 4×4 matrix:

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \\ p_{13} & p_{14} & p_{15} & p_{16} \end{pmatrix} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix} \tag{13}$$

Significantly more degrees of freedom are obtained if the transformation of an image is described by a transformation field instead of a single matrix. This is referred to as deformable registration. In this case, the position of a voxel in the image volume determines which transformation is applied to the voxel. Theoretically,

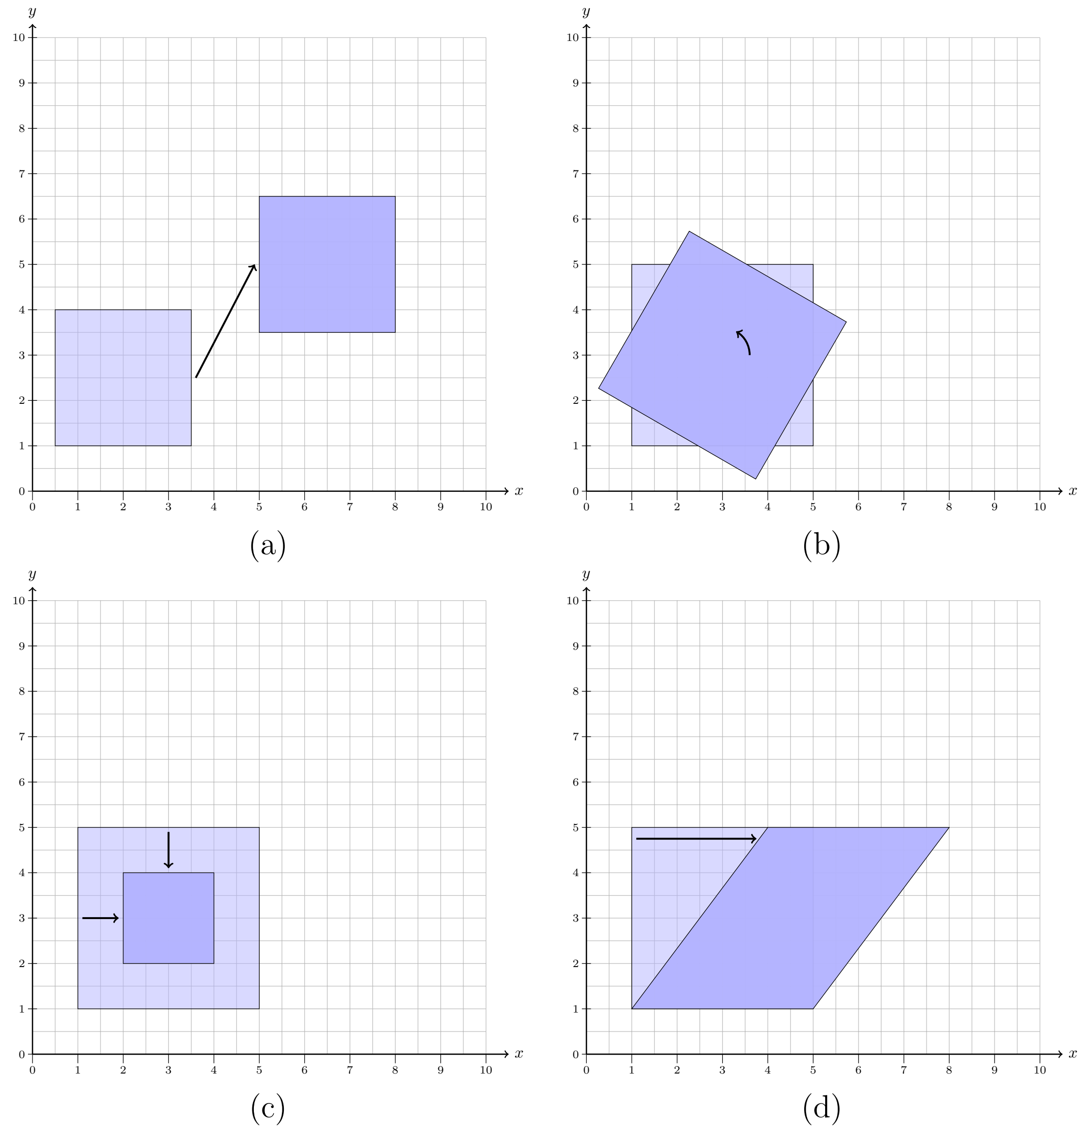

Figure 3: Illustrations of the possible transformations of an affine transformation. The light blue rectangle denotes the original image and the dark blue area the transformed image. (a) Translation, (b) rotation, (c) scaling, (d) shearing. In the case of rigid transformations, only combinations of cases (a) and (b) are permitted. [6]

you could define separate transformation parameters for each voxel. However, this quickly leads to thousands or millions of parameters. Since the transformation should be continuous between neighboring points, it is usually sufficient to define a much coarser grid of points on which the transformation is calculated and then interpolate the transformation parameters for points in between (see Figure 4). In this way, the number of parameters to be optimized can be significantly reduced. Another possibility to accelerate the deformable registration is to first determine the transformation parameters with a coarse grid with few points and then use these as a starting point for the optimization on a finer grid.

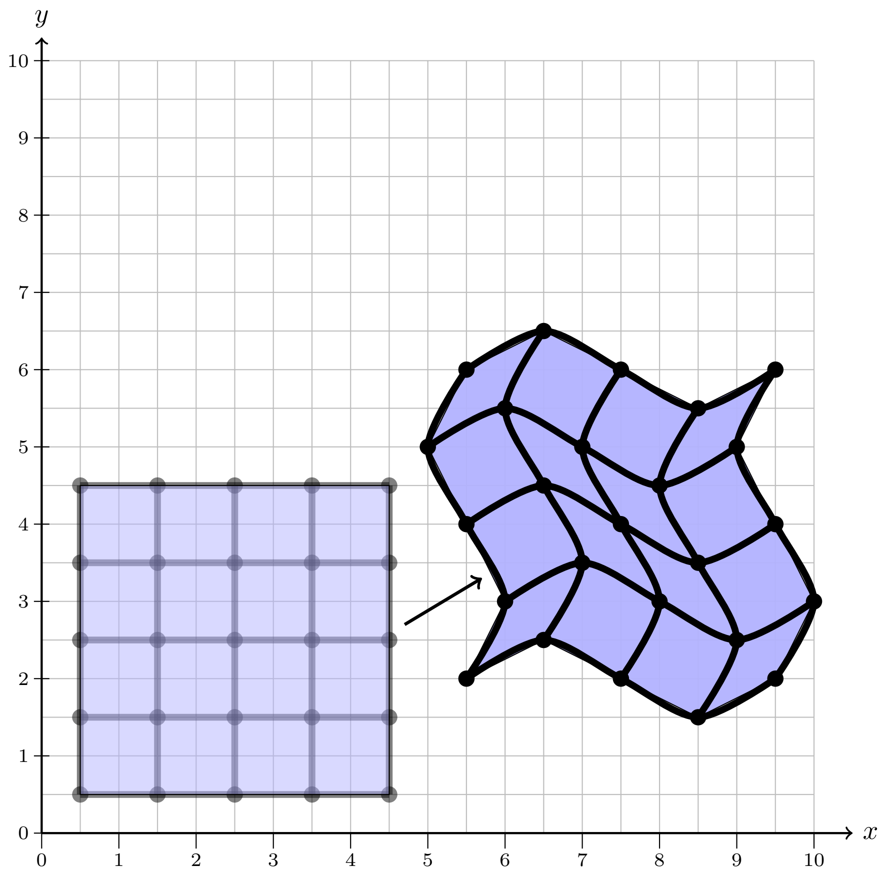

Figure 4: Representation of an elastic transformation. [6]

Execution

1 Experimental part I: Experiments with Table-top MRI

This part is planned to be performed on the first day[1]. In the end, images will be acquired that are used for image registration on the second day. Please perform the tasks marked as "Tasks" during the experiment and evaluate/discuss the results marked as "Evaluation and Discussion" afterwards at home.
In this part, all steps up to obtaining a final weighted 2D MR image are performed. As a first step, the scanner will be calibrated. Then T_1 and T_2 times are determined, which later serve as the basis for selecting appropriate TR and TE for the weighted images. Afterwards, the concept of spatial encoding will be introduced. Finally, T_1- and T_2-weighted 2D MR images are acquired, which can be used for registration in the next part.

1.1 Operating the Software

In this section, it is briefly described how the software is operated. The user interface can be seen in Figure 5. At the top left, you can select which experiment is to be carried out under "Experiment". Underneath you can start and stop the experiment. Below this, the respective parameters that can be varied in this experiment can be set in the field under "Parameters". On the right-hand side, in the upper section, you can see the current data that the device is currently recording. In the lower section on the right-hand side, you can see the pulse rate diagram of the current measurement. To save a measurement, click on "File" in the top left corner and then select "Export Data". This opens a window in which you can specify the storage location, the file name and the file type. If you want to save an image, you can select "Image" here. If you want to save data, select "Comma separated values" as the file type.

[1] Experimental part I is based on [1.]

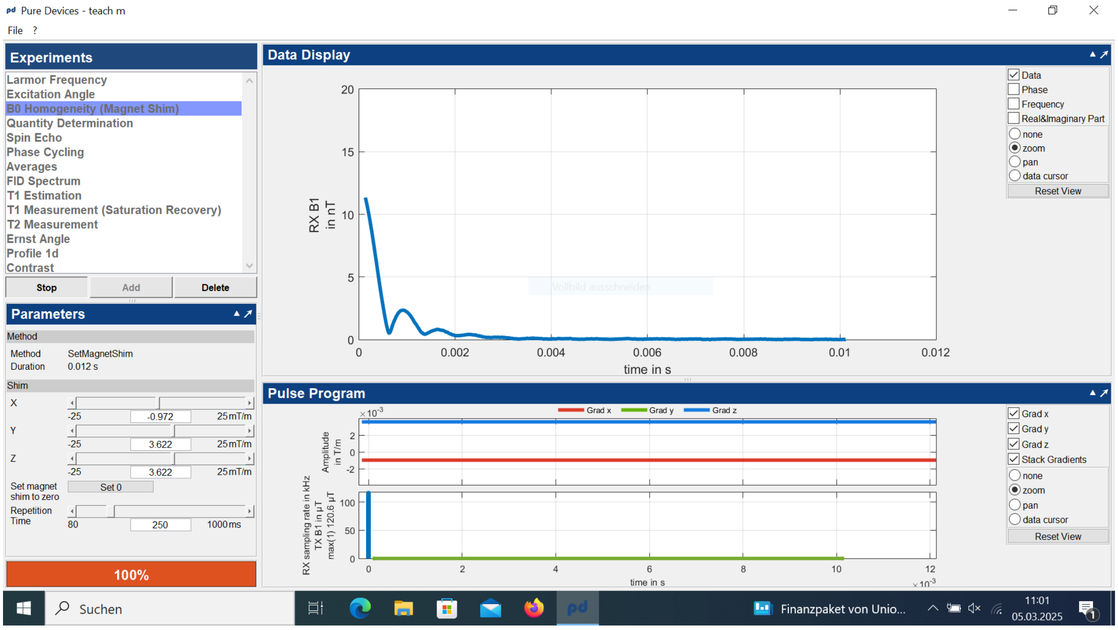

Figure 5: User interface of the software.

1.2 Calibration of the Scanner

First, the larmor frequency will be determined in order to perform resonance experiments. Then the setting for a flip angle of exactly 90° will be determined. Afterwards, shimming will be performed: There are inhomogeneities in the static magnetic field which are to be compensated by additional coils. The signal will then be investigated quantitatively, you will perform a first spin echo experiment and get a feeling how many averages are needed to get good results while the measurement time is kept as low as possible.

1.2.1 MR frequency

Objective of the experiment

The pulse sequence for this measurement contains a simple 90°-Radio Frequency (RF)-pulse for excitation of the spins. In this experiment you can freely adjust the frequency of the RF-pulse. When the frequency of the RF-pulse matches the larmor frequency the spins will get excited and rotate in the XY-plane. This induces electricity in the receiver coil and a free induction decay of the signal can be observed (Figure 6).

Preparation of the experiment

In this experiment the "10 mm Oil" sample is used. Place the test tube named "10 mm Oil" into the bore of the magnet.
In the upper left corner click on file and than options. Make sure all the boxes are not ticked except "Start program with maximized window" box.
Select the lesson MR frequency. The panel "Parameters" shows a slider for the

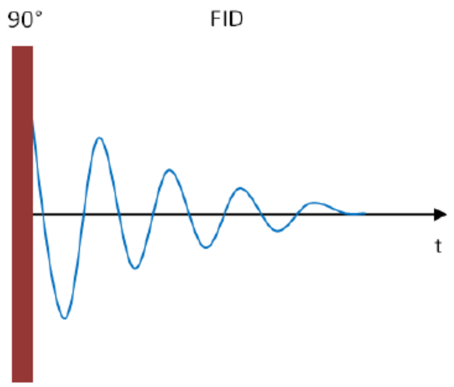

Figure 6: Pulse sequence for the experiment MR frequency.

frequency. Please make sure the real/imaginary part is chosen in the checkbox in the upper right corner. With the slider "Frequency" you can adjust the transmitter and receiver frequency of the tomography. Select "Start" for starting the measurement.

Tasks

1. Change the system frequency by varying the slider "Frequency". What are your observations? Describe them in a few sentences.
2. Now calibrate the frequency. An optimal calibration is achieved when the number of visible oscillations in the Free Induction Decay (FID) is as low as possible. Please save an image before and after the calibration and note the frequency ν.
3. How does the presence of iron influence the magnet? Adjust the frequency as mentioned in (2.). Take a piece of iron and bring it close to the magnet. What can be observed? Describe it in a few sentences and save the image for comparison.
4. What can be observed when changing the sample? Simply exchange the test tube with oil against water. Describe it in a few sentences and save the image for comparison.

Evaluation and Discussion

1. Calculate the field strength B_0 of the magnet. What is the exact field strength B_0 of the permanent magnet? Does this value match your expectations? ($\gamma_0 = 42.574 \frac{\mathrm{MHz}}{\mathrm{T}}$)
2. Compare the Images by discussing the differences between:
   - the FID before and after determining the Larmor frequency.
   - the FID without and with the influence of iron.
   - the FID with the water and oil samples.

1.2.2 MR Excitation Flipangle

Objective of the experiment

The knowledge of the 90° pulse length is one of the basic requirements for magnetic resonance experiments. A best signal can be obtained when the pulse tips the magnetization exactly into the XY-plane, resulting in a maximum of signal induced in the receiver coil. Further experiments are built on the exact knowledge of the pulse length. For determining the ideal 90° pulse length, a RF-pulse with adjustable duration is transmitted. Directly after excitation, the signal of the FID is acquired. The sequence is the same as in the experiment before; the diagram can be seen in Figure 6. The dependency between flip angle α and RF pulse duration τ is given by $\alpha = \gamma \int_0^\tau B_1(t)dt$ where γ is the gyromagnetic ratio and $B_1(t)$ the amplitude of the applied RF field.

Preparation of the experiment

In this experiment the "10 mm Oil" sample is used. Place the test tube named "10 mm Oil" into the bore of the magnet.
In the upper left corner click on file and than options. For the automatic adjustment of the frequency for the rest of the measurements tick the box "Automatically adjust frequency before measurements." and define an appropriate time interval.
Select the lesson "MR excitation flip angle". The panel "Parameters" shows a slider for the duration of the 90° pulse. With the slider "Duration" you can adjust the duration of the 90° RF-pulse. Select "Start" to start the measurement.

Tasks

1. Change the pulse duration of the 90° RF-pulse by varying the slider "Duration". What can you observe? Describe it in a few sentences.

2. What duration will give a 90° pulse and why? Please set the pulse length to an optimal 90° pulse. This setting will be used automatically for further experiments. Save two images: One with an optimal 90° pulse and one with a pulse that is not exactly 90°.

Evaluation and Discussion

1. Compare the Images by discussing the differences between:

   - the FID with an optimal 90° pulse and with a pulse that is not exactly 90°.

1.2.3 Quantity Determination

Objective of the experiment

The objective of this experiment is the investigation of the signal amplitude in relation of the substance quantity in the test tube. For this experiment a 90°

RF-pulse is used for excitation. The measurement is repeated continuously every value of TR. After each measurement the FID is displayed. In Figure 7 you can see the pulse sequence for this measurement.

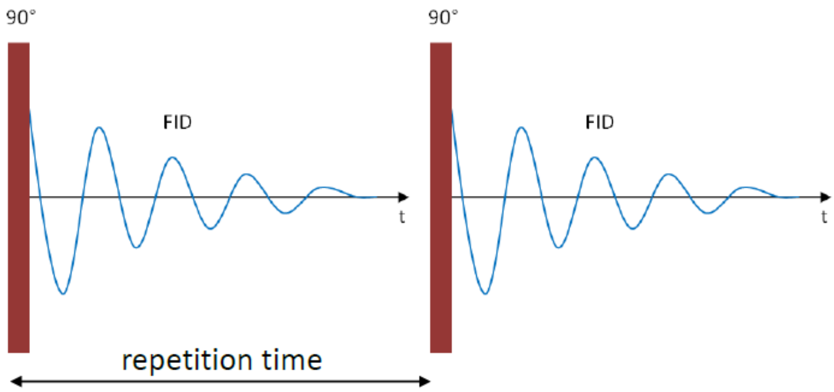

Figure 7: Pulse sequence for the experiment Quantity Determination.

Preparation of the Experiment

In this experiment, the "10 mm Oil" and "5 mm Oil" as well as the "10 mm Water" and "5 mm Water" samples are used.
First, place the test tube named "10 mm Oil" into the bore of the magnet.
Select the lesson "Quantity determination". The panel "Parameters" shows a slider for adjusting the TR of the experiment. With the slider "Repetitiontime" you can adjust the time between experiments. Set TR to 0.5 seconds.
Select "Start" to begin the measurement.

Tasks

1. Determine the amplitude of the FID directly after the 90° pulse. Please note that the duration of the 90° pulse is set to the value of the experiment 1.2.2 "MR excitation flip angle". What signal amplitude can you observe? Please save the image.

2. Repeat the measurement with the 5 mm oil test tube. What signal amplitude can you observe now? Please save the image.

3. Repeat the experiment with both water tubes. Please save both images and note both amplitudes.

4. Repeat the experiment with water and set TR to at least 5 seconds. Why should you set TR to at least 5 seconds when determining the amplitude of the water?

Evaluation and Discussion

1. Compare the signal amplitudes of the 5 mm and 10 mm tubes for all measurements performed with the same repetition time (TR=0.5) and medium (water or oil). What would you have expected, and why?

2. Compare the signal amplitudes between water and oil using the same tube diameter (5 or 10mm) and repetition time (TR=0.5). What would you have expected, and why?

3. Compare the signal amplitudes of the water measurements with different repetition times (TR) but the same tube diameter (5 or 10mm). What would you have expected, and why?

1.2.4 B_0 Imhomogeneities

Objective of the experiment

Objective of this experiment is the observation of FID duration depending on the magnet homogeneity. After a pulse the signal amplitude stays at a constant level as long as all dipoles rotate with the same frequency and without any influence of other disturbances. In practice, this is not the case since there are static as well as time-dependent disturbances of the magnetic field B_0. The simplest case of static disturbances is the inhomogeneity of the magnetic field B_0. This means that the field strength for the dipoles will vary with their location, leading to a signal cancellation. A fast decaying FID will be the result of an inhomogeneous magnet. The homogeneity of the static magnetic field can be improved by superimposing another electromagnetic field. This field is generated by additional coils called shim coils. After applying a 90° RF-pulse the decay of the FID can be extended by adjusting the field of the shim coils (see Figure 8). The best result is achieved when the decay cannot further be extended.

Preparation of the experiment

In this experiment the "10 mm Oil" sample is used. Place the test tube named "10 mm Oil" into the bore of the magnet.
Select the lesson "B_0 homogeneity (Magnet Shim)". The panel "Parameters" shows three sliders for the additional magnetic fields that will be generated by the corresponding shim coils. The unit of the shim field is mT/m. **Before you change anything, take a note of the currently set values**. Set all parameters to approximately zero mT/m before running this experiment. Select "Start" for starting the measurement.

Tasks

1. Safe a before image. Change the values for the shim field. What can you observe? Describe it in a few sentences and save an after image to show the differences.

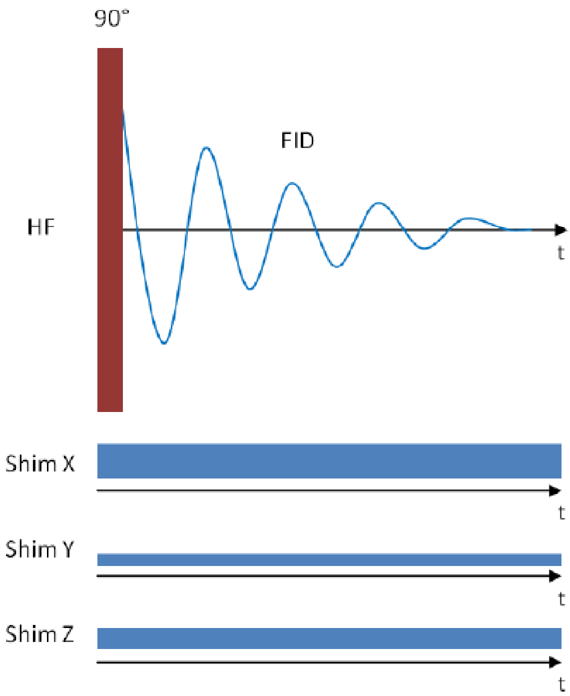

Figure 8: Pulse sequence for the experiment Quantity Determination.

2. Adjust the sliders to improve the homogeneity of the magnetic field. There may be several passes necessary to achieve a maximum homogeneity. To do so, follow the instructions below:

   (a) Set all shim values to approximately zero.

   (b) Start the procedure by adjusting slider X until the decay of the FID cannot be extended any further.

   (c) Now alter the settings for Y to extend the duration of the decay as well.

   (d) Do the same with Z.

   (e) Please start over with X and repeat this procedure until the decay cannot be improved.

   Set the values for Shim X, Shim Y and Shim Z to achieve the best possible shim. Note down the final values for Shim X, Shim Y and Shim Z. Please save an image of the final FID. These settings will automatically be used by the following experiments.

Evaluation and Discussion

1. Compare the Images by discussing the differences between:

   - the FID before and after the shimming.

1.2.5 Spin Echo

Objective of the experiment

The aim of the experiment is to generate a spin echo signal. The spins in the sample are affected by static as well as time dependent disturbances. A FID with a short decay time will be the result of these disturbances. Please refer to the results of the recent experiment B_0 homogeneity. If the disturbances are constant over time – like the inhomogeneity of the magnetic field - the phase relationship of the decaying signal and the spins will remain after the 90° RF-pulse. Now we can make use of this knowledge and rephase the signal by applying a 180° RF-pulse. The result is that the spins will "walk" into the opposite direction and generate a signal maximum. This maximum is called spin echo. The sequence for this experiment is shown in figure 9. At first the sample is excited by a 90° RF-pulse which will generate a FID. After some time a second RF-pulse 180° is applied to start rephasing the signal. The time when the spin echo signal will reach its maximum is called TE. This time depends on the position of the 180° pulse.

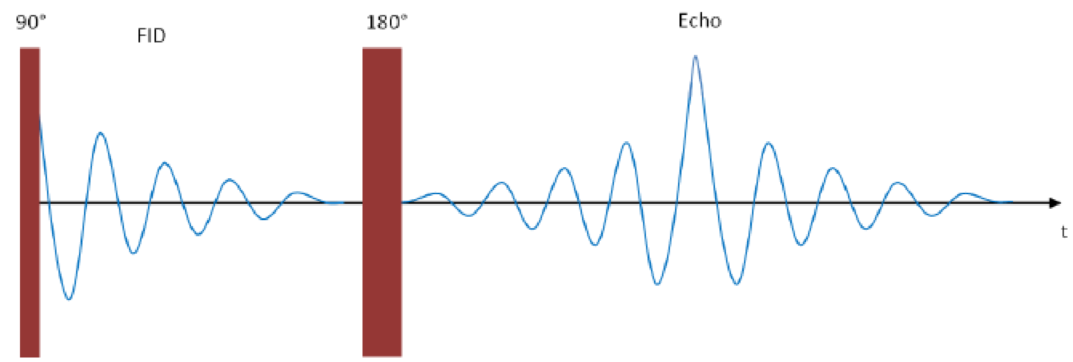

Figure 9: Pulse sequence for the experiment Spin Echo.

Preparation of the experiment

In this experiment the "10 mm Oil" sample is used. Place the test tube named "10 mm Oil" into the bore of the magnet. Select the lesson "Spin Echo". The panel "Parameters" shows a slider for the duration of the 90° pulse and one for the second pulse. Furthermore, there is an option to adjust the TE. With the slider "Duration of P1 (90°)" you can adjust the duration of the 90° pulse. The slider "Duration of P2 (180°)" is for adjusting the duration of the second RF-pulse (180°). As the name implies "Echotime" is for varying TE. Select "Start" for starting the measurement.

Tasks

1. Change the duration of the second pulse by varying the slider "Duration of P2 (180°)". What can you observe? Describe it in a few sentences and save

at least two images to show the differences. Adjust the control "Duration of P2 (180°)" to obtain an optimum spin echo.

2. Now vary the pulse duration of the 90° pulse. How does this affect the FID and the spin echo? Describe it in a few sentences and save at least two images to show the differences.

3. What can you observe when altering TE by varying the slider "Echotime"? (Note: spin echo amplitude) Describe it in a few sentences and save at least two images to show the differences.

The last setting for the 180° RF-pulse is automatically retained when running other experiments. It is advised to stop the measurement while the best values are chosen.

Evaluation and Discussion

1. Compare the Images by discussing the differences between:

   - the echo when changing the duration of the 180° pulse.
   - the echo when changing the duration of the 90° pulse.
   - the echo when changing TE.

1.3 T_1 and T_2 Measurements

In this section the T_1 and T_2 relaxation times of water and oil are determined. These times are needed afterwards to estimate the measurement parameters TR and TE for the differently weighted images.

1.3.1 T_1 Estimation

Objective of the experiment

Aim of this experiment is to estimate the spin-lattice relaxation time T_1 of oil and water. This can be done by comparing the signal amplitude of FIDs that are acquired with different TR (compare Figure 10).

Preparation of the experiment

In this experiment the "10 mm Water" sample and the "10 mm Oil" sample are used. Place the test tube named "10 mm Oil" into the bore of the magnet. Select the lesson "T_1 estimation". The panel "Parameters" shows a slider "Repetition-time" for TR of this experiment. With the slider "Repetitiontime" you can adjust the duration between consecutive measurements. Select "Start" for starting the measurement.

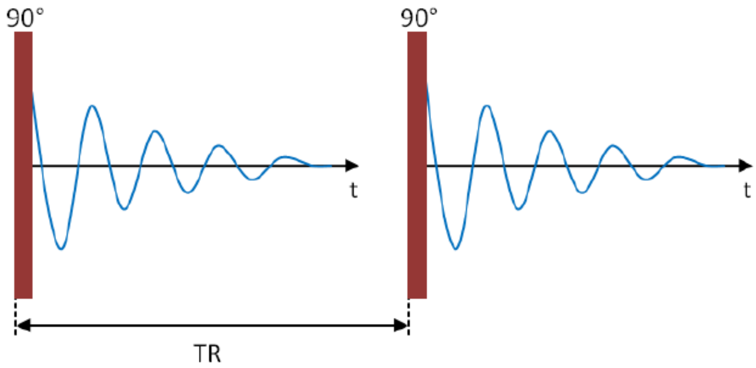

Figure 10: Pulse sequence for the experiment T_1 estimation.

Tasks

1. Change the time between both pulses by varying the slider "Repetitiontime". What can you observe? Describe in a few sentences.
2. Set TR to 15 seconds and note the value of the signal amplitude. Please save the image.
3. Reduce TR to a value where the signal amplitude has reached half of its maximum. It is recommended to proceed as follows: At first reduce TR time from 15 to 2 seconds. If the resulting FID amplitude is below half of the maximum TR has to be increased again. Repeat this procedure with smaller steps until you have achieved a proper value. Note the value of TR and save the image.
4. Calculate the spin-lattice relaxation time T_1 of oil using the value of TR.
5. Now repeat steps 2 and 3 with the "10 mm water" sample and calculate the relaxation time T_1 for water.

1.3.2 T_1 Determination

Objective of the experiment

In this experiment the spin-lattice relaxation time T_1 of various samples should be determined. The pulse sequence for this experiment is shown in Figure 11. After a Dwell time (DT) past the first 90° RF- pulse a second 90° pulse is applied. The second pulse is applied during the relaxation time of the sample. The maximum signal amplitude of the FID after the second excitation pulse will be automatically measured and processed by the computer. After the acquisition of the maximum of the second FID the measurement will pause for the time defined by TR. This will allow the sample to relax. This measurement will be repeated a predefined number of steps. With every step the time between the two 90° pulses is increased by the number of the current measurement multiplied with the predefined DT. The direct link between DT (duration between two pulses) and the signal intensity can be used to calculate the relaxation time T_1. For determination of the relaxation time T_1 this measurement can be applied to any sample. The relaxation time T_1 is a substance-specific value.

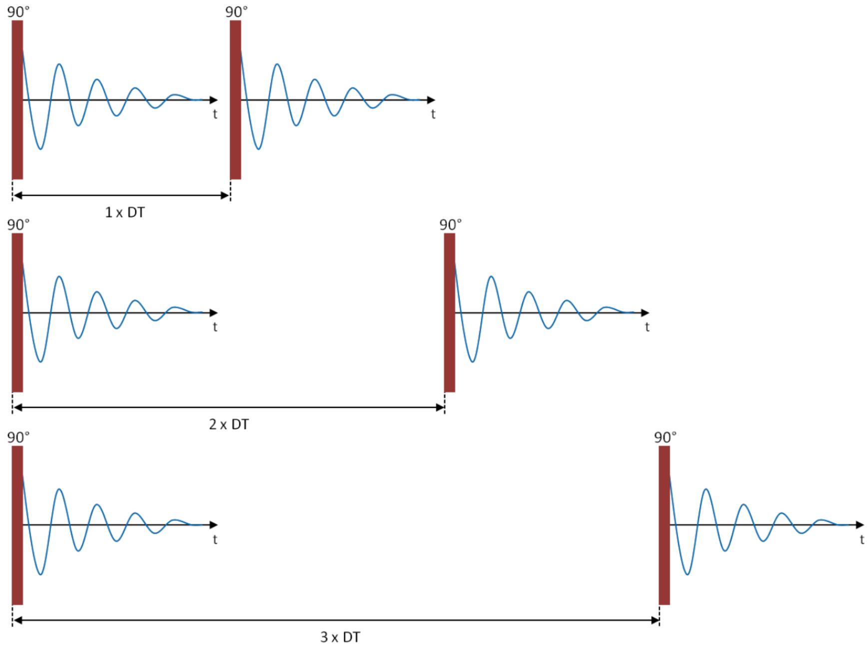

Figure 11: Pulse sequence for the experiment T_1 determination.

Preparation of the experiment

In this experiment the "10 mm Water" sample and the "10 mm Oil" sample are used. Place the test tube named "10 mm Oil" into the bore of the magnet. Select the lesson "T_1 Measurement (Saturation Recovery)". The panel "Parameters" shows a slider "Pause time" for TR of this experiment. Further there is a slider "Recovery time increment" for setting the value of DT. The control "Number of points" is for defining the number of measurements. With the slider "Pause time" you can adjust the duration of a single measurement. The slider "Recovery time increment" is for adjusting the value of DT – the time between the two 90° pulses. The slider "Number of points" varies the number of consecutive experiments.

Tasks

1. Set TR to the 3-fold of the T_1 value estimated in experiment T_1 estimation. Set the number of experiments to 24. Change the slider value for DT to an eigth of the estimated T_1 value. Now start the measurement.

2. Please click in the lower left field "Parameters" on "Fit T_1". A fit is performed and plotted as an overlay on the data and the corresponding T_1 value is shown in the plot. Please save the image **and the data**.

3. Repeat this experiment for the water sample (test tube 10mm Water) and save the image **and the data**.

Evaluation and Discussion

1. Perform a T_1 fit (cf. Equation 3) to the data (both water and oil) and discuss its quality, i.e. calculate the errors from the corresponding entry of the covariance matrix and use them to discuss how well the fit matches the data.

1.3.3 T_2 Determination

Objective of the experiment

The aim of this experiment is to measure the spin-spin relaxation time T_2 of different samples. This is can be done by comparing the signal amplitude of FIDs that are acquired with different TR (compare Figure 12). To do this a series of spin echoes also known as spin echo train is used. The pulse sequence for this experiment is shown in Figure 12. After the first spin echo a second 180° refocusing pulse is applied which will generate another spin echo at the time 2TE. Last one will show smaller amplitude than the previous spin echo. This is due to the relaxation time T_2. The maximum signal intensity of each echo is detected. After the data is acquired a curve fitting algorithm can be applied to calculate the relaxation time T_2 of the sample.

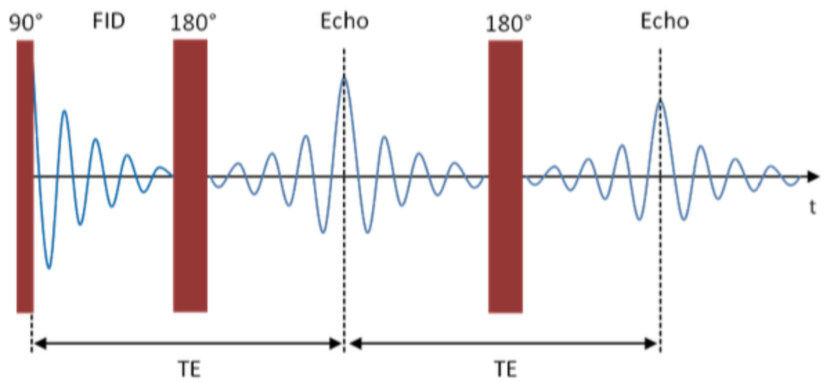

Figure 12: Pulse sequence for the experiment act2 determination.

Preparation of the experiment

In this experiment the "10 mm Water" sample and the "10 mm Oil" samples are used. Place the test tube named "10 mm Oil" into the bore of the magnet. Select the lesson "T_2 Measurement". The panel "Parameters" shows a slider "Number of echoes" and "Echotime". With the slider "Number of echoes" you can set the number of echoes that will be generated in this experiment. (Note: the total time for this experiment is the product of Number of echoes and the Echotime). The slider "Echotime" is for adjusting the time between two spin echoes in the echo train.

Tasks

1. Set the number of echoes to approximately 250 and TE to 2 ms. Now start the measurement.

2. Please click in the lower left field "Parameters" on "Fit T_2". A fit is performed and plotted as an overlay on the data and the corresponding T_2 value is shown in the plot. Please save the image **and the data**.

3. Repeat this experiment for the water sample (test tube 10mm Water) with 450 number of echos and TE 8 ms. Save the image **and the data**.

Evaluation and Discussion

1. Perform a T_2 fit (cf. Equation 3) to the data (both water and oil) and discuss its quality, i.e. calculate the errors from the corresponding entry of the covariance matrix and use them to discuss how well the fit matches the data.

1.4 Frequency Encoding

1.4.1 Profile

Objective of the experiment

Aim of this experiment is to measure a profile of the sample. The spatial encoding of the signal is done by applying additional magnetic fields which are generated by the gradient coils. The resulting magnetic field is a superposition of the main field B_0 and the field generated by the gradient coils. This implies a change in frequency for each position in the magnet. Applying a frequency analysis (Fast-Fourier-Transformation - FFT) to the acquired data will allow assigning the signal to its location in the sample.

Preparation of the experiment

In this experiment the "5 mm Water" and the "5 mm Oil" sample are used. Put both test tubes into the bore of the magnet so that they are in line and parallel with the rear edge of the magnet (Z-axis). Select the lesson "Profile 1d". The panel "Parameters" shows a slider for TR and three sliders for adjusting the gradient field. With the slider "Repetition time" you can adjust the time between experiments. Changing the gradient field can be done by varying the slider for "Gradient X", "Gradient Y" and "Gradient Z". Set TR to 1 second and check that all gradient controls are set to 0. Select "Start" for starting the measurement.

Tasks

1. Change the gradients by varying the slider "Gradient X", "Gradient Y" and "Gradient Z" in turns. What can you observe? Describe it in a few sentences.

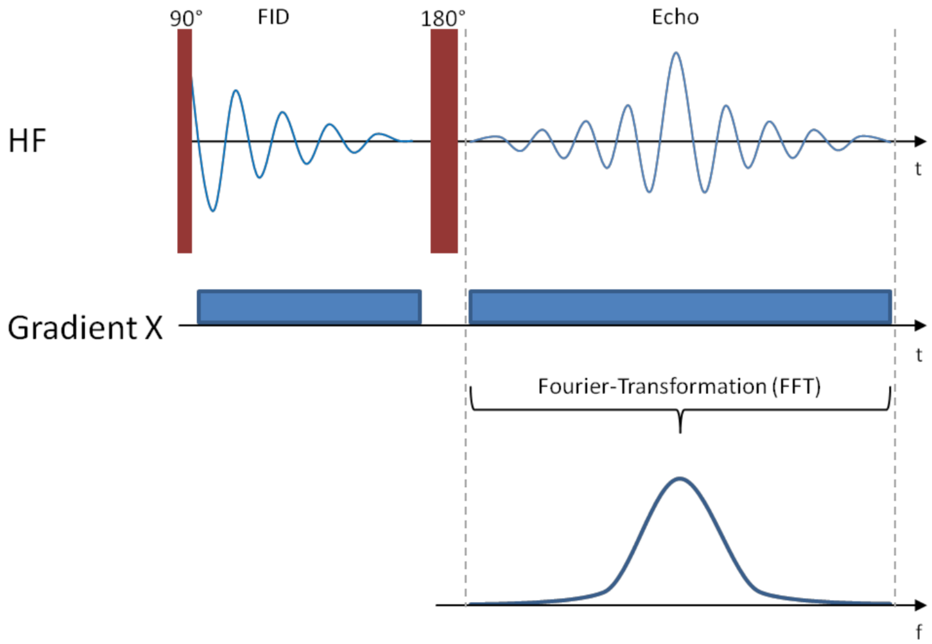

Figure 13: Illustration of the pulse sequence for the experiment Profile 1d and principle of frequency encoding.

2. Adjust the slider until you see the profile of the two test tubes. Rotate the sample by 90°. What can you observe? Describe in a few sentences and save two images, one before adjusting the gradient and one after adjusting the gradient.

3. Put the test tube back into the previous position and note the values of the sliders.

4. What is the reason for the gradient between the 90° and the 180°?

Evaluation and Discussion

1. Compare the Images by discussing the differences between:

    - the signal when adjusting the gradient.

1.4.2 T_1 Profile

Objective of the experiment

In this experiment the aim is to acquire a T_1 weighted profile of a sample. This experiment combines the spatial information from the profile of the sample (compare with previous experiment 1.4.1) with the relaxation time T_1 of the sample. To obtain the information a sequence of measurements has to be performed. A single measurement contains a 90° RF-pulse DT and the acquisition of the profile like in experiment 1.4.1. The DT will be increased after each single measurement.

This will lead to a T_1-weighted profile of the sample (compare this with section 1.3 "act1 and act2 Measurements").

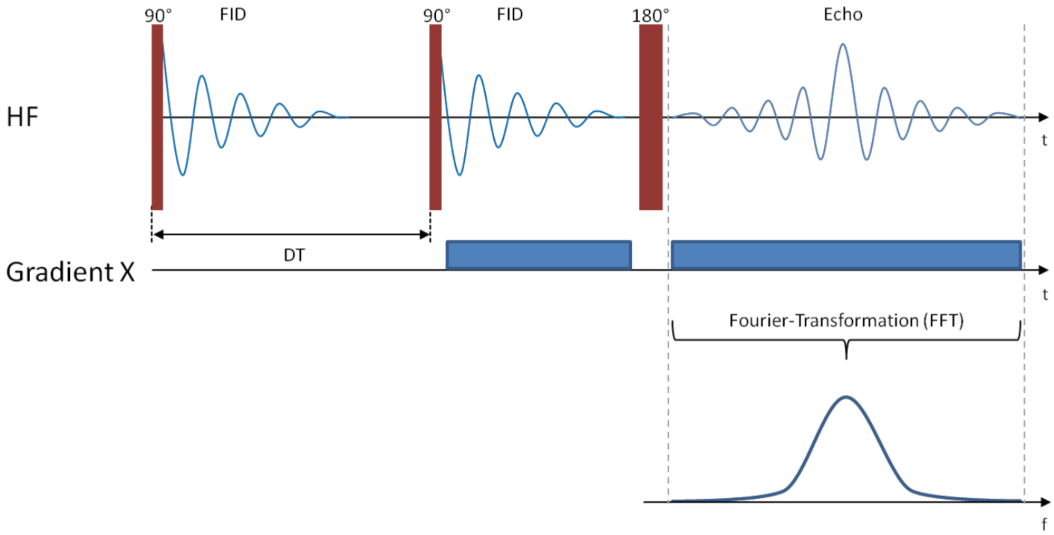

Figure 14: Pulse sequence for the experiment act1 profile.

Preparation of the experiment

In this experiment the "5 mm Water" and the "5 mm Oil" samples are used. Put both test tubes into the bore of the magnet so that they are in on line and parallel with the rear edge of the magnet. Select the lesson "T_1 profile". The panel Parameters shows a slider "Pause time" for the TR and three sliders for adjusting the gradient field. Further there is a track bar for the number of points and the time step that will be used for DT. These parameters can be accessed in this experiment. With the slider "Pause time" you can adjust the time between experiments. Changing the gradient field can be done by varying the slider for "Gradient X", "Gradient Y" and "Gradient Z". With the slider "Number of profiles", the number of measurements can be adjusted. DT equals the value of "Recovery time increment".

Tasks

1. Set TR to 1 second. Set the sliders for the gradient to the value where you obtained the profile in experiment 1.4.1 "Profile 1d".
2. Choose a suitable value for DT. Typically one chooses a tenth of the estimated T_1 value for water. Set the number of points to a value where DT of the last measurement is greater than the estimated value for T_1.
3. Now start the measurement. What can be observed? Save the resulting plot.

Evaluation and Discussion

1. Which peak belongs to the water sample and which one to the oil sample? Why?

1.4.3 T_2 Profile

Objective of the experiment

In this experiment we want to acquire a T_2 weighted profile of a sample. This experiment combines the spatial information from the profile of the sample (compare with previous experiment 1.4.1) with the relaxation time T_2 of the sample. To obtain the information a sequence of measurements has to be performed. A single measurement contains a 90° RF-pulse, DT and the acquisition of the profile like in experiment 1.4.1. DT will be added to TE after each single measurement. This will lead to a T_2-weighted profile of the sample (compare this with section 2 – relaxation times

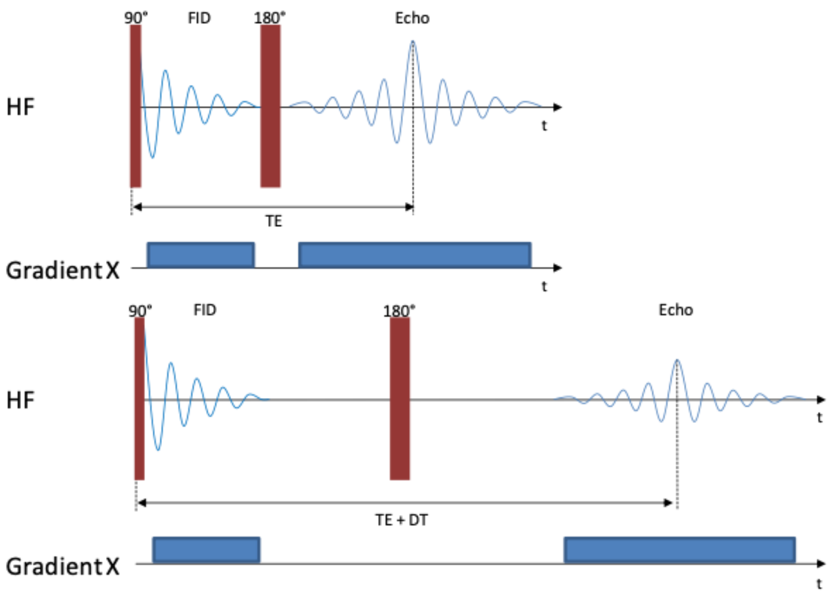

Figure 15: Pulse sequence for the experiment act2 profile.

Preparation of the experiment

In this experiment the "5 mm Water" and the "5 mm Oil" samples are used. Put both test tubes into the bore of the magnet so that they are in line and parallel with the rear edge of the magnet. Select the lesson "T_2 profile". The panel "Parameters" shows a slider "Pause time" for TR and three sliders for adjusting the gradient field. Further there is a slider for the number of points and the time step (DT). With the slider "Pause time" you can adjust the time between experiments. Changing the gradient field can be done by varying the slider for "Gradient X", "Gradient Y" and "Gradient Z". With the track bar "Number of profiles", the number of echoes to acquire can be adjusted. DT equals the value of "Echo time increment".

Tasks

1. Set TR to 2 second. Set the sliders for the gradient to the value where you obtained the profile in experiment 1.4.1 Profile.
2. Choose a suitable value for DT. Typically one chooses a tenth of the estimated T_2 value for oil. Set the number of points to a value where DT of the last measurement is greater than the estimated value for T_2.
3. Now start the measurement. What can be observed? Save the resulting plot.

Evaluation and Discussion

1. Which peak belongs to the water sample and which one to the oil sample? Why?

1.5 Acquisition of Weighted 2D Images

In this section, the second dimension is now also encoded using phase coding to create weighted 2D images.

1.5.1 Spin Echo 2D

Objective of the experiment

In this experiment the aim is to acquire a 2D magnetic resonance image of a sample. The acquisition is a two-dimensional cross section of the sample based on the spin echo method. In this measurement the parameters can be set equivalent to clinical magnetic resonance tomography devices.

Preparation of the experiment

In this experiment, both 5mm test tubes containing water and oil are used. At first, place both 5 mm test tubes into the bore of the magnet so that they are in on line and parallel with the rear edge of the magnet. Select the lesson "Spin Echo 2D". The panel "Parameters" shows the following controls:

- "Datapoints (Read)": The number of data samples acquired per phase step in read direction.
- "Datapoint (Phase)": The number of phase-encoding steps in the phase direction.
- "Slice orientation": The plane of the cross section through the sample.
- "Image Size (Read)": The gradient strength of the read gradient. The greater this value the bigger the size of the sample in the read direction.

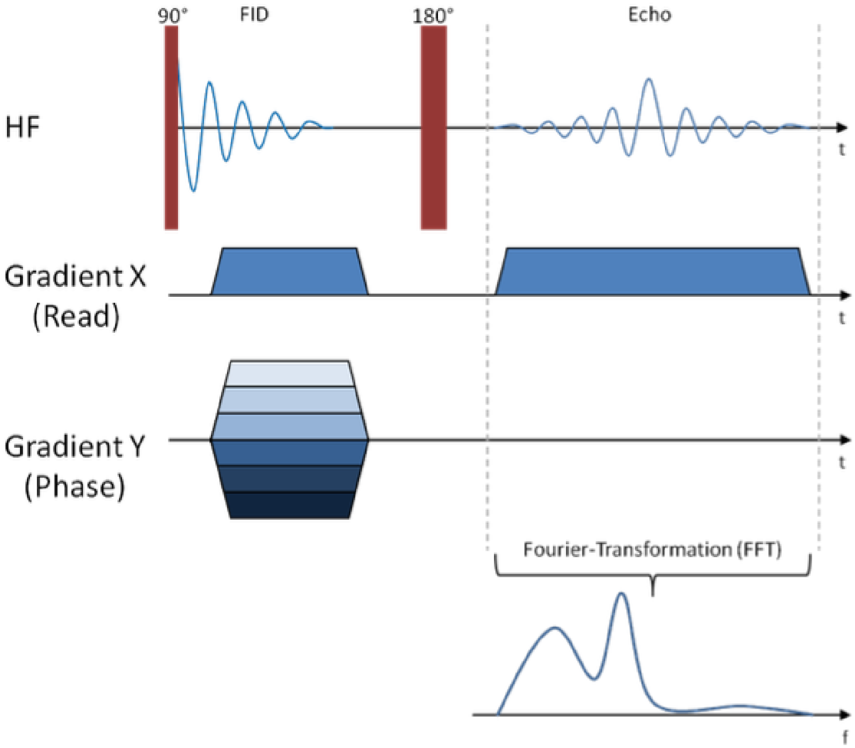

Figure 16: Pulse sequence for the experiment Spin Echo 2D.

- "Image Size (Phase)": The strength of the gradient of the phase encoding gradient. The greater this value the bigger the size of the sample in the encoding phase direction.
- "RepetitionTime": The repetition time TR for each phase encoding step.
- "Averages": The number of averages.
- "Echotime": The echo time TE.

Tasks

1. For the first magnetic resonance image, choose the parameter as follows:
   - "Datapoints (Read)": 32
   - "Datapoints (Phase)": 32
   - "Slice orientation": X-Z-slice
2. The resolution is kept low in this part so that different settings can be tried and the measurement time is still reasonable. Later the resolution will be increased to take the final images. Estimate the gradient strength for the read and phase gradient. Compare with experiment 1.4 "Profile 1d". Further set appropriate values for an T_1 weighted image for TR and TE.

3. Repeat the measurement until you obtain an image of reasonable quality. Which of the tubes in the image belongs to water and which to oil? Explain in a few sentences and save the acquired image as well as the image data.

4. Now change the values for TR and TE to perform a T_2 weighted image and repeat the measurement. Repeat the measurement until an image of reasonable quality is obtained. Which of the tubes in the image belongs to water and which to oil? Explain in a few sentences and save the acquired image.

5. Now the final images will be obtained. For this, increase the number of datapoints to 64 and the phasesteps to 64. Repeat the measurements with the settings of the T_1 and T_2 weighted images. Please save the images **and the data** (they are needed on the next day for image registration).

2 Experimental Part II: Registration in Monomodal and Multimodal Image Data

In this experiment, we examine two similarity metrics—Mean Squares Difference (MSD) and Mattes Mutual Information (MI)—using 2D image data from Task 5 of Section 1.5 acquired with the table-top MRI system. We analyze how these metrics vary with different rigid transformations by comparing fixed images (the original data) with moving images (rigidly transformed versions). We first consider the monomodal case (images with identical contrast) and then extend the analysis to the multimodal case (images with different contrasts). The registration results are evaluated by comparing the similarity and transformation parameters with known ground-truth values.
On the desktop you can find the "FP" folder. If it does not already exist create a duplicate of the "FP95_base_new" folder and rename it to "FP_day_month_year". Open "Visual Studio Code". On the left sidebar, click the "Explorer" icon (it looks like two overlapping files). Then click the "Open Folder" button and navigate to the folder you just created. Here you will find all the files necessary for the experimental part II and III.
This experiment is implemented in Python using the ITK library. To simplify the task, the file "FP_Aufgabe2_functions.py" contains four core functions—please do not modify this file. An example script, "FP_Aufgabe2_example.py", is provided to help you understand how to use these functions. The empty file "FP_Aufgabe2.py" is intended for your own code.
Here is a description of the core functions mentioned previously:

```
load_and_save_data(inputpath, outputpath)
```

This function loads CSV data exported from the table-top MRI system and saves it as `outputpath.dcm` and `outputpath.png`.

```
transform_image(inputpath, outputpath, shift_x, shift_y, rotation_angle)
```

This function reads a DICOM image and rigidly transforms it using the specified translation values (`shift_x` and `shift_y`) and `rotation_angle`. The result is saved as `outputpath.png` and `outputpath.dcm`.

```
compare_image_similarity(fixed_image_path, moving_image_path)
```

This function reads the fixed and moving DICOM images and prints the similarity values computed using the MSD and MI metrics.

```
perform_registration(fixed_image_path, moving_image_path,
    registered_image_path, metric_template)
```

This function registers the moving DICOM image to the fixed DICOM image using the chosen metric template (`itk.MeanSquaresImageToImageMetricv4` or `itk.MattesMutualInformationImageToImageMetricv4`). The registered image is saved as `registered_image_path.png` and `registered_image_path.dcm`, and the final transformation parameters along with the similarity value are printed.

2.1 Measuring Similarity

In the monomodal scenario, MRI images acquired with the same contrast are compared. Comparing an image to itself should yield the highest similarity. In contrast, similarity should decrease when the image is compared to a shifted or rotated version of itself. Note that the highest similarity does not necessarily correspond to the highest metric value.
This analysis is then extended to the multimodal case, using MRI images acquired with different contrasts. The T_1-weighted image is used as the fixed image, while the T_2-weighted image is transformed. The goal is to investigate how each similarity metric behaves in this context.

Tasks

1. Implement code to systematically vary rigid transformation parameters (translation and rotation) in the monomodal case. Use the `transform_image` function.
2. Add the calculation of the corresponding MSD and MI values to your code using the `compare_image_similarity` function.
3. Save the results to a file and visualize the metric values as a function of the transformation parameters.
4. Repeat the above steps for the multimodal case. Vary the transformation parameters and record the corresponding metric values.
5. Save and plot the results to analyze the metric behavior in the multimodal scenario.

2.2 Registering Images

In the monomodal scenario, the acquired image is used as the fixed image, while the transformed images from the subsection 2.1 serve as moving images.
In the multimodal scenario, the T_1-weighted image (fixed) and the T_2-weighted image (moving) are used. The registration is evaluated using both similarity metrics in each case.
The quality of the registration can be analyzed qualitatively by examining the registered images or quantitatively by comparing the similarity between the registered image and the fixed image. In the monomodal case, the similarity value should match the value of the image to itself for the same region. In the multimodal case, it should equal the similarity value between the original (non-transformed) T_1-weighted and T_2-weighted images for that region.

Tasks

Save the transformation parameters obtained from the registration and compare them with the known ground truth values to assess the accuracy of the registration. Additionally, use the metric value as a supplementary indicator of correctness. Both the transformation parameters and the metric value can serve as a basis for your discussion. Analyze the quality of all the performed registrations both qualitatively and quantitatively. Do not forget to take notes.

1. Select one of the acquired images and three corresponding rotated images (monomodal), between 0 and 90 degrees. Perform the registration using both the MSD and MI metrics.

2. Identify one rotation angle where both metrics yield good results, and another where only one metric performs reliably.

3. Analyze the quality of the registrations performed both qualitatively and quantitatively.

4. Now perform two separate registrations using the acquired image and two translated versions: one shifted by 20 pixels in the x-direction and the other by 20 pixels in the y-direction.

5. Select three multimodal image pairs and register them using both MSD and MI.

2.3 Evaluation

1. Include the similarity metric plots in your report. Discuss the results and explain the conditions under which each metric performs best in both monomodal and multimodal registration scenarios.

2. Analyze the limitations of each similarity metric. Describe the challenges or constraints associated with each metric and how they may affect registration performance.

3. Explain why the similarity value obtained using the Mutual Information (MI) metric is negative.

4. Include the registration results in your report. Evaluate the limitations of each metric and identify which types of transformations lead to inaccuracies or registration failures.

5. Discuss potential reasons for registration failures or inaccuracies.

3 Experimental part III: Deformable registration

The last part of the experiment demonstrates how a deformable registration works. For this purpose, an image of a white rectangle on a black background and an image of a white circle on a black background are generated. The aim of the registration is to deform the image of the rectangle so that the deformed image represents a circle.
We also use the ITK library in Pyhton for this part of the task. Instead of the rigid transformation, a spline-based transformation is selected for the deformable registration in this part of the experiment. Since significantly more parameters have to be optimized here than with a rigid transformation, it also makes sense to choose a different optimizer that can quickly optimize a large number of parameters. The LBFGS optimizer is used for this task. Because this is a monomodal problem, a metric can be used that directly compares the image intensities.
The main property to be considered here is the influence of the number of grid points of the transformation grid on the registration result, as well as the runtime of the registration. For this purpose, the number of grid points is to be varied and the results documented. The transformed images are saved directly as an image file.
Once a transformation field has been determined that satisfactorily transforms the rectangle in the moving image into a circle, this transformation field will be examined. In order to visualize the effect of the field in the individual areas of the image, the previously calculated transformation field is applied to an image with a chessboard pattern. Discuss the result. If you still have time and want to experiment with deformable registrations, you can also try out other geometric shapes as input data for the registration.
You will be provided with a code template "FP_Aufgabe3.py" with 5 functions, which will be explained in the following. Please do not change these functions:

```
create_cirlce_image(size, radius)
```

This function generates a rectangular black image with the size specified by size as a tuple with a circle in the middle of that rectangle with the specified radius. Make sure that the radius is smaller than half of the size.

```
create_plane_image(size, square_size)
```

This function generates a rectangular black image with the size specified by size as a tuple with a white square in the middle with the size specified by square_size. Make sure that square_size is smaller than the size.

```
create_chess_image(size, square_size_chess)
```

This function generates an image of the size specified by size, that has the shape of a chessboard, i.e. there is an alternating pattern of white and black squares. The size of the squares can be set by the input argument square_size_chess. Please make sure that size is a multiple of square_size_chess.

```
register_images(fixed_image, moving_image, number_gridnotes)
```

This function registered two images. You need to specify fixed_image and moving_image which both can be the output from the previous 3 functions. The input number_gridnotes defines the number of grid notes along each dimension of the image domain. This means, if you set this parameter to e.g. 30, the final grid will have 30x30 gridnotes. The output of this function is an image transformation, which can than be used to apply this transformation to another image with the "apply_registration_transform" function.

```
apply_registration_transform(fixed_image, moving_image, transformation)
```

This function applies the previously found registration between fixed_image and moving_image to the moving_image. The output is an image.

```
save_image(image, name)
```

This function can be used to save an image. The name and folder can be specified with the variable "name".

3.1 Tasks and Evaluation

1. Start with an image size of (100,100).
2. Choose an appropriate radius, size of the square, and size of the chess pattern.
3. Vary the number of gridnotes that define the transformation between 5 and 100 in steps of 5.
4. Document the time needed for the registration, the registration result (i.e. the images for a visual and qualitative discussion), and apply the determined transformation fields to the chessboard image.
5. Save the circular image, the image of the rectangle, the registered image, and the chessboard image for every setting of the number of gridnotes.

3.2 Discussion

1. Plot the time needed for registration as a function of the gridnotes.

2. Discuss your results and state which number of gridnotes provides the best compromise between the registration result and the registration time.

Question Catalog

- Where is image registration used in medicine?
- What types of transformation are used in image registration?
- Why does it make a difference for the choice of metric whether you have a monomodal or a multimodal problem?
- What is Mutual Information?
- What does the Sum of Squared Differences (SSD) metric describe?
- What physical phenomenon underlies the generation of the MRI signal?
- What is the Larmor frequency?
- What are T_1 and T_2 describing?
- What are the Bloch equations describing?
- How does image weighting in MRI work?
- How does spatial encoding in MRI work?

References and Literature

1. Experiments Module: MR basics, Pure Devices GmbH, December 2012
2. MRI questions website (https://mriquestions.com/index.html)
3. Weishaupt, D., Köchli, V. D., Marincek, B., Wie funktioniert MRI?, ISBN: 978-3-642-41616-3
4. D. L. G. Hill, P. G. Batchelor, M. Holden, D. J. Hawkes, Medical image registration, Physics in Medicine and Biology 46(3) (2001) R1-R45
5. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, P. Suetens, Multimodality image registration by maximization of mutual information, IEEE Transactions on Medical Imaging 16(2) (1997) 187-198
6. J. P. W. Pluim, J. B. A. Maintz, M. a. Viergever, Mutualinformation-based registration of medical images: a survey, IEEE Transactions on Medical Imaging 22(8) (2003) 986-1004
7. ITK - Segmentation & Registration Toolkit (https://itk.org)